



A Holistic View of AI-driven Network Incident Management

Pouya Hamadani[†], Behnaz Arzani[§], Sadjad Fouladi[§], Siva Kesava Reddy Kakarla[§],
Rodrigo Fonseca[#], Denizcan Billor^{||}, Ahmad Cheema^{||}, Edet Nkposong^{||}, Ranveer Chandra[§]
[†]MIT, [§]Microsoft Research, [#]Azure Systems Research, ^{||}Microsoft

Abstract

We discuss the potential improvement large language models (LLM) can provide in incident management and how they can overhaul the ways operators conduct incident management today. We propose a holistic framework for building an AI helper for incident management and discuss the several avenues of future research needed to achieve it.

We thoroughly analyze the fundamental requirements the community should consider when designing such helpers. Our work is based on discussions with operators of a large public cloud provider and their prior experiences both in incident management and with attempts to improve the incident management experience through various forms of automation.

CCS Concepts

• Networks → Network management.

Keywords

Large Language Models, Incident Management

ACM Reference Format:

Pouya Hamadani, Behnaz Arzani, Sadjad Fouladi, Siva Kesava Reddy Kakarla, Rodrigo Fonseca, Denizcan Billor, Ahmad Cheema, Edet Nkposong, and Ranveer Chandra. 2023. A Holistic View of AI-driven Network Incident Management. In *The 22nd ACM Workshop on Hot Topics in Networks (HotNets '23)*, November 28–29, 2023, Cambridge, MA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3626111.3628176>

1 Introduction

“No design works unless it embodies ideas that are held common by the people for whom the object is intended.”
—Adrian Forty

Frequent *incidents*—failures that compromise the reliability of essential services—pose a challenge for large-scale cloud operations [20, 24]. Prior research has focused on the Incident Management (IM) process which finds and mitigates these

incidents quickly [1, 10, 13, 22, 26, 46]. We present a network operations view of how Large Language Models (LLMs) can accelerate and improve the IM process. We discuss the requirements that any LLM-based solution for IM must meet.

When automation or customers report an incident, the incident manager assigns it to an On-Call Engineer (OCE) who investigates the issue, mitigates it (*e.g.*, moves traffic away from the problematic area or takes a failed device offline), finds the root cause, and fixes it. The mitigation step is the most important: it stops the incident’s impact on customers. Providers view Time to Mitigation (TTM) as the main indicator of efficiency and strive to keep it within 30 minutes (for more details on the IM process, see prior work [13, 22, 24]).

The IM process is time-consuming, exhausting, and stressful for the OCE because the set of possible root causes and monitoring data is large. OCEs also need to understand a diverse set of systems and their components and go through multiple hours of training to develop enough expertise to understand how different components interact.

A prevailing question is: should we utilize LLMs to improve the IM process [1, 13, 46] *i.e.*, create an LLM-powered OCE-helper? LLMs can ingest large datasets, generate structured insights from unstructured text, automate workflows, and generate human-like written content [8]: they introduce an unexplored design continuum in systems research. LLMs such as GPT-3.x [7] and PaLM [14] have already helped with tasks such as code generation [12], formal verification [21], database query generation [31, 48]. But LLMs struggle with long-term planning [8], and are prone to hallucination [8, 41].

Our answer to this question—whether one *should* use LLMs in this context—is nuanced: due to the unstructured format of IM, universal OCE-helpers are infeasible without LLMs, *but* cannot solve the end-to-end problem directly. We investigated feasible OCE-helpers and the design requirements. This led us to a modular design where we incorporate LLMs to help in different stages of the IM workflow. Our design bears a similar insight to goal-driven LLM agents [37, 47, 55].

While the implications may not be as severe, IM process shares similarities to medical diagnosis. Operators (medical practitioners) want to mitigate incidents (treat symptoms), understand what caused the incident, and fix (cure) the problem. The space of possible hypotheses (underlying diseases) is large, and they cannot test each and every one. Instead, they pick a likely hypothesis, do targeted tests, and reassess their initial



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike International 4.0 License.

HotNets '23, November 28–29, 2023, Cambridge, MA, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0415-4/23/11.

<https://doi.org/10.1145/3626111.3628176>




Figure 1: The three modules in our framework: hypothesis formation, hypothesis testing, and mitigation planning.

guess—doing so from incident summaries (symptoms) is non-trivial. They also have to be careful in how they apply mitigations (treatments) to avoid further harm to the system (patient)—the risks involved mean an OCE-helper cannot replace an operator (medical practitioner) outright. This is why we design the OCE-helpers as assistants instead of replacements for OCEs.

We interviewed OCEs with 5+ years of experience and found three key requirements for any OCE-helper, which we describe through examples based on real incidents. Initial work on such helpers shows promise [1, 13], but to realize their full potential, we have to take a holistic view of IM and identify how helpers can meet the fundamental requirements described in this work.

We found an OCE-helper should be: *iterative, reliable and safe, and adaptive*. An iterative helper forms and tests hypotheses and recommends mitigation strategies through a multi-step process. A reliable and safe helper does not cause extra harm to the network, and *adaptive* helper can evolve and adjust as operators deploy new components, make configuration changes, and discover new issues. We discuss these requirements in depth.

