Talking to routers with LLMs

- Jonathan Saha

jonathan.saha@outlook.com

What I wanted

- I will ask, "Hey AI, what is wrong with my network?"
- LLMs will ssh into my network devices and check everything
- LLM will give me a summarized and human readable answer.

What does the workflow look like

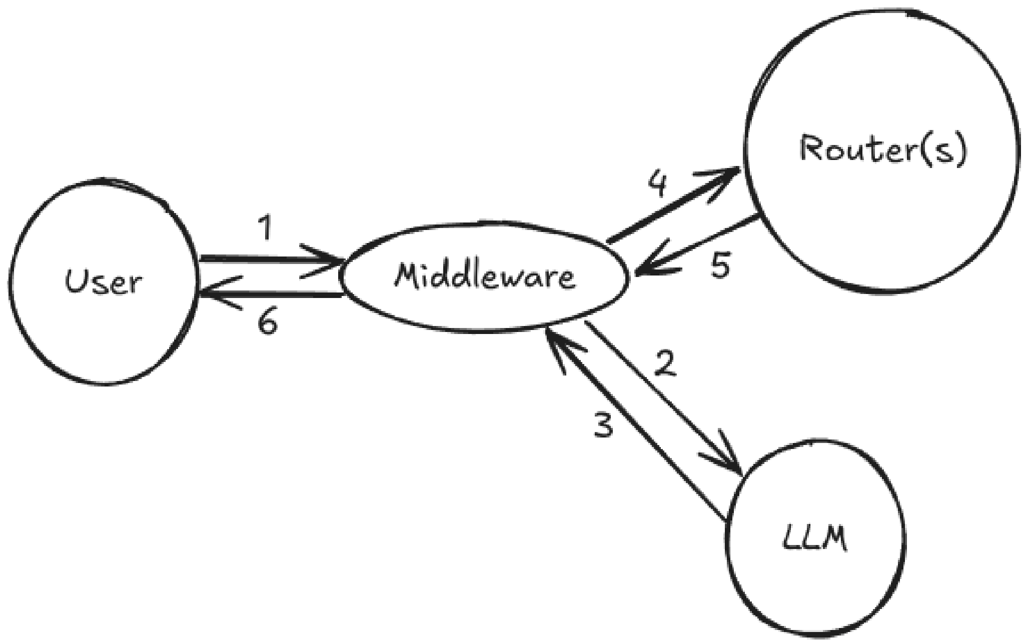

1. User asks a question.
2. The Middleware sends it to an LLM.
3. LLM returns an appropriate command to middleware.
4. Middleware sends that command to router.
5. Router gives feedback.
6. User gets the feedback.

The setup

Software:

- Containerlab
  - SR Linux
- Ollama
- Deepseek
  - R1:7b
- OpenWebUI, optional
- Middleware in python

Hardware:

- AMD 1700X (8c, 16t)
- 16 RAM
- Nvidia 1060

Steps

- Set up a two-router lab in Container lab, make sure devices are ssh enabled.
- Set up Ollama and Deepseek.
- Make a middleware in Python that can ssh into the routers and talk with the LLM via api.

Challenges I faced and their solutions

- Poor coding skill.
  - Vibe coding.
- Dumb AI, or my initial plan was too complex than I realized.
  - Later I opted for Gemini 2.
- LLM often returned improperly formatted responses.
  - Extensive filtering needed in middleware. Moving to Gemini helped.
- Giving LLM the context or training the AI.
  - Had to make my own dataset to. Used AI for that too!

