NETARENA: DYNAMICALLY GENERATED LLM BENCH-MARKS FOR NETWORK APPLICATIONS

Anonymous authors

000

001

002003004

010 011

012 013

014

016

018

019

021

025

026

027 028 029

031

032

033

034

035

037

040

041

042

043

044

046

047

051

052

Paper under double-blind review

ABSTRACT

As large language models (LLMs) expand into high-stakes domains like network system operations, evaluating their real-world reliability becomes increasingly critical. However, existing benchmarks risk contamination due to static design, show high statistical variance from limited dataset size, and fail to reflect the complexity of production environments. We introduce NETARENA, a dynamic benchmark generation framework for network applications. NETARENA features a novel abstraction and unified interface that generalizes across applications, effectively addressing the challenges of dynamic benchmarking posed by the diversity of network tasks. At runtime, users can generate unlimited queries on demand. NETARENA integrates with network emulators to provide execution-time feedback on correctness, safety, and latency. We demonstrate NETARENA on three representative applications and find that (1) it significantly improve statistical reliability among LLM agents (confidence interval overlap reduced from 85% to 0), (2) agents achieve only 13–38% average performance (as low as 3%) for large-scale, realistic queries, (3) it reveals finer-grained behaviors missed by static, correctness-only benchmarks. NETARENA also enables use cases such as SFT and RL fine-tuning on network system tasks. Code is available anonymously at https://anonymous. 4open.science/r/netarena iclr2026-BE94/README.md.

1 Introduction

While LLMs (OpenAI, 2024; Google, 2024a; Meta, 2024; Yang et al., 2024) have rapidly advanced agent capabilities across general tasks (OpenAI, 2025; Sager et al., 2025), many existing benchmarks focus on simplified settings that do not fully capture the demands of real-world deployments. Network and system applications offer a compelling alternative: they are high-stakes and require LLMs to reason under constraints like partial observability and operational risk. From data center capacity planning (Mani et al., 2023) to root cause analysis (Chen et al., 2024) and policy synthesis (Wang et al., 2024a; Sharma & Yegneswaran, 2023), these tasks require not just correctness, but robustness and efficiency, making them an ideal stress test for LLM agents.

Despite the growing interest and advancements, rigorously evaluating LLM agents in network and systems remains an open challenge. These tasks often involve large-scale infrastructure with complex domain-specific logic, but current benchmarks rely on *manually curated queries and ground truths* by domain experts. This labor-intensive process has resulted in fewer than 300 queries in recent benchmarks (Mani et al., 2023; Chen et al., 2025) even after months of effort. Such small, static benchmarks introduce critical limitations: they are prone to statistical bias, vulnerable to data contamination (Zhu et al., 2023), and raising concerns about generalizability. For instance, an agent that succeeds on one task may fail entirely when the topology, location, or context shifts. Moreover, static datasets struggle to surface rare but important edge cases, which are crucial for robust evaluation yet infeasible to enumerate manually.

A natural way to address above challenges is to **dynamically generate queries** within the benchmark, as explored in recent work (Zhu et al., 2023; Zhang et al., 2024c; Yu et al., 2024). These methods typically construct symbolic graphs to synthesize diverse problem instances, focusing on domains such as arithmetic, logic, and program synthesis. However, existing dynamic generation methods do not generalize well to network and system applications. First, unlike well-defined mathematical tasks, networking problems often lack a deterministic structure, making the traditional synthesis of queries and ground truths inadequate (Zhu et al., 2023; Zhang et al., 2024c; Yu et al., 2024).

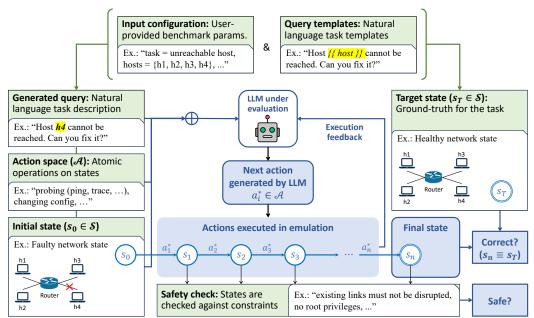


Figure 1: NETARENA introduces a unified *state-action* abstraction for network and system applications, integrating with real-world emulators to generate dynamic queries and ground truths, and enabling automated correctness and safety evaluation.

Troubleshooting routing misconfigurations, for instance, is rarely a linear, one-step process. Agents must engage in multi-turn interaction, collecting diagnostics, identifying root causes, and issuing fixes in context. Second, success in realistic network system tasks requires more than producing correct outputs: agents must avoid harmful side effects. Solutions must respect domain-specific constraints such as safety and latency, which synthetic data often fail to capture. For example, in a network with thousands of hosts, a misconfigured command can disable healthy paths and nodes, leading to cascading disruptions and service outages. Every command execution demands careful reasoning, and changes are only acceptable when both necessary and precisely scoped.

In this paper, we present NETARENA, a novel framework for dynamic benchmark generation for real-world network applications (Figure 1). NETARENA introduces a new evaluation paradigm, deploying agents in interactive, executable system environments to assess their capabilities through realistic, dynamically generated queries. Our key technical contributions are as follows:

- We define a **unified interface** for abstracting network applications based on explicit *state* and *action* spaces. This formalism supports dynamic query and ground truth generation (through executable state transitions), and enable controlled complexity scaling.
- By integrating with high-fidelity network emulators (e.g., Mininet (2022), Kubernetes (Google, 2024b)), Netarena enables **automatic**, **dynamic**, **and multi-turn verification** on LLM-generated actions, covering correctness, safety, and latency under deployment-like conditions.
- In NETARENA, users only need to specify high-level configurations (e.g., query count, complexity, task type). NETARENA **dynamically generates diverse evaluation sets** through stochastic sampling, which ensures broad coverage while reducing the risk of data contamination.

We instantiate NETARENA in three representative network applications: datacenter capacity planning, routing misconfiguration, and microservice policy troubleshooting. We evaluate five agents based on GPT-40 and QWen-72B models. Results on more LLMs and agents will appear on the leaderboard ¹. Our key findings are:

• Agent performance is strikingly low. Average correctness across tasks is only 24%, and even the best model stays below 60%. Small benchmarks (<200 queries) show high variance (average correctness rises to 38%), making statistical comparisons unreliable. NETARENA enables automated

¹This site shows a static, anonymized snapshot of the leaderboard for paper review: website link. The live website is in active use and will be shared upon publication.

Benchmark	Scale	Correctness(95% CI)	Safety/Latency	Contamination Risk	Generalizability
NeMoCopilot	33	94% [-]	N/A	High (Static)	Low (Management)
AI4OpsLab	48	59% [-]	N/A	High (Static)	Low (DevOps)
NetConfEval	3200	100% [-]	N/A	High (Static)	Low (Configuration)
NETARENA (Ours)	9,250 (unlimited)	44% [0.01, 0.14]	35% / 18s	Low (Dynamic)	High (Management, K8s, Routing, etc)

Table 1: NETARENA dynamically generate unlimited size of benchmarks for diverse network applications, supports more scalable and robust evaluation than existing benchmarks.

large-scale evaluation (e.g., >4000 queries), reducing confidence interval overlap between agents from as high as 85% to 0%, significantly improving evaluation reliability (§4.1).

- Correctness alone is insufficient. With metrics spanning correctness, safety, and latency, NETARENA exposes tradeoffs in agents' behavior. Some models produce correct answers that violate system constraints and are unsafe, while others act conservatively, preserving safety but failing to resolve issues within acceptable latency (§4.2).
- Supervised fine-tuning (SFT) behaves inconsistently. For correctness, SFT models often overfit to the complexity of their training data, and only the model trained on data spanning all levels of task difficulty generalizes well. For safety, the simplest-level SFT model surprisingly generalizes best across all tasks, outperforming those trained on harder levels. With controlled variation of task difficulty and multi-dimensional metrics, NETARENA enables such fine-grained analysis (§4.3).

Beyond evaluation, we discuss use cases for NETARENA: training reward models for on-policy reinforcement learning, and generating targeted adversarial queries to probe model weaknesses (§5).

2 RELATED WORK

General Benchmarks for Evaluating LLMs. A growing body of work has focused on benchmarking the reasoning and autonomy of LLM agents (Nathani et al., 2025; Chang et al., 2024; Gao et al., 2023; Ribeiro & Lundberg, 2022; Kiela et al., 2021; Ma et al., 2021; Lambert et al., 2024; Li et al., 2024a; 2023; Lei et al., 2023; Liang et al., 2022; Wang et al., 2024b; Yu et al., 2023; 2024; Zhong et al., 2023; Huang et al., 2024; Hendrycks et al., 2020; 2021; Starace et al., 2025; Bogin et al., 2024). For example, CORE-Bench (Siegel et al., 2024) aggregates reproducibility tasks from 90 papers to test agents' ability to rerun experiments. RE-Bench (Wijk et al., 2024) compares agent solutions on open-ended ML research tasks to those from expert engineers. MLE-Bench (Chan et al., 2024) converts 75 Kaggle competitions into agent benchmarks for leaderboard-based ML engineering. SWE-Bench (Jimenez et al., 2023) requires agents to resolve GitHub issues. Although these benchmarks are valuable in their own domains, they do not capture realistic network tasks that require deployment level reliability.

LLM Evaluation in Network and Other System Domains. In networking, LLMs have been used to generate graph-based network code (Mani et al., 2023), synthesize configuration files (Wang et al., 2024a), and support fault localization and remediation (Roy et al., 2024). LLMs have also been applied to extract protocol specifications (Sharma & Yegneswaran, 2023), reproduce networking experiments (Xiang et al., 2023; Kotaru, 2023), and evaluate system operations (Chen et al., 2025; Jha et al., 2025). Beyond networking, AIOpsLab (Chen et al., 2025) introduces 48 tasks for assessing agent performance in DevOps scenarios. WebVoyager (He et al., 2024) and WebArena (Zhou et al., 2023b) benchmark LLM agents through real-world website interaction tasks, while OSWorld (Xie et al., 2024) evaluates 369 operating system tasks. The Berkeley Function-Calling Leaderboard (BFCL) (Berkeley, 2025) measures agents' ability to correctly invoke APIs. Existing benchmarks in these domains are static and rely on expert-driven manual curation, which constrains their scalability and raises significant concerns regarding data contamination (detailed comparisons in Table 1).

Dynamic Benchmark Generation. A line of work aims to reduce benchmark contamination risk (Balloccu et al., 2024; Bender et al., 2021; Chen et al., 2021; Deng et al., 2023; Dong et al., 2024; Golchin & Surdeanu, 2023; Jacovi et al., 2023; Jiang et al., 2024; Li et al., 2024b; Li, 2023; Li & Flanigan, 2024; Oren et al., 2023; Roberts et al., 2023; Sainz et al., 2023; Shi et al., 2023; Zhang et al., 2024a; Zhou et al., 2023a; Li et al., 2025; Sun et al., 2024; Zhang et al., 2024c). DyVal (Zhu et al., 2023) and DyVal-2 (Zhu et al., 2024) introduces agents to generate and judge cognitively diverse variations of reasoning tasks. KIEval (Yu et al., 2024) dynamically conducts multi-turn

(a) Constructive: datacenter capacity planning (b) Reactive: routing misconfiguration Figure 2: Real-world network application examples.

knowledge-based interactions. LatestEval (Li et al., 2024b) constructs test sets from freshly published material to avoid overlap with model pretraining. Dysca (Zhang et al., 2024b) dynamically evaluating Vision LLMs by leveraging synthesis images. These dynamic benchmarks focus on general reasoning tasks, but are difficult to adapt to networking, where ground truth depends on system execution and cannot be reliably auto-generated.