We propose a framework that meets these key requirements (§4) and explain the practical considerations and insights that led us to it. This framework (Figure 1) is analogous to OCE’s natural thought process and includes three modules: hypothesis former, hypothesis tester, and mitigation planner. The hypothesis former offers plausible explanations for an incident. The hypothesis tester aims to test these hypotheses, querying available monitors and tools and interpreting the results. The mitigation planner creates a mitigation plan based on the validated hypotheses, which the OCE can execute at their discretion.

This ambitious framework needs further research to realize it in areas such as: (1) network-specialized LLMs to serve as backbones to these modules; (2) network-specialized embedding models that help knowledge-retrieval methods (e.g., vector databases); and (3) structural approaches that breakdown complex tasks in network management (e.g., isolate links) to LLM-digestible steps. The community needs to approach problems such as risk estimation and safety [2, 38] differently and think through how to: (1) analytically evaluate the end-to-end risk of the LLM proposals; (2) define abstractions between LLMs and existing tools to gracefully balance the trade-off between automation and safety; and (3) create verifiable LLM-based tools (e.g., text-to-SQL [48]).

Our goal for this paper is to motivate the community to further explore this space and think about the implications such solutions have for network management and beyond.

2 Foundational Principles: Design

We analyzed real incidents across various teams, services, and time-frames and interviewed OCEs with several years of experience, with a focus on their day-to-day operations, pain points, and where they thought automation could help. We studied prior work, explored the different designs for OCE-helpers, and then ran through what-if scenarios where we took a specific design for an OCE-helper, replayed historical incidents and observed how this hypothetical helper would integrate into the IM workflow. Three main themes emerge which we believe are foundational principles any such design should meet: iterative prediction, reliable and safe, and adaptive. These principles seem obvious; yet prior works do not fully meet them.

Iterative Prediction. It is simple to formulate this problem as a Machine Learning (ML) prediction problem which entices us to train a one-shot feed-forward model that predicts root cause or mitigation, conditioned on past incidents and telemetry we carefully engineer: prior work [1, 13] takes the (predefined) incident information (e.g., title, summary, and the auto-generated digest of system health and telemetry) and predicts the root cause or mitigation without any additional context. These approaches are promising: they work well with incidents similar to those resolved in the past, or where operators have fleshed out clean monitoring signals. But they fail with more complex (often more impacting¹) or novel incidents.

Challenging incidents deviate from known patterns, are sparse, and often happen due to a complex chain of events. There is not enough information in the initial incident summary of many networking incidents to find the right mitigation or root cause in one-shot—operators need to make multiple informed attempts to test various hypotheses, *safely* mitigate the incident, zero in on the root cause, and resolve the problem. The one-shot approach falls short in these cases, and such a design is inherently too restrictive for the IM workflow.

Consider the `Casc-1` incident from a report on Google Cloud incidents [24] (Figure 2). At the time, Google’s network was comprised of two Wide Area Networks (WANs), namely B2 and B4. During a network upgrade, a transient configuration inconsistency (event 1) caused more than one cluster to observe B4 with several IP prefixes (event 2). The traffic controller that managed traffic across B2 and B4 mistakenly interpreted that as B4’s failure (event 3) and rerouted all B4 traffic through B2 (event 4). This led to an overload (event 5) and packet loss (event 6). A one-shot predictor has to infer event 1 after it observes event 6, and a naïve OCE-helper that

¹These works don’t report accuracy per severity of incidents they studied and this discussion is based on our experience in IM.




Figure 2: Iterative vs. one-shot design. One-shot predictors need to make logical leaps with no extra information. Iterative chaining mimics the natural thought process of OCEs.

only observes predefined information would decide it is either a device failure (*e.g.*, a switch) or a transient increase in traffic.

An OCE checks both of these suggestions first, but would quickly reassess after they gather more evidence that shows both are incorrect: high traffic utilization of B2 (event 5) and no traffic for B4. They would hypothesize and confirm that traffic is intentionally rerouted away from B4 (event 4). This cascade of hypothesis, testing, and reassessment would continue until they discover the inconsistency and the bug. We believe an OCE-helper should behave the same, *i.e.*, only an OCE-helper that hypothesizes, tests, and re-evaluates its decisions in a feedback loop can successfully resolve such complex incidents.

Reliable & Safe. The OCE-helper should not make expensive mistakes [20] on its first day: risk assessment cannot be an afterthought. Even smaller mistakes can drive operators away—if they can’t reliably show when drastic cases happen, they will bypass the helper with the mere threat of a mistake [22]: we need *safe* (should not worsen the situation) and *reliable* (should not have inconsistencies that introduce mistakes) helpers.

Helpers need to include risk mitigating mechanisms both baked in—where it provides a reason for why it arrived at a particular response (these mechanisms often improve LLM responses [52])—and on the outside, *e.g.*, a wrapper around the helper which statistically analyzes the mitigation the helper proposes [38]. We should not just analyze what helpers suggest: we should explicitly design the helper to take in feedback from all risk assessors, and search for actions with lower risks.

If we intend to eventually automate OCE-helpers, confidence and risk measures are non-negotiable. We discuss why prior work does not enable such a solution and the research the community needs to fill the gap in §4.4.

Adaptive. A core challenge with IM is that network components’ software and hardware evolve rapidly—often teams don’t explicitly coordinate with each other when they update components they own. Such evolution invalidates prior incidents and mitigation strategies, and uncoordinated changes lead to new incidents [5, 24]. Such “new” incidents are complex and have the highest negative impact, but also where we see the most potential for helpers to improve the IM workflow and where we believe research in this space should focus.