Result

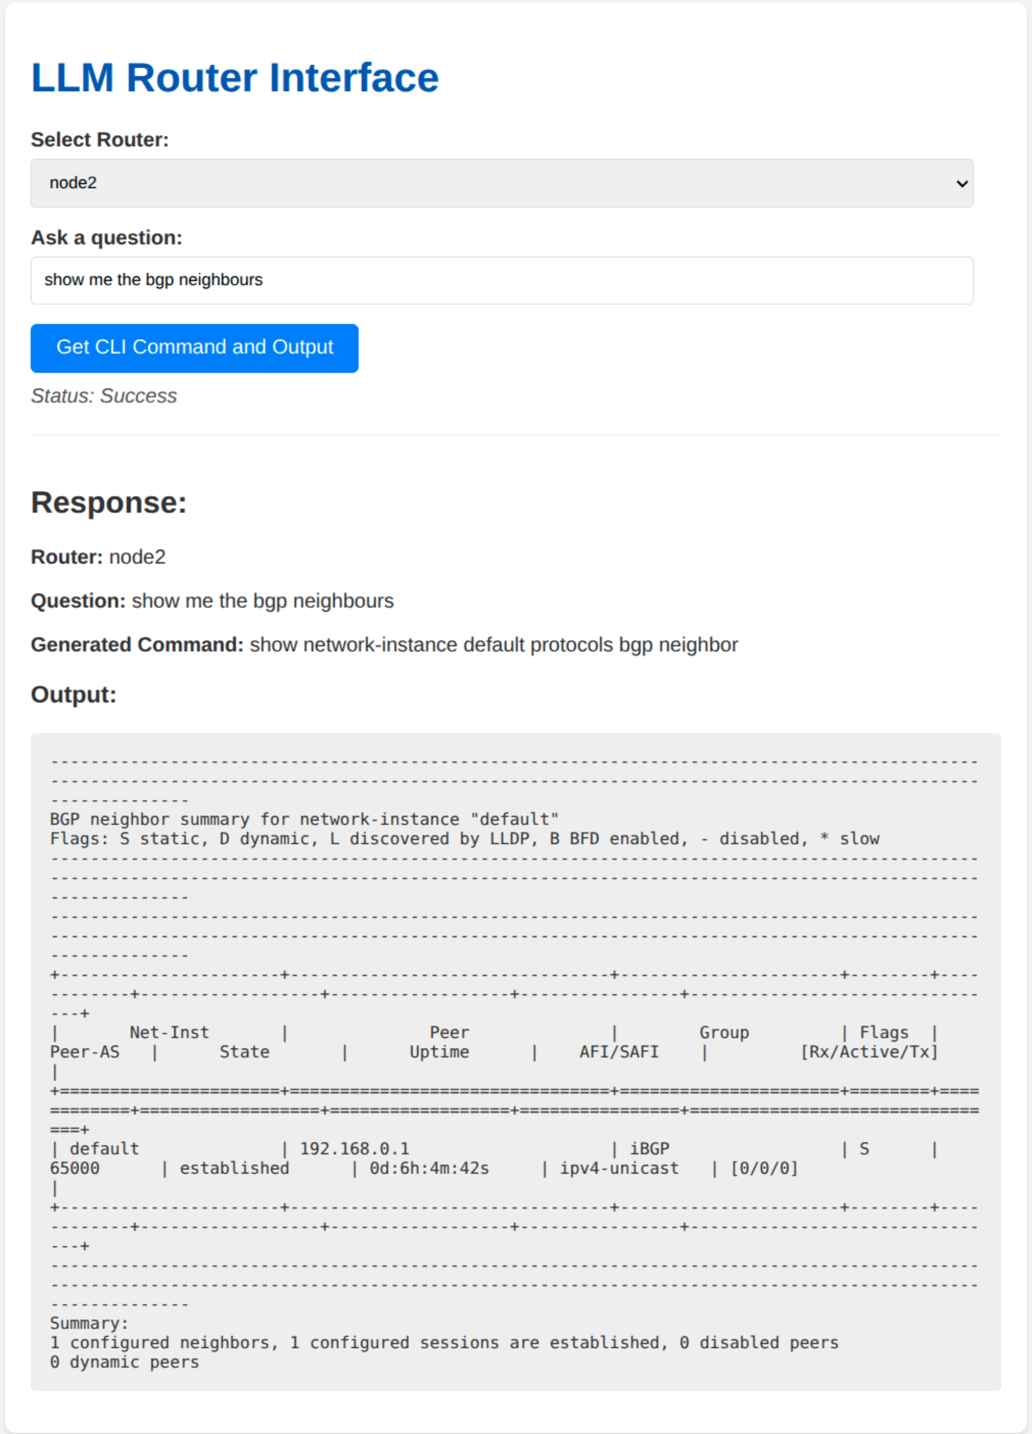

LLM Router Interface

Select Router:

node2

Ask a question:

show me the bgp neighbours

Get CLI Command and Output

Status: Success

Response:

Router: node2

Question: show me the bgp neighbours

Generated Command: show network-instance default protocols bgp neighbor

Output:

```
---------------------------------------------------------------------------------------------
---------------------------------------------------------------------------------------------
--------------
BGP neighbor summary for network-instance "default"
Flags: S static, D dynamic, L discovered by LLDP, B BFD enabled, - disabled, * slow
---------------------------------------------------------------------------------------------
---------------------------------------------------------------------------------------------
--------------
---------------------------------------------------------------------------------------------
---------------------------------------------------------------------------------------------
--------------
+----------------------+--------------------------------+----------------------+--------+----
--------+------------------+------------------+----------------+-----------------------------
---+
|       Net-Inst       |              Peer              |        Group         | Flags  |
Peer-AS   |      State       |      Uptime      |    AFI/SAFI    |         [Rx/Active/Tx]
|
+======================+================================+======================+========+====
========+==================+==================+================+=============================
===+
| default              | 192.168.0.1                    | iBGP                 | S      |
65000      | established      | 0d:6h:4m:42s     | ipv4-unicast   | [0/0/0]
|
+----------------------+--------------------------------+----------------------+--------+----
--------+------------------+------------------+----------------+-----------------------------
---+
---------------------------------------------------------------------------------------------
---------------------------------------------------------------------------------------------
--------------
Summary:
1 configured neighbors, 1 configured sessions are established, 0 disabled peers
0 dynamic peers
```