3 NETARENA

162

163

164

165

166

167

168

169

170

171

172

173

174 175

176

177

178

179

181

182

183

185 186

187

188

189 190

191

192

193

194

196

197

199

200

201

202

203

204

205

206

207208

209

210

211

212

213

214

215

In this section, we summarize two representative network tasks that can potentially be automated via LLM-based agents, present a unified pipeline for dynamic query and ground truth generation, and describe how NETARENA automatically verifies each step of an agent's output (Figure 1).

3.1 LLM-BASED NETWORK APPLICATION TASKS

Real-world network applications require distinct forms of interaction between LLM agents and the task environment. We highlight two representative classes of these tasks:

- Constructive tasks require structured solutions to well-specified queries. These resemble "white-box" settings where the query expresses a clear intent, and agents must generate a policy update that fulfills that intent. In such cases, the LLM agent modify the current network state into a target state using interpretable operations. For example, in datacenter capacity planning (Figure 2a), users may request to find the optimal placement for new switches. The agent must synthesize a valid solution by composing various operations (e.g., add, rank, update). Although these tasks have deterministic outcomes, verifying correctness is challenging. It requires operational checks to ensure the solution adheres to policy constraints (e.g., the bandwidth must be over a minimum threshold). Appendix B.3 describes a full example.
- Reactive tasks involve diagnosing faults and issuing repairs in under-specified, evolving environments. These resemble "black-box" settings, where the query specifies the problem but it is unclear how to fix it. In such cases, the LLM agent must iteratively observe, hypothesize, and act to identify the correct solution. For example, in a routing misconfiguration task (Figure 2b), users may state, "this link is down, help me fix it." The agent must perform multiple steps of information gathering and action, such as inspecting interface states, identifying missing routes, and applying targeted fixes. Since these tasks involve multi-turn interaction, evaluation cannot rely on predefined solutions but must assess whether the intended outcome (e.g., restored connectivity) is achieved without introducing new risks at each turn. Appendix C.3 provides a detailed example.

3.2 A Unified Abstraction for Generating Network Benchmarks

Above real-world network application tasks go beyond static input-output matching, making scalable evaluation far more complex. Effective benchmarking requires principled task modeling that enables systematic generation of queries and ground truths. To address this, we define a unified pipeline with a general abstraction that consistently generates diverse queries and ground truths across applications under a single evaluation strategy.

State Transition Process. While network applications differ in detailed objectives, they often share a foundational structure: they operate over an underlying network/system topology (graph), and each

step of the interaction involves analyzing or modifying the state of this topology. Each application's task objective can be modeled as a finite state transition system (S, A, E), where S is the set of system states, A the set of atomic action functions, and E the application-specific execution function. Each action $a_t \in A$ is parameterized by task-specific operands θ_t and represented as $a_t(\theta_t)$.

A benchmark query defines an execution episode that begins from an initial state $s_0 \in \mathcal{S}$ and applies a sequence of T such parameterized actions $\{a_0(\theta_0), a_1(\theta_1), \dots, a_{T-1}(\theta_{T-1})\}$:

 $s_{t+1} = \mathcal{E}(s_t, a_t(\theta_t)), \quad \text{for } t = 0, \dots, T - 1$ (1)

To instantiate a new application in NETARENA, developers only need to define the state space \mathcal{S} (e.g., a routing topology with connectivity status) and the action space \mathcal{A} (e.g., IP (u) indicating a link-level IP error, where u is the operand link).

Query and Ground Truth Generation. We distinguish between the two types of tasks based on how queries and their corresponding ground truths are generated:

(1) Constructive tasks. These tasks start from a known initial state s_{init} , and the goal is to generate a sequence of actions that deterministically transition the system to a specified target state s_T . The ground truth is given by a predefined action sequence $\mathcal{A}^* = \{a_0^*, a_1^*, \dots, a_{T-1}^*\}$, which yields s_T under the execution function:

$$\mathcal{E}(s_0, \mathcal{A}^*) \triangleq \left(\mathcal{E}(\cdot, a_{T-1}^*(\theta_{T-1}^*)) \circ \cdots \circ \mathcal{E}(\cdot, a_1^*(\theta_1^*)) \circ \mathcal{E}(\cdot, a_0^*(\theta_0^*)) \right) (s_0) = s_T \tag{2}$$

Here, $\mathcal{E}_{a(\theta)}$ denotes applying a parameterized action to a state, and \circ represents functional composition. This captures the cumulative transformation of the system through sequential actions.

When generating a new query, NETARENA samples an initial state s_0 and a set of action \mathcal{A}^* , with each action's operand dynamically drawn from a large space. By executing these actions on s_0 , NETARENA produces the intended goal state s_T . A natural language template then converts $(s_0, \mathcal{A}^*, s_T)$ into a prompt for the LLM agent. The agent's correctness is evaluated by comparing its resulting state to s_T , and, when available, its predicted action sequence to \mathcal{A}^* to assess the agent's reasoning process.

(2) Reactive tasks. These tasks begin from a faulty network state s_{faulty} , generated by applying a hidden fault injection sequence $\mathcal{A}_{\text{inj}} = \{a_0^{\text{inj}}(\theta_0^{\text{inj}}), \dots, a_K^{\text{inj}}(\theta_K^{\text{inj}})\}$ to an originally healthy state s_0 :

$$s_{\text{faulty}} = \mathcal{E}(s_0, \mathcal{A}_{\text{inj}})$$
 (3)

Here, the fault injection sequence A_{ini} is hidden from the LLM agent.

When generating a new query, NETARENA injects errors into the original healthy state s_0 , producing a faulty state s_T that serves as the query input. A natural language template then describes s_T , and the LLM agent is tasked with recovering the system to s_0 . Unlike constructive tasks, multiple valid recovery paths may exist for a single faulty state. Correctness is therefore judged by whether the agent restores the system to s_0 , rather than by matching the specific injected action sequence \mathcal{A}_{ini} .

3.3 REALISTIC AGENT EVALUATION VIA EMULATOR INTEGRATION

A key challenge in deploying LLM agents for network applications is the uncertainty of their real-world behavior, including side effects, security risks, and inefficiencies. Testing these behaviors in real production environments is infeasible due to risk and cost (Chkirbene et al., 2024). To address this, NETARENA integrates directly with high-fidelity network emulators, providing controlled and reproducible evaluation under realistic conditions with diverse system metrics.

Comprehensive Performance Metrics from Emulator Integration. We embed agent evaluation in high-fidelity network emulators, which provide the closest deployment-like environments without the risks of production. These emulators are widely adopted in both academia and industry (e.g., Google, Alibaba, etc) for testing large-scale systems. Agent actions are executed end-to-end, and their effects are validated through emulator feedback. For instance, in a routing task, Mininet (2022) (Test-of-Time Award at SIGCOMM 2020) can verify whether connectivity is restored and whether new risks emerge, such as inadvertently disabling functional links. This integration forms the basis for evaluating agents across three core performance metrics.

• Correctness. We evaluate correctness by comparing the final network state produced by the LLM agent \hat{s}_{LLM} to the ground-truth state s^T , defined as s_T for constructive tasks and s_0 for reactive tasks. Correctness is defined as:

$$CORRECT(Q) = \mathbb{I}(\hat{s}_{LLM} \equiv s_T^*) \tag{4}$$

Here, $\hat{s}_{LLM} \equiv s_T^*$ denotes application-specific state equivalence, where equivalence may be syntactic (e.g., graph isomorphism) or functional (e.g., restored connectivity), depending on tasks.

• Safety. Safety evaluates whether LLM-generated actions satisfy task constraints C_Q , including structural invariants (e.g., no cross-layer violations) and operational guarantees (e.g., no unauthorized changes, no service disruption). For multi-turn execution with states s_0, s_1, \ldots, s_T , safety is checked at each step:

$$SAFE_{all}(Q) = \mathbb{I}\left(\forall t \in [1, T], \ s_t = \mathcal{E}(s_{t-1}, \hat{a}_{t-1}(\hat{\theta}_{t-1})) \ \land \ s_t \models \mathcal{C}_Q\right)$$

$$(5)$$

This formulation decouples final correctness from per-step safety, enabling fine-grained evaluation of agent behavior throughout the execution.

• Efficiency. Efficiency captures how quickly and compactly an agent completes the task. We track how many commands the agent issues and the end-to-end latency from query issuance to task completion. Efficient agents solve tasks with fewer commands and minimal latency, which is especially critical in time-sensitive scenarios like failure recovery.

With this design, NETARENA dynamically generates queries via randomized sampling in each evaluation round, reducing the risk of data contamination and ensuring agents are consistently tested on diverse, unseen tasks. Guidance on extending NETARENA to new applications is in Appendix §A.

4 EXPERIMENTS

Setup. We create five common LLM-based agents using two base models: GPT-40 (OpenAI, 2024) and QWen2.5-72B (Yang et al., 2024). Each model is paired with two prompting strategies: Chain-of-Thought (CoT) (Wei et al., 2022), Few-shot (Brown et al., 2020). We also evaluate a more advanced agent with ReAct (Yao et al., 2023) on GPT-40.

Representative Applications. To demonstrate generality, we implement three network applications:

- Capacity Planning (CP). Agents are evaluated on structured planning tasks over a realistic datacenter topology, based on Google's multi-layer abstraction (Mogul et al., 2020). Tasks include estimating bandwidth and modifying device configurations, spanning 12 action types (e.g., add, update, rank). Safety is checked by enforcing structural constraints and ensuring bandwidth meets minimum thresholds. Latency is measured as end-to-end solution time. (Details in §B)
- Routing Misconfiguration (Routing). Agents diagnose and repair dynamic faults (e.g., broken links or invalid forwarding rules) in Mininet (2022). This involves issuing diagnostic commands, interpreting outputs, and applying fixes to restore connectivity. Safety checks whether modifications improve the state without introducing new issues, and latency is measured by the number of iterations to resolution. (Details in §C)
- *Microservice Policy Deployment (K8s)*. Agents troubleshoot misconfigured Kubernetes network policies in Google's open-source microservice demo (Google, 2024b), aiming to restore valid interservice communication. Tasks involve identifying incorrect ports or overly restrictive rules. Safety metric follows the Routing definition, evaluating whether changes are necessary and effective, while latency counts the steps required for resolution. (Details in §D)

4.1 REDUCING CONFIDENCE INTERVAL OVERLAP WITH LARGER QUERY SIZE

A key advantage of NETARENA's dynamic generation is its ability to evaluate agents on large, diverse query sets, improving the statistical reliability of comparisons. Since correctness and safety are binary outcomes (i.e., pass/fail per query), we compute confidence intervals using the standard error of the mean (SEM) for a Bernoulli distribution: SEM = $\sqrt{\frac{\hat{p}(1-\hat{p})}{N}}$, where \hat{p} is the empirical success rate and N is the number of queries. We report 95% confidence intervals as $\hat{p} \pm 1.96 \cdot \text{SEM}$.