Consider a failure reported in AWS Direct Connect Tokyo [44] (Figure 3). This service provides low-latency consistent tunneling to the Tokyo region. New software and a perfect storm of




Figure 3: One-shot methods need to learn similar incidents. Iterative methods only need to learn incremental changes.

events led to this failure, where customers saw latency spikes and packet loss. Operators deployed a new protocol that reacted faster to network failures. This protocol had an unknown defect triggered by a specific set of packet headers and payloads. Months later, traffic from one customer happened to fit that pattern, and several network devices failed to forward traffic. Operators first removed these devices to mitigate the incident, but the same failure popped up again on other devices. Operators eventually realized the culprit, disabled the protocol, and resolved the defect. This was a unique incident, partly because the events that led to it involved new and untested protocols. No amount of historical incidents could supply a helper with the knowledge to mitigate such an incident.

It is difficult to adapt to new incidents because there is little “learn-able” data in such sparse novel scenarios: we know the changes, but are unaware what impact they may cause until they happen. We cannot use end-to-end approaches (*e.g.*, one-shot helpers) for new incidents since they need end-to-end samples the change causes. An iterative helper that explicitly *reasons* about individual components and how operators change them can gradually build towards the right mitigation. Operators only need to update this helper with the new behavior of the system and not its impact. This helper could evolve almost as quickly as OCEs to any changes in the infrastructure, guidelines, and incidents, and does not need end-to-end samples.

An OCE-helper that safely assists with incidents needs to satisfy these principles. While they are stringent, and there is no current automated OCE-helper that satisfies them, we have designed a framework that addresses all of them systematically (§4). We first describe how we can verify and measure whether any such framework can satisfy these requirements.

3 Foundational Principles: Evaluation

It is hard to evaluate solutions in this space. Prior works’ metrics (*e.g.*, F_1 score, semantic similarity) paint an incomplete picture of the impact these helpers have on metrics operators care about: how fast they resolve the incident and whether they can do so safely—alone, they will not reveal the helper’s ability to meet our requirements. Instead, we need to track: (1) TTM; and (2) the overheads the helper’s mistakes induce (new incidents caused by wrong mitigation plans, SLA violations, etc).

We can A/B test a helper to conduct an end-to-end evaluation of the target system, *i.e.*, randomly assign incidents to either a

helper-enhanced intervention or to a helper-free control group. This is the most robust evaluation we can get and with enough cases and statistical tests, we can compare the two groups at any statistic (e.g., average TTM). This is nonetheless challenging, because of the high variability in the nature of each incident, among OCE’s expertise, and other confounding factors.

It is hard to estimate the overhead of the helper’s mistakes—the complex interactions between the system state and the mitigation make it hard to model; A/B tests are one mechanism to do so but they are also costly and invasive. Many incidents (e.g., customer-reported ones) are diverse and heavy-tailed: we need a long test span to collect enough measurements.

We can instead ‘replay’ historical incidents and compare the replayed TTM to the original; we can use this to scale up the measurements. We can query telemetry retroactively, but this will not work if the mitigation the helper suggests differs from the one the operator used in the original incident.

We can naïvely solve this problem and report the TTM savings for incidents where the mitigations match and the fraction where this was not the case. Can we do better? We can find past incidents where operators used the same mitigation the helper proposed and estimate the impact of that mitigation on the TTM distribution—this measure will be, by definition, approximate: the semantics of the specific incidents may be different which will impact the TTM. We need further research to formalize how we derive this estimate, and what to condition it on.

Operators also need *efficient* OCE-helpers—they cannot cost more than the SLA violations the incidents induce. The infrastructure cost to train (or fine-tune) and run inference constitute the *system cost* of the helpers—research [17, 45, 50] which makes ML training and inference more efficient can help reduce this cost. Work on OCE-helpers should report these costs.

Research on OCE-helpers should also report management costs. Each time operators have to research how to adapt to new incidents or maintain a reliable and performant solution they add to this cost. We considered an OCE-helper that automates well-structured incidents that have a clear resolution strategy—in these cases OCEs follow a detailed Troubleshooting Guide (TSG) to mitigate and resolve the incident and the LLM can do the same—but the cost outweighs the benefit: to make this (seemingly simple) solution work we had to integrate LLMs with monitoring APIs, put guard-rails to minimize damages that results from mis-predictions, and even carefully design prompts to make sure the LLM exactly follows the TSG. We can achieve the same goal with a hard-coded Python script (in fact, operators already do)! A change in the TSG necessitates a change in both solutions, and the cost would not amortize.

4 Framework

We present a framework that meets the foundational principles we discussed. It consists of three modules: a hypothesis former, a hypothesis tester, and a mitigation planner. These modules

replicate an OCE’s thought process when resolving an incident. We discuss where LLMs fits within this framework.

4.1 Our Perspective

We next discuss our design perspective to provide context for our design choices. These choices are not necessary (unlike the principles in §2) but are based on our firsthand experiences.

OCE-centric design. IM workflows are too chaotic and intricate to fully automate. We believe the helper should act as a “copilot” and suggest the next steps but keep the OCE in the driver’s seat. The OCE remains responsible for the incident and can decide to use helpers only when they deem it useful. This makes it easier for operators to deploy the solution and helps prevent the negative impacts of the helper on the IM workflow when it inevitably makes mistakes.

Decentralized extensibility. 100+ independent networking teams work on IM—each team has its own categories of incidents, root causes, guidelines, tools, and documentation. It is impractical for all teams to follow the same formula [34] (especially if they have ossified legacy workflows). Instead, we believe in a design where each team can modify how the helper impacts their workflow but not interfere with other teams—we want to enable operators to distributedly manage the helper.

Modular design. We split the helper into independent modules—similar to how, early on, researchers split the networking stack into separate layers (*i.e.*, OSI [18]). A well-designed monolithic system can outperform a modular one—e.g., co-optimized application and transport layers outperform the modularized stack [29]—but a modularized design accelerates early research and helps find solutions quickly.

Asset reuse. OCEs use many tools: device health monitors, packet loss detectors, link/device isolation tools, etc. It took operators many years of research and experience to create them, and we designed our framework to leverage them.

4.2 The Case for LLMs

LLMs have a unique ability to help us achieve a holistic OCE-helper design, and perhaps a fully automated solution:

They can reason (kind of). Resolving incidents (or in general, debugging faults) boils down to inferring the *chain of causal relationships* that lead to the incident (e.g., Figure 2). Operators logically understand the components—e.g., switch failures lead to packet loss—which they use to make these inferences—and backtrack through the possible chain and confirm or reject prior states in the chain to hone in on the correct cause.

We see a similar causal deduction in LLMs [52]: they *parrot* the cause and effect relationships they learned through training on vast corpora of text; e.g., we asked GPT-4 why a VM experienced packet loss, and it provided an exhaustive list including congestion, device failure, misconfigurations, bugs, *etc.* Since OCEs also learn such deduction through training, we can get more detailed responses if we fine-tune LLMs on the same

documents and if we look at the mitigation history of previous incidents—this is in line with prior works’ observations [1].

We can embed causal deductions statistically in a black-box [3, 4, 22], *e.g.*, statistically model the causal link between device failures and packet loss. But such models are usually not tractable (operators make them explicit and narrow to make them tractable): a monolithic black box that models all causal relationships needs to observe all telemetry and predict all possibilities. We have too much telemetry to collect or feed to a black-box and the larger input/output sizes increase sample complexities—they need too many training samples. IM involves rare incidents and is structurally sample-starved.

They are technically well-suited to the problem. Unlike prior ML solutions which have fixed input spaces, we can input anything from normal text, to numbers, to system logs, to loose descriptions, or instructions to an LLM. LLMs also quickly and cheaply instance-optimize, *i.e.*, specialize to a specific instant and problem, with in-context learning and can integrate with external tools [43] and plugins [40].

But LLMs are not magic. Computational complexity grows quadratically with token count, and we have input limits (currently 32K tokens \approx 24K words in GPT-4).² Mainstream LLMs only accept text, and we need clever tricks for multi-modal inputs (*e.g.*, images, graphs, time series). We still can’t use them for long-term planning and reasoning [8, 41]. LLMs are unreliable—they are powerful enough to enable OCE-helpers, but imperfect on their own: they cannot replace the OCE. As researchers, we have to find how best to use them, and we discuss our initial work on this in the next section.

4.3 Overview

The framework we propose has three “LLM agents” (Figure 1): (1) the hypothesis former produces bite-sized hypotheses and describes possible root causes or mitigations in each step; (2) the hypothesis tester takes a hypothesis as input and generates a procedure to verify it; and (3) the mitigation planner creates a detailed mitigation plan. These modules need not use the same techniques: some may use a simple LLM while others more complex goal-driven LLM-agents [47, 55]. Each module presents several suggestions to the OCE who then approves a select few. OCEs can pre-approve certain suggestions that have high confidence and low risk. This strict approval process avoids compounding mistakes—one early irrelevant hypothesis can lead the LLM astray for long durations.

In this design, we generate hypotheses and tests in a loop and eventually produce a mitigation plan (the design shadows the OCE). The end-to-end helper workflow is as follows:

Hypothesis former. The hypothesis former takes in the current context and produces a hypothesis small enough that we can

easily verify—its only goal is to help the operator get one step closer to recognizing the root cause or mitigation.

Take an example incident where a link drops packets: the hypothesis former hypothesizes that either a link is faulty or overloaded, or the monitoring pipeline is broken. The OCE has to choose which hypothesis to test. The LLM must produce a confidence score and an explanation to help novice OCEs decide.

Hypothesis tester. The hypothesis tester receives these choices and generates a plan to validate each: in our earlier example, the tester can query monitors to check the link utilizations and validate whether any link is overloaded. The tester can access the operator’s monitoring infrastructure (*e.g.*, link utilizations), LLM-based tools (*e.g.*, one that checks whether there is an ongoing incident with similar symptoms), or even produce manual steps the OCE can execute (*e.g.*, it can tell the OCE to ask the customer for a packet trace). The role the LLM plays in this step is to identify what tools can test the hypothesis and how to interpret the results from the queries and either accept or reject the hypothesis. The OCE will verify both steps.

Mitigation planner. The OCE then invokes the mitigation planner which takes the *tested* hypothesis as input and produces a mitigation plan that the OCE then can trigger. We explicitly *only* allow the OCE to start the mitigation process because: (1) mitigations can impose drastic costs on cloud tenants, and the responsibility of premature mitigation should fall on the OCE; (2) the hypothesis forming and testing is meant to decrease OCE (not the helper’s) uncertainty. This adds an extra layer of risk protection without impacting TTM.