As shown in Figure 3, small query sets (CP:100, Route:150, K8s:150) produce wide error bars and overlapping intervals, making it difficult to distinguish agent performance. For example, on CP:100, GPT+ReAct overlaps by more than 50% with both QWen+CoT and QWen+Fewshot. Scaling to larger query sets with NETARENA (e.g., CP:5000) eliminates this overlap, revealing GPT+ReAct as the clear winner. Beyond accuracy, larger benchmarks also expose agents to richer task variations, reducing overfitting and enabling more robust generalization analysis.

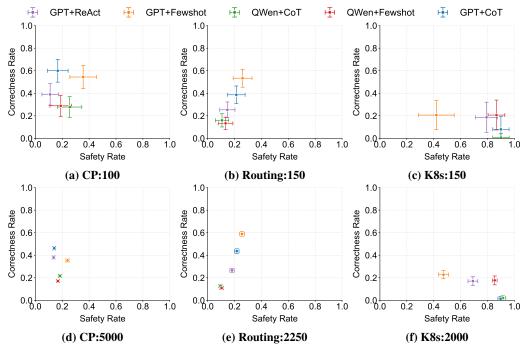


Figure 3: Increasing query size improve statistical confidence of comparisons.

Failure Type	Capacity Planning	Routing	Microservice Policy
Safety Violation	LLM adds a switch with 0 connection, violating datacenter constraints.	LLM assigns incorrect IPs, breaking existing network connectivity.	LLM deletes a running pod, causing immediate service disruption.
Control Logic Error	LLM calculates capacity at the wrong layer; it did not traverse the topology. LLM misorders commands, assigning an IP before activating the interface.		LLM uses incorrect patch, despite examples showing alternatives.
Operational Error	LLM hallucinates nonexistent node attributes, triggering execution failures.	LLM uses incorrect systemctl, trying to edit full system services.	LLM omits the namespace in a Kubernetes command, causing it to fail.

Table 2: Example failure types per application. Even with complete prompt context, LLMs frequently produce incorrect outputs.

The two-dimensional graph also highlights the importance of evaluating both correctness and safety, especially when agents have similar correctness rates. For example, on K8s:150, GPT+ReAct and QWen+Fewshot show comparable correctness. But on K8s:2000, GPT+ReAct exhibits a noticeably lower safety rate, making it riskier in practice.

4.2 FINE-GRAINED EVALUATION VIA COMPLEXITY-AWARE BREAKDOWN

While aggregate metrics are useful, they often hide critical weaknesses. A model may appear strong on simple queries but fail on compositional tasks or violate safety constraints in complex scenarios. To surface these issues, NETARENA applies complexity control during benchmark generation, annotating each query with action types and difficulty levels (Tables 4, 5, 6 in the Appendix), enabling finegrained analysis of correctness, safety, and latency.

As shown in Figure 4, this breakdown reveals clear differences across applications. In datacenter capacity planning (CP), both correctness and safety drop sharply as complexity grows, particularly for GPT-40+Few-shot. This reflects higher variance in complex tasks, where retrieval-style pattern matching fails to generalize and autonomous reasoning is required. We also find that all agents consistently struggle with add operations, which demand strict structural constraint satisfaction when introducing new nodes.

In routing and K8s tasks, GPT-40+Few-shot shows stronger performance in multi-turn, stateful diagnosis. Safety analysis, however, uncovers divergent behaviors. GPT-40 agents are often too aggressive, issuing unsafe fixes such as removing ingress+change protocol (RI+CPR)" in the K8s

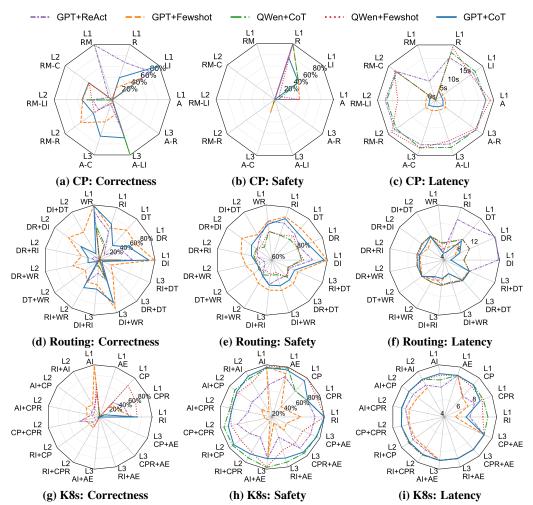


Figure 4: Breakdown across complexity levels (L1–L3, gets harder counter-clockwise). Agent strengths varies by application, metric, and task complexity.

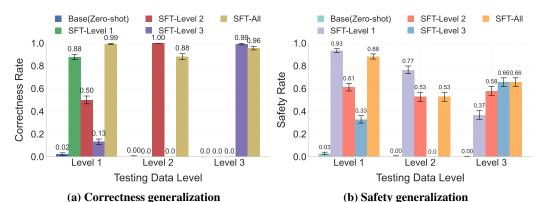
environment. Other models lean too conservative, frequently failing to act even when safe resolutions exist. Latency patterns vary as well: some agents resolve simple issues efficiently, while others generate long, redundant command sequences even for basic ingress fixes.

By moving beyond single-number scores, NETARENA provides a deeper view of each agent's planning strategies, failure modes, and generalization boundaries. This analysis not only strengthens benchmarking reliability but also offers actionable insights for improving LLM's deployment readiness. Table 2 provides further failure examples of agents in each application.

4.3 EVALUATING ROBUSTNESS VIA SUPERVISED FINE-TUNING

For constructive tasks where intermediate solutions can be automatically generated, NETARENA enables large-scale labeled data creation for supervised fine-tuning (SFT). To study generalization, we fine-tune four Qwen-7B models on different subsets of datacenter capacity planning queries: Level-1 (800), Level-2 (600), Level-3 (600), and a mixed dataset spanning all levels (2000). Each model is then evaluated across all levels using correctness and safety metrics (Figure 5).

Correctness results reveal clear signs of overfitting: Models excel on their training level but collapse on others; for instance, the SFT-Level-2 model achieves perfect correctness on Level-2 yet fails on Level-3. Only the mixed-level model generalizes, maintaining over 0.96 correctness across all levels. Interestingly, safety scores remain more stable. Even when correctness drops, models often preserve structural validity, suggesting that safety constraints transfer more easily across task complexities. For example, the SFT-Level-2 model fails on Level-3 correctness but still achieves reasonable safety, reflecting partial generalization of constraint adherence.



(a) Correctness generalization (b) Safety generalization Figure 5: Supervised Fine-tuned (SFT) model performance at different complexity level.

Overall, NETARENA provides a practical framework for producing scalable SFT examples for constructive type of network applications, which demonstrably boosts model performance. However, it is crucial to rigorously evaluate the SFT model's ability to generalize to unseen task types and configurations to prevent drawing overconfident conclusions.

5 CONCLUSION AND OTHER USE CASES OF NETARENA

We present NETARENA, a unified pipeline for dynamic LLM evaluation in real-world network applications. By abstracting tasks into *state-action* form, NETARENA enables controllable query generation with automatic ground-truth derivation. Its integration with high-fidelity emulators supports execution-time validation across correctness, safety, and latency. As such, NETARENA offers a practical foundation for developing, evaluating, and debugging LLM agents in safety-critical network domains. NETARENA also enables two prominent use cases below.

5.1 Enabling Post-Training RL in NetArena's Environments

Many reactive tasks lack step-level ground truth, making supervised fine-tuning infeasible and positioning RL as a natural alternative. However, RL requires reliable environments that **support interactive execution and feedback for reward computation**. NETARENA addresses this gap by integrating emulators that automatically generate step-wise feedback in response to agent actions,, enabling structured RL training and evaluation.

To test this, we fine-tune a QWen2.5-0.5B model (the largest feasible with our GPUs) using GRPO from TRL (HuggingFace, 2025) in a Mininet routing environment. Rewards include –100 for invalid commands, +10 for valid diagnostics, and +100 for correct fixes. The RL-finetuned model does not fully solve routing issues but consistently produces valid Mininet commands, outperforming its zero-shot baseline. This shows RL can shape useful policies even at small scales, though model capacity itself may limit its improvement in this experiment.

This experiment showcases that NETARENA can be used as a RL training environment. Beyond post-training RL, NETARENA also enables closed-loop self-improvement: as agents refine their reasoning traces and generate higher-quality action sequences, those traces can be incorporated into future RL episodes to improve agents.

5.2 PROBING AGENTS WITH ADVERSARIAL EXAMPLES

Understanding agent failure modes is critical for deployment. Unlike static benchmarks that cover a narrow task slice, NETARENA can dynamically generate adversarial test cases targeting specific weaknesses. By analyzing error types, failure modes, and inconsistent reasoning traces, we can identify task configurations that consistently degrade performance.

For instance, NETARENA exposes fine-grained control knobs (e.g., topology size, failure types, task complexity) that can be tuned to explore the model's capability boundaries. We propose RL or heuristic-guided sampling to iteratively generate harder queries based on prior failures. Over time, this adversarial loop reveals critical limitations, giving developers concrete insights into where generalization breaks and where extra training or safeguards are required.

Reproducibility statement. All experiments and results presented in this paper are fully reproducible using the open-source NETARENA framework available at the annoymous link: https://anonymous.4open.science/r/netarena_iclr2026-BE94/README.md.

The complete codebase includes three network applications: Capacity Planning (MALT) for datacenter topology management, Routing for network configuration in Mininet environments, and Kubernetes (K8s) for microservices orchestration using Google's Online Boutique demo.

A complete Docker environment (Dockerfile) with all dependencies pre-installed ensures consistent reproduction across different systems, including GPU support via the NVIDIA Container Toolkit for local model inference. The supplementary materials contain detailed parameter configurations, prompt templates, evaluation metrics implementation, and statistical analysis code that generated the confidence interval comparisons and performance breakdowns shown in our results figures.

Each application provides automated execution scripts under the experiments folder: (run_app_malt.sh, run_app_route.sh, run_app_k8s.sh) that reproduce our experimental configurations.

NETARENA supports multiple LLM agents, including GPT-40, Qwen2.5-72B, Gemini-Pro, Claude-4, and other open-source models. It offers prompt templates for Chain-of-Thought (CoT), Few-shot learning, and ReAct strategies, enabling seamless integration with various LLMs. Detailed integration instructions are provided in each application's README (app-k8s/README.md, app-malt/README.md, app-route/README.md).

REFERENCES

Simone Balloccu, Patrícia Schmidtová, Mateusz Lango, and Ondřej Dušek. Leak, cheat, repeat: Data contamination and evaluation malpractices in closed-source llms. <u>arXiv preprint arXiv:2402.03927</u>, 2024.

Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT), 2021.

UC Berkeley. Berkeley function-calling leaderboard. https://gorilla.cs.berkeley.edu/leaderboard.html, 2025.