The mitigation planner generates a list of possible actions (*e.g.* de-isolate link or reroute traffic) and the associated risk. The OCE selects which to use. The OCE has to use various tools to execute each action and then presents the results to the module, which decides whether the mitigation was successful.

Risk assessment is central to the third module. We highlight two important risk assessments in §2: (1) an *internal*, qualitative analysis, which the LLM produces based on thought chains and reason (*e.g.*, if the de-isolated link is faulty, packet loss may increase); (2) an *external*, quantitative analysis, which a white-box algorithm produces based on network principles (*e.g.*, 23% risk that de-isolating link A disrupts a microservice). We find these measures necessary as they are complementary views of risk, and vital to both experienced and novice OCEs.

Design benefits. Teams can independently contribute and develop the system. They can add new tools or remove deprecated ones. OCEs have first-hand experience and know best what tool or mitigation implementation is cost-effective and which tools to update to support new incident patterns.

While these modules work best in tandem, their tasks are separate. Operators benefit from each individual module, and OCEs can fill the gaps for modules in development. This allows a phased execution and lowers the high barrier to deployment.

²Ongoing trends suggest higher token lengths and specialized smaller models, which would benefit our applications.

A modular design eases adaptation. In an end-to-end model, each change requires a full prompt redesign or model fine-tuning. With a modular design, we can surgically upgrade the relevant module. We can adapt LLM-based modules by: (1) fine-tuning [1] which pays an up-front cost; (2) in context learning [13] which is faster to develop but cannot accept tasks with large contexts because of limited prompt size.

4.4 Research Directions

Our community needs to research several new directions to materialize this framework, which are:

Network expertise. We need to evaluate, and if necessary, improve the LLM’s performance on complex systems and network problems (LLMs can diagnose simple problems but not complex ones [47, 55]). To fine-tune for networking knowledge, we can use IM data: (1) OCE training documents, slides, and videos (which are multi-modal, high-quality, and low volume); (2) TSGs (which are high-volume, detailed, and sometimes stale); (3) incident logs and OCE communications (which are high-volume, low-quality and unstructured); (4) direct LLM chatting with OCEs (which are high-quality and expensive); (5) IM tools and their documentation.

Network-focused Embeddings. We have a limited token size per query and have to carefully pick what information is relevant to the current prompt. Frameworks currently embed pieces of text to a vector, and store the tuple in a vector database. At runtime, the vector database runs an approximate nearest-neighbor search on the database and appends top results to the prompt. But embeddings come from generic models trained on non-network specific data [23], and also not trained specifically to guide prompts. Future work can train network-specific embedding models that guide LLM prompts. Recent work may serve as a good starting point [57].

Intelligent planning. LLMs cannot plan long-term [8]. We require the framework to generate and test small hypotheses that are easy to plan and test but also benefit from better planning which requires more study. Methods such as chain-of-thought [28, 52, 54] show promise but rely on intuitive ways to structure the LLM response. There is no barrier to entry for this research, and in fact, the systematic mindset in our community may help us find better solutions. Recent work provides a good starting point for the design [37, 47].

Risk assessment. We propose three lines of research. First, automated qualitative reasoning with LLMs: LLMs can support their conclusions with logical [28, 52] or common-sense [6] arguments if we use them in a multi-step process with “short steps”. But researchers have not yet built a reasoning engine for network and cloud systems. In the short-term, a network expert LLM that understands key network concepts (*e.g.*, routing, wired and wireless media, congestion, RDMA), cloud services (VMs, WANs, Clos topologies, SDN, NICs, Programmable

switches), and can deduce the interaction between these different components, at least in steps can help address this.

Second, we can use reliable and white-box analytical models of risk that apply to specific scenarios. Prior work has looked at analytical risk assessment [2, 38, 53, 58] but they fall short: (1) they don’t measure risk directly (*e.g.*, they measure the increase in link B’s utilization if we isolate link A but the risk operators care about is the impact on the applications that used link A before isolation and link B after); (2) they consider a small set of mitigations compared to the full breadth of what operators can use (*e.g.*, they consider isolating the faulty link but do not model migrating the affected VMs or partially re-routing traffic); and (3) none consider whether the mitigation itself may cause an incident (*e.g.*, whether the mitigation may trigger a race condition that would cause further harm to the network).

Third, and perhaps the most exciting path, is how we can merge these two perspectives on confidence and risk; LLM-based qualitative and high-level assessments, and explicit models with quantified analysis. At a high level, an LLM-based agent decides what is at risk given some mitigation, generates a causal graph of components, and uses analytical tools to quantify the risks, given the causal model.

Toolbox abstraction. The boundary between modules and tools is unclear. There is a trade-off; tools can provide low-level telemetry and let the main module parse the information, or tools could parse low-level telemetry and provide the main module with high-level insights. With the former, we have to verify the LLM-based module parsed the data correctly, and with the latter, we have to manually design the tools.

Verifiable LLM-based tools. LLMs can help with code generation [12] and SQL queries [31, 48]. But we need to verify the outputs they generate if we want to use them in an automated pipeline. This requires research in: (1) formal verification techniques [21] to prove correctness; (2) repairing responses with consistency checks such as Haskell QuickCheck [15].

5 Related Works

LLMs. Prompt-engineering research has surged recently and is necessary to utilize LLMs to their full potential. Broadly, these techniques include: approaches attempt to improve the reasoning quality of LLMs through chain-of-thought and scratch-padding [28, 39, 52, 54], keeping the LLM self-consistent [27, 32, 42, 47, 51], etc. [9, 16, 19, 25, 33, 43, 56].