Ben Bogin, Kejuan Yang, Shashank Gupta, Kyle Richardson, Erin Bransom, Peter Clark, Ashish Sabharwal, and Tushar Khot. Super: Evaluating agents on setting up and executing tasks from research repositories. arXiv preprint arXiv:2409.07440, 2024.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. Advances in neural information processing systems, 33:1877–1901, 2020.

Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio Starace, Kevin Liu, Leon Maksin, Tejal Patwardhan, et al. Mle-bench: Evaluating machine learning agents on machine learning engineering. arXiv:2410.07095, 2024.

Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. A survey on evaluation of large language models. <u>ACM Transactions on Intelligent Systems and Technology</u>, 2024.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code, 2021.

- Yinfang Chen, Huaibing Xie, Minghua Ma, Yu Kang, Xin Gao, Liu Shi, Yunjie Cao, Xuedong Gao, Hao Fan, Ming Wen, et al. Automatic root cause analysis via large language models for cloud incidents. In Proceedings of the Nineteenth European Conference on Computer Systems, pp. 674–688, 2024.
 - Yinfang Chen, Manish Shetty, Gagan Somashekar, Minghua Ma, Yogesh Simmhan, Jonathan Mace, Chetan Bansal, Rujia Wang, and Saravan Rajmohan. Aiopslab: A holistic framework to evaluate ai agents for enabling autonomous clouds. arXiv preprint arXiv:2501.06706, 2025.
 - Zina Chkirbene, Ridha Hamila, Ala Gouissem, and Unal Devrim. Large language models (llm) in industry: A survey of applications, challenges, and trends. In 2024 IEEE 21st International Conference on Smart Communities: Improving Quality of Life using AI, Robotics and IoT (HONET), pp. 229–234. IEEE, 2024.
 - Chunyuan Deng, Yilun Zhao, Xiangru Tang, Mark Gerstein, and Arman Cohan. Investigating data contamination in modern benchmarks for large language models. arXiv preprint arXiv:2311.09783, 2023.
 - Yihong Dong, Xue Jiang, Huanyu Liu, Zhi Jin, and Ge Li. Generalization or memorization: Data contamination and trustworthy evaluation for large language models. arXiv:2402.15938, 2024.
 - Irena Gao, Gabriel Ilharco, Scott Lundberg, and Marco Tulio Ribeiro. Adaptive testing of computer vision models. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 2023.
 - Sergiu Gatlan. Facebook: Outage caused by faulty routing configuration changes. https://www.bleepingcomputer.com/news/technology/facebook-outage-caused-by-faulty-routing-configuration-changes/, 2021.
 - Shahriar Golchin and Mihai Surdeanu. Time travel in llms: Tracing data contamination in large language models. arXiv preprint arXiv:2308.08493, 2023.
 - Google. Introducing gemini, your new personal ai assistant. https://gemini.google/assistant/?hl=en, 2024a.
 - Google. Microservices demo: Online boutique. https://github.com/GoogleCloudPlatform/microservices-demo, 2024b.
 - Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu, Yong Dai, Hongming Zhang, Zhenzhong Lan, and Dong Yu. Webvoyager: Building an end-to-end web agent with large multimodal models. arXiv preprint arXiv:2401.13919, 2024.
 - Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. arXiv:preprint arXiv:2009.03300, 2020.
 - Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. In <u>Advances in Neural Information Processing Systems (NeurIPS)</u>, 2021.
 - Yuzhen Huang, Yuzhuo Bai, Zhihao Zhu, Junlei Zhang, Jinghan Zhang, Tangjun Su, Junteng Liu, Chuancheng Lv, Yikai Zhang, Yao Fu, et al. C-eval: A multi-level multi-discipline chinese evaluation suite for foundation models. In Advances in Neural Information Processing Systems, 2024.
 - HuggingFace. Trl transformer reinforcement learning. https://github.com/huggingface/trl, 2025.
 - Alon Jacovi, Avi Caciularu, Omer Goldman, and Yoav Goldberg. Stop uploading test data in plain text: Practical strategies for mitigating data contamination by evaluation benchmarks. In <u>Proceedings</u> of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2023.

- Saurabh Jha, Rohan Arora, Yuji Watanabe, Takumi Yanagawa, Yinfang Chen, Jackson Clark, Bhavya Bhavya, Mudit Verma, Harshit Kumar, Hirokuni Kitahara, et al. Itbench: Evaluating ai agents across diverse real-world it automation tasks. arXiv preprint arXiv:2502.05352, 2025.
 - Minhao Jiang, Ken Ziyu Liu, Ming Zhong, Rylan Schaeffer, Siru Ouyang, Jiawei Han, and Sanmi Koyejo. Investigating data contamination for pre-training language models. <u>arXiv preprint</u> arXiv:2401.06059, 2024.
 - Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues? <u>arXiv preprint</u> arXiv:2310.06770, 2023.
 - Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, et al. Dynabench: Rethinking benchmarking in nlp. In Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), 2021.
 - Manikanta Kotaru. Adapting foundation models for operator data analytics. In <u>Proceedings of the 22nd ACM Workshop on Hot Topics in Networks</u>, pp. 172–179, 2023.
 - Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, et al. Rewardbench: Evaluating reward models for language modeling. arXiv preprint arXiv:2403.13787, 2024.
 - Fangyu Lei, Qian Liu, Yiming Huang, Shizhu He, Jun Zhao, and Kang Liu. S3eval: A synthetic, scalable, systematic evaluation suite for large language models. arXiv:2310.15147, 2023.
 - Changmao Li and Jeffrey Flanigan. Task contamination: Language models may not be few-shot anymore. In Proceedings of the AAAI Conference on Artificial Intelligence, 2024.
 - Xiang Li, Yunshi Lan, and Chao Yang. Treeeval: Benchmark-free evaluation of large language models through tree planning. arXiv preprint arXiv:2402.13125, 2024a.
 - Xiang Li, Yunshi Lan, and Chao Yang. Treeeval: Benchmark-free evaluation of large language models through tree planning. In <u>Proceedings of the AAAI Conference on Artificial Intelligence</u>, volume 39, pp. 24485–24493, 2025.
 - Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. Alpacaeval: An automatic evaluator of instruction-following models, 2023.
 - Yucheng Li. Estimating contamination via perplexity: Quantifying memorisation in language model evaluation. arXiv preprint arXiv:2309.10677, 2023.
 - Yucheng Li, Frank Guerin, and Chenghua Lin. Latesteval: Addressing data contamination in language model evaluation through dynamic and time-sensitive test construction. In Proceedings of the AAAI Conference on Artificial Intelligence, 2024b.
 - Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. Holistic evaluation of language models. arXiv preprint arXiv:2211.09110, 2022.
 - Zhiyi Ma, Kawin Ethayarajh, Tristan Thrush, Somya Jain, Ledell Yu Wu, Robin Jia, Christopher Potts, Adina Williams, and Douwe Kiela. Dynaboard: An evaluation-as-a-service platform for holistic next-generation benchmarking. In <u>Advances in Neural Information Processing Systems</u>, 2021.
 - Sathiya Kumaran Mani, Yajie Zhou, Kevin Hsieh, Santiago Segarra, Trevor Eberl, Eliran Azulai, Ido Frizler, Ranveer Chandra, and Srikanth Kandula. Enhancing network management using code generated by large language models. In Proceedings of the 22nd ACM Workshop on Hot Topics in Networks, pp. 196–204, 2023.

- Meta. Build the future of AI with Meta Llama 3. https://llama.meta.com/llama3, 2024.
- Mininet. Mininet: An instant virtual network on your laptop (or other pc). https://mininet.org/, 2022.
 - Jeffrey C Mogul, Drago Goricanec, Martin Pool, Anees Shaikh, Douglas Turk, Bikash Koley, and Xiaoxue Zhao. Experiences with modeling network topologies at multiple levels of abstraction. In 17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20), pp. 403–418, 2020.
 - Deepak Nathani, Lovish Madaan, Nicholas Roberts, Nikolay Bashlykov, Ajay Menon, Vincent Moens, Amar Budhiraja, Despoina Magka, Vladislav Vorotilov, Gaurav Chaurasia, et al. Mlgym: A new framework and benchmark for advancing ai research agents. <u>arXiv preprint arXiv:2502.14499</u>, 2025.
 - OpenAI. Hello gpt-4o. https://openai.com/index/hello-gpt-4o/, 2024.
 - OpenAI. Computer-using agent. https://openai.com/index/computer-using-agent/, 2025.
 - Yonatan Oren, Nicole Meister, Niladri Chatterji, Faisal Ladhak, and Tatsunori B Hashimoto. Proving test set contamination in black box language models. arXiv preprint arXiv:2310.17623, 2023.
 - Marco Tulio Ribeiro and Scott Lundberg. Adaptive testing and debugging of nlp models. In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL), 2022.
 - Manley Roberts, Himanshu Thakur, Christine Herlihy, Colin White, and Samuel Dooley. To the cutoff... and beyond? a longitudinal perspective on llm data contamination. In <u>Proceedings of the Twelfth International Conference on Learning Representations (ICLR)</u>, 2023.
 - Devjeet Roy, Xuchao Zhang, Rashi Bhave, Chetan Bansal, Pedro Las-Casas, Rodrigo Fonseca, and Saravan Rajmohan. Exploring Ilm-based agents for root cause analysis. <u>arXiv:2403.04123</u>, 2024.
 - Pascal Sager, Benjamin Meyer, Peng Yan, Rebekka von Wartburg-Kottler, Layan Etaiwi, Aref Enayati, Gabriel Nobel, Ahmed Abdulkadir, Benjamin F. Grewe, and Thilo Stadelmann. AI agents for computer use: A review of instruction-based computer control, GUI automation, and operator assistants. CoRR, abs/2501.16150, 2025. doi: 10.48550/ARXIV.2501.16150. URL https://doi.org/10.48550/arXiv.2501.16150.
 - Oscar Sainz, Jon Campos, Iker García-Ferrero, Julen Etxaniz, Oier López de Lacalle, and Eneko Agirre. Nlp evaluation in trouble: On the need to measure llm data contamination for each benchmark. In Findings of the Association for Computational Linguistics (ACL), 2023.
 - Prakhar Sharma and Vinod Yegneswaran. Prosper: Extracting protocol specifications using large language models. In <u>Proceedings of the 22nd ACM Workshop on Hot Topics in Networks</u>, pp. 41–47, 2023.
 - Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi Chen, and Luke Zettlemoyer. Detecting pretraining data from large language models. <u>arXiv preprint</u> arXiv:2310.16789, 2023.
 - Zachary S Siegel, Sayash Kapoor, Nitya Nagdir, Benedikt Stroebl, and Arvind Narayanan. Corebench: Fostering the credibility of published research through a computational reproducibility agent benchmark. arXiv preprint arXiv:2409.11363, 2024.
 - Giulio Starace, Oliver Jaffe, Dane Sherburn, James Aung, Jun Shern Chan, Leon Maksin, Rachel Dias, Evan Mays, Benjamin Kinsella, Wyatt Thompson, et al. Paperbench: Evaluating ai's ability to replicate ai research. arXiv preprint arXiv:2504.01848, 2025.
 - Liangtai Sun, Yang Han, Zihan Zhao, Da Ma, Zhennan Shen, Baocai Chen, Lu Chen, and Kai Yu. Scieval: A multi-level large language model evaluation benchmark for scientific research. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 38, pp. 19053–19061, 2024.

- Changjie Wang, Mariano Scazzariello, Alireza Farshin, Simone Ferlin, Dejan Kostić, and Marco Chiesa. Netconfeval: Can Ilms facilitate network configuration? Proceedings of the ACM on Networking, 2(CoNEXT2):1–25, 2024a.
 - Siyuan Wang, Zhuohan Long, Zhihao Fan, Zhongyu Wei, and Xuanjing Huang. Benchmark self-evolving: A multi-agent framework for dynamic llm evaluation. arXiv preprint arXiv:2402.11443, 2024b.
 - Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. <u>Advances in</u> neural information processing systems, 35:24824–24837, 2022.
 - Hjalmar Wijk, Tao Lin, Joel Becker, Sami Jawhar, Neev Parikh, Thomas Broadley, Lawrence Chan, Michael Chen, Josh Clymer, Jai Dhyani, et al. Re-bench: Evaluating frontier ai r&d capabilities of language model agents against human experts. arXiv preprint arXiv:2411.15114, 2024.
 - Qiao Xiang, Yuling Lin, Mingjun Fang, Bang Huang, Siyong Huang, Ridi Wen, Franck Le, Linghe Kong, and Jiwu Shu. Toward reproducing network research results using large language models. In Proceedings of the 22nd ACM Workshop on Hot Topics in Networks, pp. 56–62, 2023.
 - Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh J Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, et al. Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments. <u>Advances in Neural Information Processing Systems</u>, 37:52040–52094, 2024.
 - An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. <u>arXiv preprint</u> arXiv:2412.15115, 2024.
 - Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In <u>International Conference on Learning Representations (ICLR)</u>, 2023.
 - Dingli Yu, Simran Kaur, Arushi Gupta, Jonah Brown-Cohen, Anirudh Goyal, and Sanjeev Arora. Skill-mix: a flexible and expandable family of evaluations for ai models. In <u>Proceedings of the Twelfth International Conference on Learning Representations (ICLR)</u>, 2023.
 - Zhuohao Yu, Chang Gao, Wenjin Yao, Yidong Wang, Wei Ye, Jindong Wang, Xing Xie, Yue Zhang, and Shikun Zhang. Kieval: A knowledge-grounded interactive evaluation framework for large language models. arXiv preprint arXiv:2402.15043, 2024.
 - Hugh Zhang, Jeff Da, Dean Lee, Vaughn Robinson, Catherine Wu, Will Song, Tiffany Zhao, Pranav Raja, Dylan Slack, Qin Lyu, Sean Hendryx, Russell Kaplan, Michele Lunati, and Summer Yue. A careful examination of large language model performance on grade school arithmetic, 2024a.
 - Jie Zhang, Zhongqi Wang, Mengqi Lei, Zheng Yuan, Bei Yan, Shiguang Shan, and Xilin Chen. Dysca: A dynamic and scalable benchmark for evaluating perception ability of lvlms. arXiv:2406.18849, 2024b.
 - Zhehao Zhang, Jiaao Chen, and Diyi Yang. Darg: Dynamic evaluation of large language models via adaptive reasoning graph. arXiv preprint arXiv:2406.17271, 2024c.
 - Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. Agieval: A human-centric benchmark for evaluating foundation models. arXiv preprint arXiv:2304.06364, 2023.
 - Kun Zhou, Yutao Zhu, Zhipeng Chen, Wentong Chen, Wayne Xin Zhao, Xu Chen, Yankai Lin, Ji-Rong Wen, and Jiawei Han. Don't make your llm an evaluation benchmark cheater. <u>arXiv</u> preprint arXiv:2311.01964, 2023a.
 - Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, et al. Webarena: A realistic web environment for building autonomous agents. arXiv preprint arXiv:2307.13854, 2023b.

Kaijie Zhu, Jiaao Chen, Jindong Wang, Neil Zhenqiang Gong, Diyi Yang, and Xing Xie. Dyval: Dynamic evaluation of large language models for reasoning tasks. In The Twelfth International Conference on Learning Representations, 2023.

Kaijie Zhu, Jindong Wang, Qinlin Zhao, Ruochen Xu, and Xing Xie. Dynamic evaluation of large language models by meta probing agents. arXiv preprint arXiv:2402.14865, 2024.

No.	Step	Description
1	Emulator Selection	Select a high-fidelity emulator that matches the target network application. Examples: Kubernetes (microservices), Mininet (routing/switching).
2	Topology Generator	Implement a topology creation function. Input parameters: number of nodes, edge density, topology constraints (e.g., connectivity). Output: a full topology object deployable in the emulator.
3	Task Type Selection	Choose the task category: Constructive tasks – LLMs propose configurations; Reactive tasks – LLMs respond to injected errors or alerts.
4	Operation Library (Constructive)	Define basic operations representing atomic configuration steps. These should be modular and composable, and parameterizable (e.g., AddNode, LinkSwitch). Used to generate complex configuration queries.
5	Error Injection (Reactive)	Define typical error types for reactive diagnosis tasks (e.g., link failure, misconfiguration). These errors are injected into the environment to test LLM diagnostic capabilities.
6	Evaluation Metrics	Define evaluation criteria: Correctness – does the LLM achieve the intended effect?; Safety – does it interrupt the existing service?; Latency – how fast does the LLM generate valid responses?

Table 3: Construction steps for creating new applications in NETARENA.

A EXTEND NETARENA TO NEW APPLICATIONS

NETARENA is designed for easy extensibility across diverse network and system domains through a standardized API. To add a new application, users define three key components: (1) the application's *state* and *action* spaces, (2) a backend emulator or simulator that connects to NETARENA's evaluation framework (such as Mininet or Kubernetes), and (3) three core metrics customized to the task. Table 3 presents this process in a clear step-by-step format. By lowering the barrier to incorporating realistic, execution-grounded tasks, NETARENA becomes a generalizable platform for benchmarking LLM agents in real-world infrastructure reasoning scenarios. We encourage the community to extend NETARENA to new protocols, topologies, and deployment settings, supporting robust and reproducible agent evaluation at scale.

B DETAILS FOR DATACENTER CAPACITY PLANNING

The first application focuses on applying LLM agents for datacenter capacity planning—a critical aspect of network lifecycle management that optimizes resource utilization and minimizes provisioning costs. Effective capacity planning depends on accurate network topology representations at different abstraction levels. High-level abstractions help network operators evaluate overall bandwidth needs between data centers, while detailed low-level views enable engineers to manage individual device configurations and connections efficiently.

Based on Google's multi-layer topology abstraction (Mogul et al., 2020) and publicly available datasets, we build a simulation environment modeling a realistic datacenter topology with 5,493 nodes and 6,424 edges, encompassing 10 distinct device types such as packet switches, ports, and chassis. Nodes have attributes like physical port capacity and are interconnected following hierarchical constraints reflective of real datacenter structures. To enable dynamic interaction between LLM agents and representative capacity-planning tasks, we define 12 operational types (e.g., 'update,' 'add,' 'count,' 'rank'), each supporting a wide range of diverse and detailed queries. For instance, an 'add' operation may involve simple tasks, like attaching a new port to a switch, or complex scenarios, such as integrating a new packet switch into an aggregation block. Each query instance is generated dynamically and randomly via i.i.d. sampling, reducing data contamination risk. LLM agents generate executable code for each query, which is evaluated against dynamically generated ground-truth code execution results. A query is considered successfully resolved if the agent's generated code executes correctly, and without introducing any security vulnerabilities or violations of the datacenter hierarchical constraints.

B.1 APPLICATION ENVIRONMENT

To enable LLM agents to interact with data center capacity planning, we adopt Google's multi-layer topology abstraction (Mogul et al., 2020) and use their released datasets as the foundational topology. We also develop a simulator to evaluate the performance of deployed capacity planning algorithms across key metrics.

Topology. The capacity planning topology consists of 5,493 nodes and 6,424 edges. Each node represents different device types within the data center (e.g., packet switch, port, chassis), with a total of 10 distinct node types. Each node may also have attributes; for instance, a port may have a physical capacity attribute, which specifies its capacity. These nodes are interconnected by edges, where, for example, an edge between a packet switch and a port indicates that the switch contains the corresponding port. To simulate the abstraction of a real data center, not all nodes can be arbitrarily connected. Figure 8 illustrates the hierarchical dependencies between the nodes.

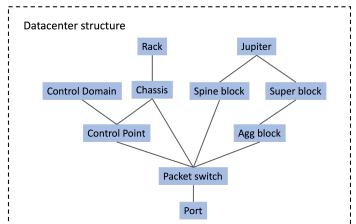


Figure 6: Hierarchical dependencies between nodes in the datacenter topology.

B.2 Dynamic Benchmark Creation

Basic Operations and Operands. We implement six fundamental operations: *add*, *count*, *update*, *remove*, *list*, and *rank*. Each operation defines a specific action, while the operands represent the datacenter entities or elements the operations act upon. A user query is interpreted as a combination of these operations applied to dynamically generated operands based on the query type. For example, consider the query: "If we add a new packet switch to each chassis, what will be the new total capacity on all chassis?" This query can be broken down into the following steps: (1) *List* all chassis nodes. (2) *Add* a new packet switch node to each chassis node. (3) *Count* the total capacity on updated chassis nodes.

Query Complexity Control. Using the basic operations, queries can be dynamically constructed by combining them with appropriate operands. To manage complexity, we categorize queries based on two factors: the number of operations involved and the type of control sequence utilized. Control sequences can vary in structure, including sequential combinations of multiple actions, conditional *If-Else* statements, loops such as *For-loops* followed by sequential actions, or *For-loops* combined with *If-Else* statements. The operands associated with these basic operations are determined dynamically for each query, allowing the framework to remain flexible and adaptable to diverse query types and scenarios.

Ground Truth Generation. Dynamically generating ground truth algorithms is the most complex process. To accomplish this, we first manually implement the algorithm for each basic operation, supported by dynamic operands, as functional modules. These modules are then combined into templates to generate more detailed queries and their corresponding ground truth. For instance, when the randomly selected operation types are "add" and "count", the system randomly selects new node types and parent nodes. A unique node name is created, and a corresponding natural language query is formed: "Add *child_node_name* to *parent_node_name*. Count the *child_node* in *parent_node_name* in the updated graph." The system subsequently generates a Python function

by combining the "add" and "count" algorithms into a new function, which represents the ground truth. This function includes the steps of adding the selected node to the graph, linking it to its parent node, and performing a counting query on the updated graph.