Incident Management. Multiple work improve and understand the IM workflow either by describing examples from large production environments [24], monitoring solutions to help diagnosis [30, 35, 36, 49], or ML-based solutions to improve different parts of the OCE experience [22].

The work which focuses on using ML to improve the IM workflow itself includes [1, 10, 11, 13, 22, 26, 46]. From this set, those which involve LLMs [1, 13] laid the initial foundation to show there is a promise for LLMs in this space.

References

- [1] Toufique Ahmed, Supriyo Ghosh, Chetan Bansal, Thomas Zimmermann, Xuchao Zhang, and Saravan Rajmohan. 2023. Recommending Root-Cause and Mitigation Steps for Cloud Incidents using Large Language Models. (2023). arXiv:cs.SE/2301.03797 <https://arxiv.org/abs/2301.03797> ICSE'23.
- [2] Omid Alipourfard, Jiaqi Gao, Jeremie Koenig, Chris Harshaw, Amin Vahdat, and Minlan Yu. 2019. Risk-based planning of network changes in evolving data centers. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles*. 414–429.
- [3] Abdullah Alomar, Pouya Hamadian, Arash Nasr-Esfahany, Anish Agarwal, Mohammad Alizadeh, and Devavrat Shah. 2023. CausalSim: A Causal Framework for Unbiased Trace-Driven Simulation. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI '23)*. USENIX Association, Boston, MA, 1115–1147. <https://www.usenix.org/conference/nsdi23/presentation/alomar>
- [4] Behnaz Arzani, Kevin Hsieh, and Haoxian Chen. 2021. Interpret-able feedback for AutoML systems. (2021). arXiv:cs.LG/2102.11267 <https://arxiv.org/abs/2102.11267>
- [5] Microsoft Azure. 2023. Post Incident Review (PIR) – Azure Networking – Global WAN issues. (2023). <https://azure.status.microsoft/en-us/status/history/>, Accessed: 2023-06-26.
- [6] Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. COMET: Commonsense Transformers for Automatic Knowledge Graph Construction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Florence, Italy, 4762–4779. <https://doi.org/10.18653/v1/P19-1470>
- [7] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. (2020). arXiv:cs.CL/2005.14165 <https://arxiv.org/abs/2005.14165>
- [8] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. Sparks of Artificial General Intelligence: Early experiments with GPT-4. (2023). arXiv:cs.CL/2303.12712 <https://arxiv.org/abs/2303.12712>
- [9] Harrison Chase. 2022. LangChain. (Oct. 2022). <https://github.com/hwchase17/langchain> <https://github.com/hwchase17/langchain>, Accessed: 2023-06-26.
- [10] Junjie Chen, Xiaoting He, Qingwei Lin, Hongyu Zhang, Dan Hao, Feng Gao, Zhangwei Xu, Yingnong Dang, and Dongmei Zhang. 2019. Continuous incident triage for large-scale online service systems. In *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 364–375.
- [11] Junjie Chen, Shu Zhang, Xiaoting He, Qingwei Lin, Hongyu Zhang, Dan Hao, Yu Kang, Feng Gao, Zhangwei Xu, Yingnong Dang, et al. 2020. How incidental are the incidents? characterizing and prioritizing incidents for large-scale online service systems. In *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering*. 373–384.
- [12] 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 Hebgren 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. 2021. Evaluating Large Language Models Trained on Code. (2021). arXiv:cs.LG/2107.03374 <https://arxiv.org/abs/2107.03374>
- [13] Yinfang Chen, Huaibing Xie, Minghua Ma, Yu Kang, Xin Gao, Liu Shi, Yunjie Cao, Xuedong Gao, Hao Fan, Ming Wen, Jun Zeng, Supriyo Ghosh, Xuchao Zhang, Chaoyun Zhang, Qingwei Lin, Saravan Rajmohan, and Dongmei Zhang. 2023. Empowering Practical Root Cause Analysis by Large Language Models for Cloud Incidents. (2023). arXiv:cs.SE/2305.15778 <https://arxiv.org/abs/2305.15778>
- [14] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. PaLM: Scaling Language Modeling with Pathways. (2022). arXiv:cs.CL/2204.02311 <https://arxiv.org/abs/2204.02311>
- [15] Koen Claessen and John Hughes. 2000. QuickCheck: a lightweight tool for random testing of Haskell programs. In *Proceedings of the fifth ACM SIGPLAN international conference on Functional programming*. 268–279.
- [16] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training Verifiers to Solve Math Word Problems. (2021). arXiv:cs.LG/2110.14168 <https://arxiv.org/abs/2110.14168>
- [17] Meghan Cowan, Saeed Maleki, Madanlal Musuvathi, Olli Saarikivi, and Yifan Xiong. 2022. MSCCL: Microsoft Collective Communication Library. *arXiv preprint arXiv:2201.11840* (2022).
- [18] John D Day and Hubert Zimmermann. 1983. The OSI reference model. *Proc. IEEE* 71, 12 (1983), 1334–1340.
- [19] Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. 2023. Improving Factuality and Reasoning in Language Models through Multiagent Debate. (2023). arXiv:cs.CL/2305.14325 <https://arxiv.org/abs/2305.14325>
- [20] Marisa Fernandez. 2018. Prime Day woes might have cost Amazon \$72m-\$99m in sales. (2018). <https://www.axios.com/2018/07/18/prime-day-woes-might-have-cost-amazon-from-72-99-million> <https://www.axios.com/2018/07/18/prime-day-woes-might-have-cost-amazon-from-72-99-million>, Accessed: 2023-06-26.
- [21] Emily First, Markus N. Rabe, Talia Ringer, and Yuriy Brun. 2023. Baldu: Whole-Proof Generation and Repair with Large Language Models.

- (2023). arXiv:cs.LG/2303.04910 <https://arxiv.org/abs/2303.04910>
- [22] Jiaqi Gao, Nofel Yaseen, Robert MacDavid, Felipe Vieira Frujeri, Vincent Liu, Ricardo Bianchini, Ramaswamy Aditya, Xiaohang Wang, Henry Lee, David Maltz, Minlan Yu, and Behnaz Arzani. 2020. Scouts: Improving the Diagnosis Process Through Domain-Customized Incident Routing. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM '20)*. Association for Computing Machinery, New York, NY, USA, 253–269. <https://doi.org/10.1145/3387514.3405867>
- [23] Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2022. SimCSE: Simple Contrastive Learning of Sentence Embeddings. (2022). arXiv:cs.CL/2104.08821 <https://arxiv.org/abs/2104.08821>
- [24] Ramesh Govindan, Ina Minei, Mahesh Kallahalla, Bikash Koley, and Amin Vahdat. 2016. Evolve or Die: High-Availability Design Principles Drawn from Google's Network Infrastructure. In *Proceedings of the 2016 ACM SIGCOMM Conference*. <http://dl.acm.org/authorize.cfm?key=N19254>
- [25] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. REALM: Retrieval-Augmented Language Model Pre-Training. (2020). arXiv:cs.CL/2002.08909 <https://arxiv.org/abs/2002.08909>
- [26] Jiajun Jiang, Weihai Lu, Junjie Chen, Qingwei Lin, Pu Zhao, Yu Kang, Hongyu Zhang, Yingfei Xiong, Feng Gao, Zhangwei Xu, et al. 2020. How to mitigate the incident? an effective troubleshooting guide recommendation technique for online service systems. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1410–1420.
- [27] Geunwoo Kim, Pierre Baldi, and Stephen McAleer. 2023. Language Models can Solve Computer Tasks. (2023). arXiv:cs.CL/2303.17491 <https://arxiv.org/abs/2303.17491>
- [28] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2023. Large Language Models are Zero-Shot Reasoners. (2023). arXiv:cs.CL/2205.11916 <https://arxiv.org/abs/2205.11916>
- [29] Adam Langley, Alistair Riddoch, Alyssa Wilk, Antonio Vicente, Charles Krasic, Dan Zhang, Fan Yang, Fedor Kouranov, Ian Swett, Janardhan Iyengar, et al. 2017. The quic transport protocol: Design and internet-scale deployment. In *Proceedings of the conference of the ACM special interest group on data communication*. 183–196.
- [30] Yuliang Li, Rui Miao, Changhoon Kim, and Minlan Yu. 2016. FlowRadar: A Better NetFlow for Data Centers. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*. USENIX Association, Santa Clara, CA, 311–324. <https://www.usenix.org/conference/nsdi16/technical-sessions/presentation/li-yuliang>
- [31] Aiwei Liu, Xuming Hu, Lijie Wen, and Philip S. Yu. 2023. A comprehensive evaluation of ChatGPT's zero-shot Text-to-SQL capability. (2023). arXiv:cs.CL/2303.13547 <https://arxiv.org/abs/2303.13547>
- [32] Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. Self-Refine: Iterative Refinement with Self-Feedback. (2023). arXiv:cs.CL/2303.17651 <https://arxiv.org/abs/2303.17651>
- [33] Microsoft. 2023. LangChain. (Feb. 2023). <https://github.com/microsoft/semantic-kernel> <https://github.com/microsoft/semantic-kernel>, Accessed: 2023-06-26.
- [34] Jeffrey C Mogul, Drago Goricane, Martin Pool, Anees Shaikh, Douglas Turk, Bikash Koley, and Xiaoxue Zhao. 2020. Experiences with Modeling Network Topologies at Multiple Levels of Abstraction.. In *NSDI*. 403–418.
- [35] Masoud Moshref, Minlan Yu, Ramesh Govindan, and Amin Vahdat. 2014. DREAM: Dynamic Resource Allocation for Software-defined Measurement. In *Proceedings of the ACM SIGCOMM Conference*.
- [36] Masoud Moshref, Minlan Yu, Ramesh Govindan, and Amin Vahdat. 2016. Trumpet: Timely and Precise Triggers in Data Centers. In *ACM SIGCOMM*.
- [37] Yohei Nakajima. 2023. Task-driven Autonomous Agent. (April 2023). <https://github.com/yoheinakajima/babyagi> <https://github.com/yoheinakajima/babyagi>, Accessed: 2023-06-28.
- [38] Pooria Namyar, Behnaz Arzani, Daniel Crankshaw, Daniel S. Berger, Kevin Hsieh, Srikanth Kandula, and Ramesh Govindan. 2023. Mitigating the Performance Impact of Network Failures in Public Clouds. (2023). arXiv:cs.NI/2305.13792 <https://arxiv.org/abs/2305.13792>
- [39] Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and Augustus Odena. 2021. Show Your Work: Scratchpads for Intermediate Computation with Language Models. (2021). arXiv:cs.LG/2112.00114 <https://arxiv.org/abs/2112.00114>
- [40] OpenAI. 2023. ChatGPT plugins. (2023). <https://openai.com/blog/chatgpt-plugins> <https://openai.com/blog/chatgpt-plugins>, Accessed: 2023-06-26.
- [41] OpenAI. 2023. GPT-4 Technical Report. (2023). arXiv:cs.CL/2303.08774 <https://arxiv.org/abs/2303.08774>
- [42] Debjit Paul, Mete Ismayilzada, Maxime Peyrard, Beatriz Borges, Antoine Bosselut, Robert West, and Boi Faltings. 2023. REFINER: Reasoning Feedback on Intermediate Representations. (2023). arXiv:cs.CL/2304.01904 <https://arxiv.org/abs/2304.01904>
- [43] Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language Models Can Teach Themselves to Use Tools. (2023). arXiv:cs.CL/2302.04761 <https://arxiv.org/abs/2302.04761>
- [44] Amazon Web Services. 2021. Summary of AWS Direct Connect Event in the Tokyo (AP-NORTHEAST-1) Region. (2021). <https://aws.amazon.com/message/17908/>
- [45] Aashaka Shah, Vijay Chidambaram, Meghan Cowan, Saeed Maleki, Madan Musuvathi, Todd Mytkowicz, Jacob Nelson, and Olli Saarikivi. 2023. TACCL: Guiding Collective Algorithm Synthesis using Communication Sketches. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. 593–612.
- [46] Manish Shetty, Chetan Bansal, Sai Pramod Upadhyayula, Arjun Radhakrishna, and Anurag Gupta. 2022. AutoTSG: Learning and Synthesis for Incident Troubleshooting. (2022). arXiv:cs.SE/2205.13457 <https://arxiv.org/abs/2205.13457>
- [47] Noah Shinn, Federico Cassano, Beck Labash, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language Agents with Verbal Reinforcement Learning. (2023). arXiv:cs.AI/2303.11366 <https://arxiv.org/abs/2303.11366>
- [48] Ruoxi Sun, Sercan O. Arik, Hootan Nakhost, Hanjun Dai, Rajarishi Sinha, Pengcheng Yin, and Tomas Pfister. 2023. SQL-PaLM: Improved Large Language Model Adaptation for Text-to-SQL. (2023). arXiv:cs.CL/2306.00739 <https://arxiv.org/abs/2306.00739>
- [49] Cheng Tan, Ze Jin, Chuanxiong Guo, Tianrong Zhang, Haitao Wu, Karl Deng, Dongming Bi, and Dong Xiang. 2019. NetBouncer: Active Device and Link Failure Localization in Data Center Networks. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*. USENIX Association, Boston, MA, 599–614. <https://www.usenix.org/conference/nsdi19/presentation/tan>
- [50] Weiyang Wang, Moein Khazraee, Zhizhen Zhong, Manya Ghobadi, Zhihao Jia, Dheevatsa Mudigere, Ying Zhang, and Anthony Kewitsch. 2023. TopoOpt: Co-optimizing Network Topology and Parallelization Strategy for Distributed Training Jobs. In *20th USENIX Symposium on*