B.3 LLM-BASED AGENTS USAGE

Emulator. To emulate datacenter capacity planning tasks using large language model (LLM) agents, we construct a comprehensive evaluation pipeline encapsulated in Python, which simulates real-world decision-making by executing model-generated code against a dynamic infrastructure graph and validating correctness and safety. The evaluator begins by initializing with a network graph and a specified LLM agent, alongside the appropriate prompting strategy. When a user query is issued, the LLM agent generates Python code intended to analyze or modify the network graph; this code is executed to produce a structured result. A built-in safety checker then verifies that the updated graph adheres to expected datacenter constraints, including valid node types, edge formats, topological hierarchies, bandwidth configurations, and the presence of ports on switches. In parallel, a golden (ground-truth) answer is executed and evaluated in the same manner. The framework compares LLM and ground-truth outputs using type-specific strategies—such as structural isomorphism for graphs or strict equality for lists and text—and logs detailed outcomes regarding correctness, safety violations, and execution latency. Errors such as incorrect output types, faulty logic, or unsafe graph mutations are captured and recorded in a structured JSON format for downstream analysis. This emulator not only enables robust benchmarking of LLM agents under system-level constraints but also facilitates the generation of diverse labeled data for supervised fine-tuning and safe deployment of AI agents in real datacenter environments.

Prompts. We provide the initial prompt as below.

Generate the Python code needed to process the network graph to answer the user question or request. The network graph data is stored as a networkx graph object. The Python code you generate should be in the form of a function named process_graph that takes a single input argument graph_data and returns a single object return_object. The input argument graph_data will be a networkx graph object with nodes and edges. The graph is directed, and each node has a name attribute to represent itself. Each node has a type attribute, in the format of EK_TYPE. Each node can have other attributes depending

type attribute, in the format of EK_TYPE. Each node can have other attributes depending on its type. Each directed edge also has a type attribute, which can include RK_CONTAINS or RK_CONTROL. You should check the relationships based on edges and check the name based on node attributes.

The node types have the following hierarchy:

- EK JUPITER contains EK SPINE BLOCK
- EK_SPINE_BLOCK contains EK_AGG_BLOCK
- EK_AGG_BLOCK contains EK_PACKET_SWITCH
- EK_CHASSIS contains EK_CONTROL_POINT
- EK_CONTROL_POINT contains EK_PACKET_SWITCH
- EK_RACK contains EK_CHASSIS
- EK PACKET SWITCH contains EK PORT
- EK_SPINE_BLOCK contains EK_PACKET_SWITCH
- EK_CONTROL_DOMAIN contains EK_CONTROL_POINT
- EK CHASSIS contains EK PACKET SWITCH
- EK JUPITER contains EK SUPER BLOCK
- EK_SUPER_BLOCK contains EK_AGG_BLOCK

Adding new nodes requires considering the attributes of the new node. You should also consider adding edges based on their relationships with existing nodes. The name to add on each layer can be inferred from the new node's name string. When adding new nodes, you should also add edges based on their relationship with existing nodes.

Packet switch nodes have a switch location attribute switch_loc in the node. PORT nodes have an attribute physical_capacity_bps. When calculating the capacity of a node, you need to sum the physical_capacity_bps on the PORT nodes within the hierarchy that contains this node.

B.4 EVALUATION METRICS

LLM agents are tasked with queries related to generating Python-based capacity planning algorithms. These algorithms take the original network graph as input and produce outputs that depend on the query, either direct answers or updated topologies. The simulator environment evaluates the following metrics.

- **Correctness:** The output's correctness is evaluated by comparing the LLM-generated results with the ground truth. If the outputs match exactly, the LLM's answer is labeled as correct.
- Safety: Safety ensures the structural and attribute integrity of the network graph by verifying that the LLM's output adheres to all defined constraints. This includes validating node types, edge types, and hierarchical relationships, checking the presence of mandatory attributes (e.g., physical capacity for PORT nodes), ensuring no isolated nodes exist, and enforcing connectivity rules between related components.

• Latency: Latency measures the execution time of the Python code generated by the LLM. This metric excludes the time spent on LLM prompting and response, focusing solely on the effectiveness and runtime performance of the generated solution.

C DETAILS FOR ROUTING MISCONFIGURATIONS

The second application focuses on applying LLM agents for network routing configuration troubleshooting—an essential network management task that involves quickly identifying and resolving routing-related issues such as misconfigured paths, link failures, and network congestion. Effective troubleshooting is critical in practice, as unresolved routing problems can cause severe performance degradation or outages. For instance, a major outage at Meta in October 2021 was caused by faulty configuration changes to backbone routers. This misconfiguration disrupted communication between data centers, leading to a global service outage and significant downtime (Gatlan, 2021).

To evaluate LLM agents on routing troubleshooting tasks, we construct a simulated network environment using Mininet, featuring a router connected to multiple switches and hosts organized into distinct subnets, enabling realistic network interactions. Within this environment, we dynamically inject various routing misconfigurations—such as incorrect forwarding rules or broken links—that disrupt connectivity and cause failures in the pingmesh (host-to-host connectivity) tests. LLM agents is tasked to perform diagnostics by executing network commands (e.g., inspecting routing tables, interface statuses, and error logs), analyzing the results, and proposing corrective actions. An evaluation query is considered successfully resolved if, after the agent's intervention, the network connectivity is restored, verified by the successful execution of the 'pingall()' method, without introducing new security issues.

C.1 APPLICATION ENVIRONMENT

To evaluate the network troubleshooting capabilities of LLMs, we designed an experimental setup where LLM behaves like a network engineer to interact with a simulated network environment. Among the different network simulation tools available, we selected Mininet for its lightweight nature, ease of use, and support for custom dynamic topologies. Within Mininet, we created dynamic network topologies and intentionally injected errors to simulate realistic network faults. The LLM was then tasked with diagnosing and resolving these issues, demonstrating its ability to perform automated troubleshooting in a controlled environment.

Topology. In Mininet, we will construct a network topology that includes a router, multiple switches, and host computers. The router is connected to all the switches and is responsible for forwarding traffic between different subnets. The switches, in turn, are connected to the hosts, enabling them to communicate within their respective subnets and interact with other devices on the network. To create various network topologies, we use two variables, num_switches and num_hosts_per_subnet, to control the setup. num_switches defines how many switches are present in the network, while num_hosts_per_subnet specifies how many hosts are connected to each switch. In our environment setup, the number of subnets and the number of hosts per subnet range from 2 to 4, making the topology dynamic and complex.

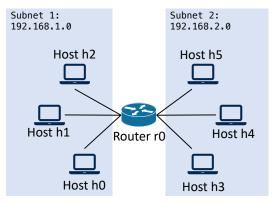


Figure 7: A simple version of routing network topology.

C.2 DYNAMIC BENCHMARK CREATION

Query Generation In the Topology section, we designed a network where all nodes are initially able to connect with one another. In our benchmark, each query represents a network state with an

error, causing some nodes to be unable to communicate with certain others. We aim to have the LLM handle these error states and attempt to resolve the issues, ensuring that the hosts in the network can reconnect and communicate with each other.

To generate a problematic network, we start with a valid network configuration and apply a series of erroneous configuration commands to introduce faults. These faulty commands are selected from one of five basic categories of network errors, each with specific details that define the nature of the error. Importantly, within each error category, the details of the error can vary, leading to different manifestations of the same type of fault. This variability within error categories is what allows our benchmark to generate a diverse range of dynamic, distinct network errors, ensuring that the LLM is tested under various realistic fault conditions.

The five basic categories of network errors are as follows. We will describe each error type in detail, along with how the details within each error category can vary, ensuring a diverse set of dynamic network faults.

Error Type 1: Disable Routing

 The error type disable_routing is designed to simulate a failure in IP forwarding, which is crucial for routing packets between different subnets. In this error scenario, we disable IP forwarding using one of several methods. These methods are varied to generate diverse network configurations, making the benchmark dynamic and comprehensive.

There are four methods to disable IP forwarding. Method 1 uses sysctl to globally disable packet forwarding across all routes, while Method 2 utilizes iptables to drop all forwarded packets, affecting communication between subnets. Method 3 applies ip rule to prohibit forwarding based on specific rules, offering finer control. Method 4, also using iptables, drops traffic from a randomly selected subnet, creating a localized failure. These methods contribute to error diversity allows the benchmark to simulate a wide range of network issues.

Error Type 2: Disable Interface

The error type disable_interface simulates a failure by interfering with a specific network interface, which is crucial for communication between devices. In this error scenario, we have three methods to destroy a healthy network by disabling some interfaces.

Method 1 uses ifconfig to bring the interface down, resulting in a complete loss of connectivity through that interface. Method 2 utilizes ip link to achieve the same outcome of disabling the interface, but with a different network management tool. Method 3 changes the Maximum Transmission Unit (MTU) of the interface, which can cause packet fragmentation or loss if the MTU is set too low, leading to communication problems. These methods contribute to error diversity by offering various levels of disruption to simulate a wide range of interface failures.

Error Type 3: Remove IP

The error type remove_ip is designed to simulate a failure by modifying or removing the IP address of a network interface. This can cause disruption in the network's ability to route packets, particularly if the interface's IP is essential for communication. The error is injected by using one of four methods to modify the IP address of the specified interface.

Method 1 flushes the IP using ip addr flush, making the interface unreachable. Method 2 assigns a random IP within the 10.0.0.0/24 subnet, which can cause conflicts or incorrect routing if the IP is already in use. Method 3 assigns a wrong subnet mask (e.g., 8, 16, 30, 31, or 32), preventing communication due to misconfigured subnetting. Method 4 assigns a duplicate IP from another subnet, causing conflicts or routing issues, with a fallback IP used if no other subnets are available.

Error Type 4: Drop Traffic to/From Subnet

The error type <code>drop_traffic_to_from_subnet</code> simulates a failure by manipulating the traffic to or from a specific subnet. This can disrupt communication between the subnet and other network components. The error is injected by using one of four methods to modify the flow of traffic to/from the subnet.

Method 1 drops all incoming and outgoing traffic to/from the subnet using iptables. Method 2 actively rejects traffic to/from the subnet using iptables rules. Method 3 blocks ICMP traffic, preventing ping communication with the subnet. Method 4 introduces network delay using to netem, adding latency to the traffic. These methods provide diverse ways to manipulate network traffic, enabling the benchmark to simulate various types of network disruptions.

Error Type 5: Wrong Routing Table

The error type wrong_routing_table simulates a failure by misconfiguring the routing table, which can disrupt the network's ability to properly forward traffic between subnets. The error is injected by modifying the routing table using one of four methods. Each method changes the routing behavior differently, enabling the benchmark to simulate various types of routing issues.

Method 1 removes an existing route and adds a new route using a different interface. Method 2 adds a route with an incorrect gateway, potentially causing traffic to be misdirected. Method 3 adds a route with a very high metric, which makes it less preferred compared to other routes, possibly causing traffic delays or incorrect routing. Method 4 creates a routing loop by adding a route that forwards traffic to another subnet, which could lead to a loop of network traffic.

Increasing Error Complexity: Combining Basic Errors

The code above demonstrates the injection of basic network errors, but to generate more complex network states, we combine multiple basic errors. Specifically, we pair two different categories of errors and inject them sequentially into the network. This approach creates more intricate and realistic network failures, providing a better test for the LLM's diagnostic capabilities. By combining multiple errors, the network's behavior becomes more complex, leading to a more challenging scenario for the LLM. The result of running the pingall command becomes more complicated, with more nodes failing to communicate. The LLM is faced with greater challenges, as it must analyze multiple potential causes, troubleshoot through various methods, and draw on a wider set of diagnostic skills. This makes the task significantly more difficult than diagnosing single errors in isolation.