- Networked Systems Design and Implementation (NSDI 23)*. 739–767.
- [51] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. Self-Consistency Improves Chain of Thought Reasoning in Language Models. (2023). arXiv:cs.CL/2203.11171 <https://arxiv.org/abs/2203.11171>
- [52] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. (2023). arXiv:cs.CL/2201.11903 <https://arxiv.org/abs/2201.11903>
- [53] Xin Wu, Daniel Turner, Chao-Chih Chen, David A Maltz, Xiaowei Yang, Lihua Yuan, and Ming Zhang. 2012. NetPilot: Automating datacenter network failure mitigation. In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*. 419–430.
- [54] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. (2023). arXiv:cs.CL/2305.10601 <https://arxiv.org/abs/2305.10601>
- [55] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. ReAct: Synergizing Reasoning and Acting in Language Models. *arXiv preprint arXiv:2210.03629* (2022).
- [56] Andy Zeng, Maria Attarian, Brian Ichter, Krzysztof Choromanski, Adrian Wong, Stefan Welker, Federico Tombari, Aavek Purohit, Michael Ryoo, Vikas Sindhwani, Johnny Lee, Vincent Vanhoucke, and Pete Florence. 2022. Socratic Models: Composing Zero-Shot Multimodal Reasoning with Language. (2022). arXiv:cs.CV/2204.00598 <https://arxiv.org/abs/2204.00598>
- [57] Shenglin Zhang, Pengxiang Jin, Zihan Lin, Yongqian Sun, Bicheng Zhang, Sibao Xia, Zhengdan Li, Zhenyu Zhong, Minghua Ma, Wa Jin, Dai Zhang, Zhenyu Zhu, and Dan Pei. 2023. Robust Failure Diagnosis of Microservice System through Multimodal Data. (2023). arXiv:cs.SE/2302.10512 <https://arxiv.org/abs/2302.10512>
- [58] Danyang Zhuo, Monia Ghobadi, Ratul Mahajan, Klaus-Tycho Förster, Arvind Krishnamurthy, and Thomas Anderson. 2017. Understanding and mitigating packet corruption in data center networks. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*. 362–375.