The process of injecting two errors works as follows:

- 1. Single Error Injection: If only one error is to be injected (i.e., errornumber equals 1), the code only injects one single basic error. We use a function called process_single_error to handle this error type and inject it into the network.
- 2. Multiple Error Injection: When injecting more than one error, errortype and errordetail are lists containing the error types and their corresponding details. Our benchmark uses a loop to pair elements from these lists and sequentially inject each error by calling process_single_error function for each pair.
- By combining different types of basic errors, the benchmark creates more complex network states that better simulate real-world network failures, providing the LLM with a more challenging test environment.

C.3 LLM-BASED AGENTS USAGE

Emulator. In the Mininet simulator, an LLM is allowed to execute commands to retrieve information about network states and analyze the results of the pingall command. The pingall command sends ICMP echo requests (ping) from every host to each other in the network and reports the results. This is typically used to verify the connectivity between all nodes in a network, ensuring that the network is running as expected. A successful pingall output indicates that all hosts can communicate with each other, while failure may suggest network issues such as misconfigurations or connectivity problems.

The LLM can also utilize several types of commands to gather information about the network, diagnose issues, and propose solutions. To gather network information and diagnose issues, the LLM can use various commands such as ifconfig, ip addr, ip link, ip route. These commands allow the LLM to examine network configurations like IP address, network interfaces, routing tables. By analyzing the output from these commands, the LLM can identify problems such as incorrect configurations and suggest corrective actions to resolve the network issues.

```
1188
            Example Usage PingAll Results:
1189
            Command: PingAll
1190
            *** Fast Ping: testing ping reachability
1191
            h1 \rightarrow X X X X X X r0
1192
            h2 \rightarrow X X X X X X r0
1193
            h3 \rightarrow X X X X X X
1194
            h4 \rightarrow X X X X X r0
1195
            h5 \rightarrow X X X X X X r0
1196
            h6 \rightarrow X X X X X T 0
1197
            r0 \rightarrow h1 h2 X h4 h5 h6
1198
            *** Results: 76% dropped (10/42 received)
1199
             Example feedback from Mininet:
            Command: ip route
             192.168.1.0/24 dev r0-eth1 proto kernel scope link src 192.168.1.1
1201
            192.168.2.0/24 dev r0-eth2 proto kernel scope link src 192.168.2.1
1202
            192.168.3.0/24 dev r0-eth3 proto kernel scope link src 192.168.3.1
1203
            192.168.4.0/24 dev r0-eth4 proto kernel scope link src 192.168.4.1
            192.168.5.0/24 dev r0-eth5 proto kernel scope link src 192.168.5.1
1205
             192.168.6.0/24 dev r0-eth6 proto kernel scope link src 192.168.6.1
1206
```

Prompts. We provide the initial prompt as below.

You need to behave like a network engineer who finds the root cause of network issues and fixes them in a routing application.

There is a Mininet network with problems in the router r0, causing the network to be partially disconnected. Some nodes cannot successfully ping other nodes. Your task is to fix these issues so that the pingall result shows all connections are successful.

I recommend using diagnostic commands to gather information about the router and network to identify the cause of the problem. Once you have sufficient information and understand the root cause, provide commands to fix the issue.

When implementing your solution, be careful not to disrupt existing connected edges - your commands should not cause previously working connections to break.

Please provide your output in JSON format with the keys 'machine' and 'command'. You can only issue one command at a time as I can only execute commands sequentially.

Important notes:

1207

1208 1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1223

1224

1225

1226

1227

1228 1229

1230

1231

1232 1233

1234

1237

1239

1240

1241

- The router's name may not be exactly r0. It may have a prefix (like p29_r0).
- The same applies to host names and interface names (e.g., p29_h1, p29_h2, p29_r0-eth1, p29_r0-eth2).
- The prefix could be anything (p29, p30, p31, etc.).
- Do not include sudo in your commands.
- You are not permitted to use the vtysh command.
- Do not use ping commands as the ping results are already provided to you.

I will provide you with the latest PingAll() feedback from the network along with your previous actions and their results to help you diagnose the problem.

C.4 EVALUATION METRIC

The performance evaluation of LLMs requires a comprehensive and diversified approach. A single metric only reflects the final success or failure of a command, but overlooks the intermediate steps in tasks such as network fault diagnosis, which often involve multiple iterations. For such tasks, it is essential to consider not only the end result, but also the necessity and efficiency of each intermediate step, which can be assessed through the number of iterations. Furthermore, evaluating whether these intermediate actions affect overall performance or even cause potential damage to the network is crucial as this represents a safety issue. By considering multiple dimensions of evaluation, we can gain a more holistic understanding of LLMs' capabilities, address their limitations, and ensure that

they are safe, reliable, and effective for real-world deployment. Here are the evaluation metrics we considered.

- Correctness: This metric focuses on whether the final command execution results in a successful outcome, such as whether a pingall command succeeds. It reflects the accuracy of the LLM's final action in diagnosing and resolving faults. A high correctness rate indicates that the LLM reliably produces solutions that lead to successful outcomes without errors.
- Safety: Safety is assessed by examining the LLM's ability to preserve network stability during diagnosis and resolution. We expect the LLM to avoid issuing commands randomly, as arbitrary configuration changes could disrupt the network. In a real-world scenario, such actions could result in severe consequences, such as network downtime. Therefore, we aim to evaluate whether the LLM issues commands responsibly, ensuring it gathers sufficient information before taking action. If a command is issued to configure the network but fails to resolve the issue, the LLM will be penalized with a lower score, reflecting the unsafe nature of its response.
- Latency: This metric measures the number of iterations the LLM requires to reach a successful resolution. Fewer iterations to resolve the issue indicate that the LLM is more efficient in troubleshooting, as it can pinpoint the root cause of the problem accurately and provide the appropriate commands. An LLM with fewer iterations demonstrates higher efficiency in resolving network issues, leading to quicker resolutions and reduced network downtime. This improved efficiency is crucial for minimizing disruption, as it shows that the LLM is able to address network problems with precision and speed, ensuring better overall performance and reliability.

D DETAILS FOR MICROSERVICE POLICY DEPLOYMENT TROUBLESHOOTING

The third application targets on troubleshooting tasks in Kubernetes and microservice policy configurations. Kubernetes (K8s), as the dominant orchestration platform, efficiently manages microservices by automating deployment, scaling, and orchestration, which significantly improves service scalability, resilience, and agility. However, misconfigured policies or faulty service deployments in Kubernetes clusters can lead to significant operational issues, including service downtime, performance bottlenecks, or critical security vulnerabilities.

To evaluate whether LLM agents can autonomously identify and resolve Kubernetes network policy misconfigurations, we create a simulation environment based on Google's publicly available microservice benchmark (Google, 2024b). Specifically, we leverage the Kubernetes-based microservice application composed of 11 services (e.g., frontend, checkout, payment) communicating via gRPC, secured by 13 distinct network policies that define permissible interactions among nodes and their ports. We dynamically generate realistic troubleshooting scenarios by injecting various types of network-policy misconfigurations. For each scenario, the LLM agent autonomously investigates the system, diagnoses the root cause, and attempts corrections to restore intended network connectivity. The LLM agent is deemed to have solved the query if it restores node communication, verified through connectivity tests, demonstrating its ability to automate Kubernetes configuration troubleshooting.

D.1 APPLICATION ENVIRONMENT

To evaluate whether LLM agents can autonomously identify and resolve Kubernetes network policy misconfigurations, we create a simulation environment based on Google's publicly available microservice benchmark (Google, 2024b). This benchmark is based on Online Boutique, a cloud-first microservices demo application. Online Boutique is a web-based e-commerce platform where users can browse products, add them to their cart, and complete purchases. For our simulation, we use a local Kubernetes cluster to emulate the Online Boutique environment, providing an ideal setup for testing Kubernetes network policies in a real-world, cloud-native scenario.

The diagram above illustrates the architecture of the Online Boutique application, which consists of several microservices. In the diagram, each service is represented by a node, and the arrows indicate one-way communication between them. The one-way access is crucial for ensuring security and maintaining the integrity of the system. By restricting communication to only one direction, we can better control the flow of data and prevent unauthorized access or interactions between microservices, which is important for preserving the confidentiality and reliability of each service.

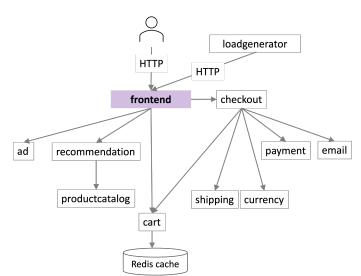


Figure 8: Structure of Google's microservice benchmark.

Our benchmark is designed to intentionally disrupt this structure by causing failures in the communication between the microservices. We will inject network misconfigurations into the cluster to simulate faults, which could lead to unauthorized access or prevent legitimate nodes from accessing the required services. The faulty cluster will then be provided to an LLM, allowing it to autonomously identify the root cause of the issues and take corrective actions, such as restoring proper access controls or resolving connectivity problems. This process will test the LLM's ability to handle complex network policies and restore the system to its desired state.

D.2 DYNAMIC BENCHMARK CENERATION

In our benchmark design, we focus on error injection within a K8s cluster, specifically targeting the network configuration layer. To emulate real-world misconfigurations and generate diagnostic queries, we systematically inject faults by automatically modifying the network configuration YAML files. Our methodology is grounded on five fundamental types of basic network errors, which can also be added sequentially to increase complexity. These basic errors involve modifications to the ingress and egress rules, protocols, and ports within the Kubernetes network policies. By introducing such errors, we can observe how network disruptions manifest and how the LLM reacts to misconfigurations, which are critical for evaluating LLM's fault detection and correction abilities.

Basic Network Policy Errors. We generate examples of wrong policy deployed in the network as below.

• Add Ingress Rule: This error type involves adding new ingress rules to the network policy, which allows inbound traffic from external or unauthorized sources. In the YAML configuration, this would involve adding a new from rule under the ingress section.

```
ingress:
    - from:
    - podSelector:
         matchLabels:
         app: frontend
```

Potential Impact: Adding an ingress rule with an external source could compromise the security boundaries of the system, allowing unauthorized access to internal services, which might lead to security breaches.

• Add Egress Rule: This error type adds a new egress rule to allow internal services to communicate with external services or networks, potentially violating security protocols. The YAML change typically involves adding a new to rule under the egress section.

```
1350 egress:
1351 - to:
1352 - podSelector:
1353 matchLabels:
1354 app: frontent
```

Potential Impact: Adding an egress rule might cause unauthorized data leaks or expose the system to connections with external entities. It can allow sensitive data to flow out of the network, risking privacy or system integrity.

• **Remove Ingress Rule:** Deleting an ingress rule can block necessary traffic from reaching internal services. In the YAML configuration, this would involve removing an existing ingress rule under the ingress section.

Potential Impact: By removing an ingress rule, valid traffic from authorized services may be blocked, leading to service downtime or failure to respond to incoming requests. This can reduce system availability and affect service reliability.

• **Change Protocol:** Changing the communication protocol (e.g., from TCP to UDP or vice versa) can disrupt inter-service communication. The YAML configuration might be modified by altering the protocol field under the ports section.

```
ports:
    - port: 9555
    protocol: UDP % Changed from TCP to UDP
```

Potential Impact: If services depend on a specific protocol for communication, changing it can result in connectivity issues. Services may fail to establish connections, causing service interruptions and degraded performance.

• Change Port: Modifying the port number used for communication between services can lead to issues such as services becoming unreachable or port conflicts. This would be represented by modifying the port value under the ports section of the YAML configuration.

```
ports:
    - port: 8080 % Changed from 9555 to 8080
    protocol: TCP
```

Potential Impact: Changing the port configuration can make services unreachable if other services still expect the old port. Port conflicts can arise if another service is already using the new port, resulting in failed connections and network instability.

Increasing Complexity through Sequential Error Injection. To increase the difficulty and complexity of network misconfigurations, we adopt a sequential error injection strategy. This approach involves identifying and modifying two separate YAML files, each representing different aspects of the network configuration. By injecting different types of errors into each file, we create a more intricate and challenging scenario for the Kubernetes cluster. Sequential injection of multiple error types forces LLMs to handle progressively more complex network misconfigurations, which can test the LLM's ability to diagnose and resolve issues more effectively.

D.3 LLM-BASED AGENTS USAGE

Emulator. In this section, we describe the setup of a emulation platform designed to facilitate interactions between LLMs and a K8s cluster. The goal of this environment is to create a testing

ground where LLMs can be used to diagnose and resolve network misconfigurations within the K8s cluster. The environment allows us to introduce various network issues, test LLM-based troubleshooting techniques, and evaluate the efficiency and accuracy of LLM interventions.

For our simulation, we utilized the kind simulator, a tool designed for running local Kubernetes clusters using Docker container "nodes". Using kind, we can simulate an entire Kubernetes cluster on a single virtual machine (VM), which provides a lightweight and easy-to-manage environment for our tests. In addition, we deploy Google's microservices demo locally, benefiting from the compact nature of the setup and the simplicity of management. Besides, the Kind simulator fully replicates the behavior of a real Kubernetes cluster, ensuring that the environment that we provide to the LLM is consistent with production-like conditions. This allows us to present the LLM with an environment that mimics real-world Kubernetes clusters, enabling accurate and reliable testing of its fault detection and resolution capabilities.

Once the local K8s cluster was successfully set up, we can inject errors to simulate network misconfigurations and connectivity issues. After the error injection, we have a function to collect connectivity failures due to configuration errors between nodes. This information is then provided to the LLM, which will help it analyze the current network state and generate a Kubernetes command to resolve the identified issue. The command is executed within the environment, and the updated connectivity status is subsequently fed back to the LLM. Equipped with this new information, the LLM interacts with the system, refining its diagnosis, and issuing additional commands if necessary. This iterative process continues until the network issues are resolved, enabling the LLM to effectively troubleshoot and fix network problems. Through this approach, the benchmark assesses the diagnostic accuracy, problem solving efficiency, and adaptability of the LLM in dynamic network environments, providing a comprehensive assessment of its performance in real-world scenarios. The following provides the implementation details for the network status check and the interaction with LLMs.

Network Status Check. We have implemented a function that automatically checks the network connectivity between nodes in the K8s cluster to identify any discrepancies between the actual communication and the expected connectivity.

At the beginning of the network status check, a debug container (which will be reused for the entire benchmarking process) is created for each pod, containing basic network tools necessary for connectivity testing. During the network status testing process, we access the debug container and use the 'nc' (Netcat) command to test whether a pod can communicate with other pods within the cluster. The results of these tests are then provided to the LLM in the form of a mismatch report, detailing which nodes have communication issues that differ from the expected connectivity.

Example Mismatch

Mismatch Summary:

frontend → adservice:9555 (Expected: True, Actual: False)

frontend → cartservice:7070 (Expected: False, Actual: True)

Explanation:

In the above mismatch results, the connectivity between pods is compared against the expected behavior. The first mismatch indicates that the 'frontend' pod was expected to communicate with the 'adservice' pod on port 9555, but the actual connectivity failed (Expected: True, Actual: False). In contrast, the second mismatch shows that the 'frontend' pod was expected not to communicate with the 'cartservice' pod on port 7070, but the actual connection was established (Expected: False, Actual: True).

LLM Interaction. To enable effective interaction between the LLM and the K8s cluster, we allow the LLM to use Kubernetes commands to troubleshoot and resolve network issues. When the LLM analyzes the network status and identifies potential problems, it generates Kubernetes commands as output. These commands are then extracted from the LLM's response and executed on the K8s cluster.

To run the generated Kubernetes commands, we use Python's 'subprocess' module, which allows us to programmatically execute shell commands. The command's output, including any errors or status messages, is captured and stored for further analysis. This enables us to monitor the LLM's actions and track the results of the commands it issues.

By interacting with the K8s cluster in this way, the LLM is able to perform network troubleshooting tasks, such as diagnosing misconfigurations and resolving connectivity issues. The model uses the output from each command to refine its diagnosis, iterating on its approach if necessary. Through this process, the LLM can progressively identify the root causes of network problems and issue further corrective commands until the issues are resolved, completing the troubleshooting process.

Prompts. We provide the initial prompt as below.

You need to behave like a network engineer who can find the root cause of network policy deployment issues and fix them in the microservices architecture.

Our microservices architecture contains following services and desired communication relationships:

- User and loadgenerator can access the frontend service via HTTP.
- frontend communicates with the following services: checkout, ad, recommendation, productcatalog, cart, shipping, currency, payment and email.
- checkout further communicates with payment, shipping, email, and currency.
- recommendation communicates with product catalog.
- cart communicates with the Redis cache for storing cart data.

Your task is to inspect the current network policies and verify if they meet the described communication patterns. If there are any mismatches, you should fix them.

How the interaction works:

- Provide one command at a time to check connectivity or node accessibility.
- Each time, I will give you the previous commands and their corresponding outputs.
- I will also provide the current connectivity status, including any mismatches between the expected and actual connectivity status.
- Use this information to identify and fix misconfigurations step-by-step.

D.4 EVALUATION METRIC.

The performance evaluation of LLMs in the K8s cluster troubleshooting requires a comprehensive and multidimensional approach. A single metric only reflects the final success or failure of a command but overlooks the intermediate steps in tasks like network fault diagnosis, which often involve multiple iterations. For tasks in a K8s cluster, it's important to evaluate not only the final result but also the necessity and efficiency of each intermediate action, which can be assessed by the number of iterations. Additionally, it's crucial to evaluate whether these intermediate actions impact overall performance or cause potential damage to the cluster, as this poses a safety risk. By considering these different dimensions, we can gain a more thorough understanding of the LLM's capabilities, identify its limitations, and ensure that it is safe, reliable, and effective in real-world K8s deployments. The evaluation metrics we considered are as follows:

- Correctness: This metric focuses on whether the final command execution results in a successful outcome. We will use the network status check in the previous section to find if the K8s cluster gets to the expected state with the LLM's help. It reflects the accuracy of the LLM's final action in diagnosing and resolving network faults in a K8s environment. A high correctness rate indicates that the LLM reliably produces solutions that lead to successful outcomes without errors.
- Safety: Safety is measured by evaluating the LLM's ability to maintain the K8s cluster's stability during the diagnosis and resolution process. This includes monitoring network status during the LLM troubleshooting process to see if the LLM's command will destroy original connectivity. For instance, an increase in pod failures or a loss of network connectivity could indicate that the LLM's commands are destabilizing the existing system state, potentially leading to larger failures in production environments.
- Latency: This metric measures the number of iterations the LLM requires to reach a successful
 resolution of the issue within a K8s cluster. Fewer iterations to resolve the issue indicate that the
 LLM is more efficient in troubleshooting, as it can more accurately pinpoint the root cause of the
 problem and generate the appropriate Kubernetes commands. An LLM that reaches a solution in
 fewer iterations demonstrates higher efficiency in addressing K8s network issues, leading to quicker
 resolutions and reduced downtime. This efficiency is essential in a production environment, where
 minimizing network disruption is critical to maintaining service availability and system reliability.

15	13	
15	14	
15	15	

	Query Example	P
	Remove jul.a4.m3.s2c8 from the graph. List the direct child nodes of jul.a4.m3 in the updated graph.	
Level 1	Rank all child nodes of type EK_CONTROL_DOMAIN with name jul.al.dom based on the physical_capacity_bps attribute.	
Level 1	List all the child nodes of ju1.a2.dom. Return a list of child node names.	
	Add a new PORT with new_EK_PORT_65 and type=EK_PORT to the node jul.a3.m2.s2c1.	
	Remove jul.a2.m2.s2c4.pl3. Count the number of nodes with type=EK_PORT under jul.a2.m2.s2c4 in the updated graph.	
Level 2	Remove jul.a4.m4.s3c8 from the graph. List the direct child nodes of jul.a4.m4.	
	Remove jul.a4.m2.s2c3. Rank the child nodes of jul.a4.m2 based on the total bandwidth.	
	Add new_EK_PACKET_SWITCH_13 to ju1.a3.m3. Count the number of type=EK_PACKET_SWITCH under ju1.a3.m3.	
Level 3	Add new_EK_PACKET_SWITCH_61 to jul.s4.dom. List the direct child nodes of jul.s4.dom in the updated graph.	
	Add my pagymm gyrmgy CR to to 1 2 Doub the shill and of to 1 2 bond on the start bondwidth	





Action Label R Α RM-C RM-LI RM-R A-C A-LI $Add \ \texttt{new_EK_PACKET_SWITCH_67} \ to \ \texttt{jul.a3.dom}. \ Rank \ the \ child \ nodes \ of \ \texttt{jul.a3.dom} \ based \ on \ the \ total \ bandwidth.$

Table 4: Action and complexity level details for queries in CP.

	Error Details	Error Label
	disable_routing	DR
	disable_interface	DI
Level 1	remove_ip	RI
	drop_traffic_to_from_subnet	DT
	wrong_routing_table	WR
	disable_routing + disable_interface	DR+DI
	disable_routing + remove_ip	DR+RI
	disable_routing + drop_traffic_to_from_subnet	DR+DT
Level 2	disable_routing + wrong_routing_table	DR+WR
	remove_ip + wrong_routing_table	RI+WR
	drop_traffic_to_from_subnet + wrong_routing_table	DT+WR
	disable_interface + drop_traffic_to_from_subnet	DI+DT
	disable_interface + wrong_routing_table	DI+WR
Level 3	remove_ip + drop_traffic_to_from_subnet	RI+DT
	disable interface + remove ip	DI+RI

Table 5: Error and complexity level details for queries in Routing.

	Error Details	Error Label
	remove_ingress	RI
	add_ingress	AI
Level 1	change_port	CP
	change_protocol	CPR
	add_egress	AE
	remove_ingress + add_ingress	RI+AI
	remove_ingress + change_port	RI+CP
Level 2	remove_ingress + change_protocol	RI+CPR
Level 2	add_ingress + change_port	AI+CP
	add_ingress + change_protocol	AI+CPR
	<pre>change_port + change_protocol</pre>	CP+CPR
	change_port + add_egress	CP+AE
Level 3	change_protocol + add_egress	CPR+AE
Level 5	remove_ingress + add_egress	RI+AE
	add_ingress + add_egress	AI+AE

Table 6: Error and complexity level details for queries in K8s.