Result

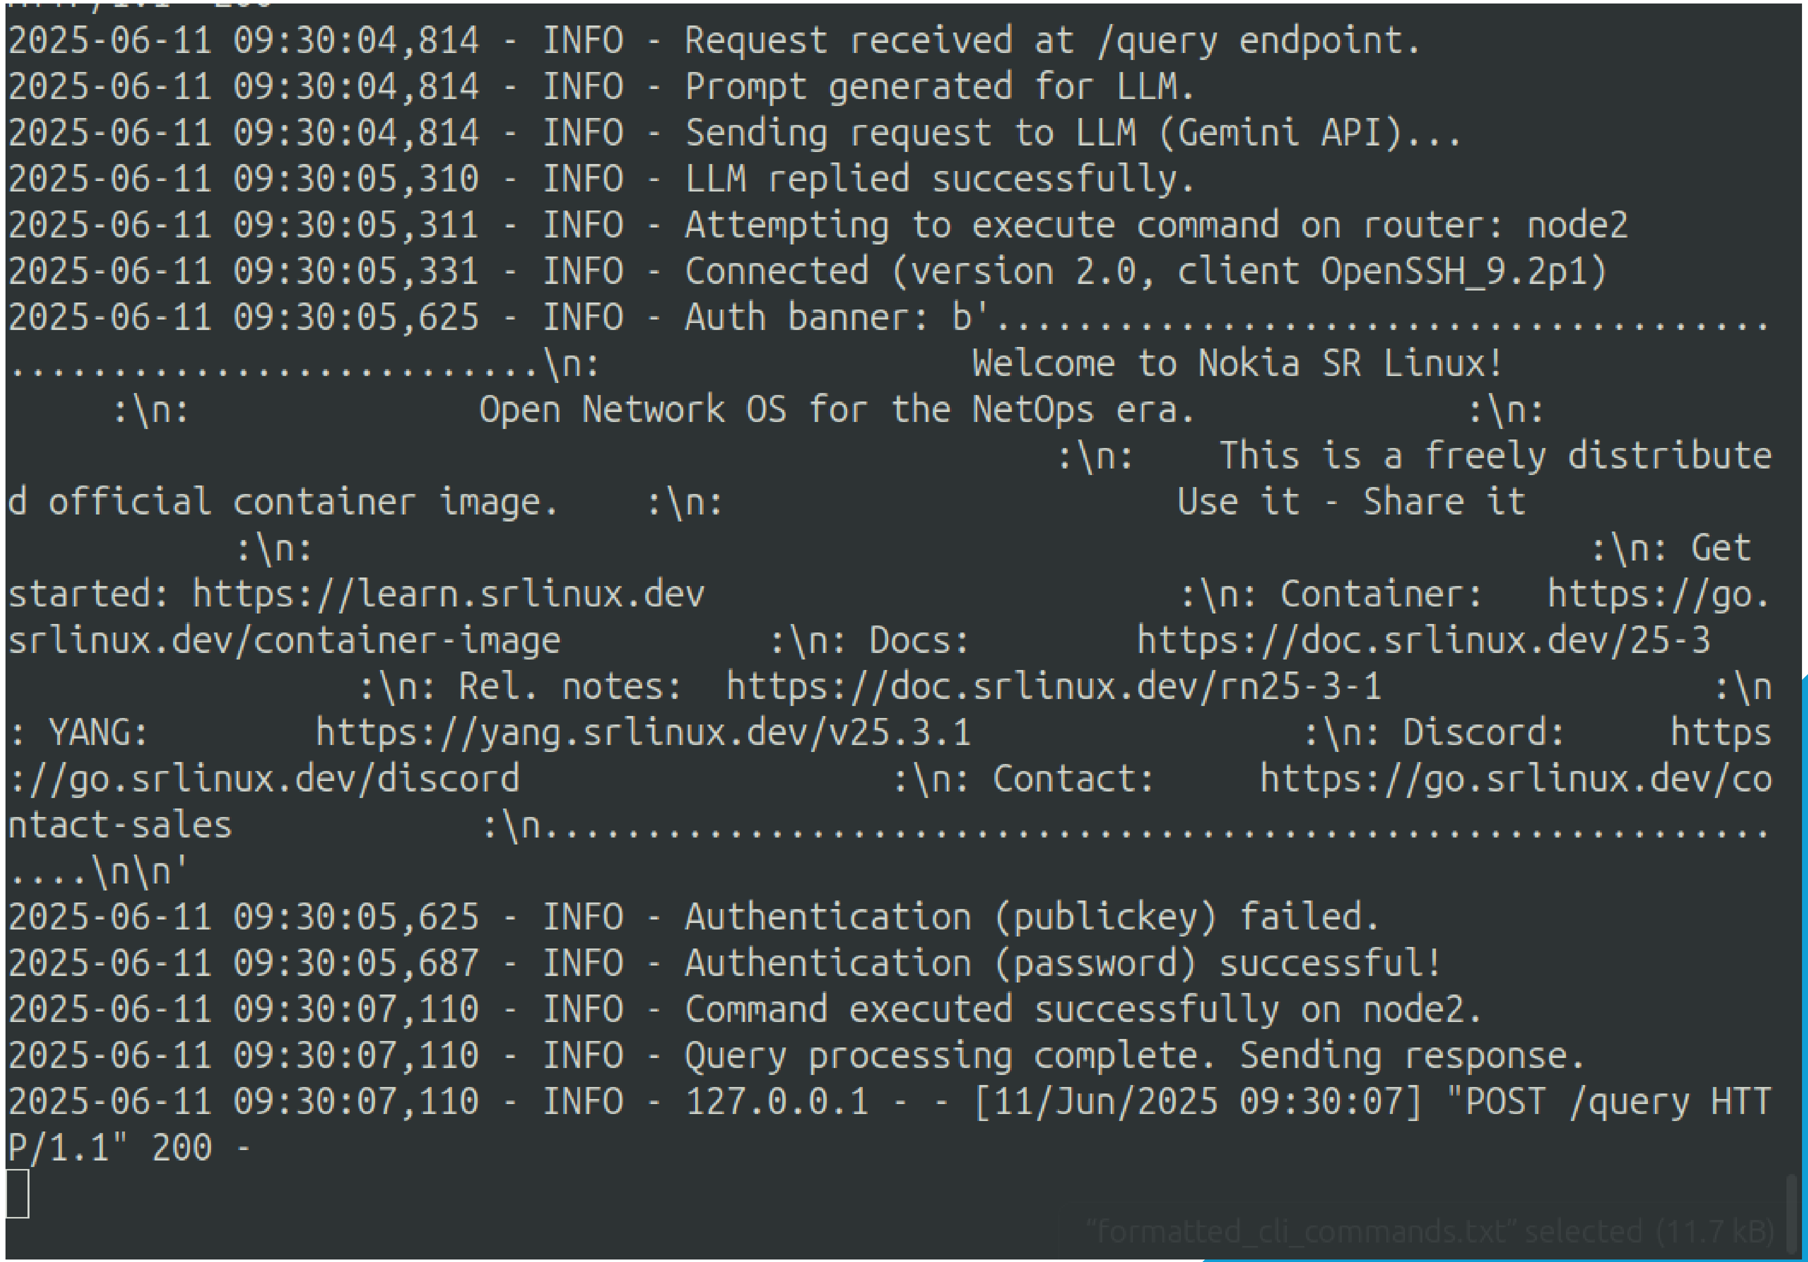

```
2025-06-11 09:30:04,814 - INFO - Request received at /query endpoint.
2025-06-11 09:30:04,814 - INFO - Prompt generated for LLM.
2025-06-11 09:30:04,814 - INFO - Sending request to LLM (Gemini API)...
2025-06-11 09:30:05,310 - INFO - LLM replied successfully.
2025-06-11 09:30:05,311 - INFO - Attempting to execute command on router: node2
2025-06-11 09:30:05,331 - INFO - Connected (version 2.0, client OpenSSH_9.2p1)
2025-06-11 09:30:05,625 - INFO - Auth banner: b'.....................................
...........................\n:                  Welcome to Nokia SR Linux!
     :\n:              Open Network OS for the NetOps era.              :\n:
                                                    :\n:    This is a freely distribute
d official container image.    :\n:                     Use it - Share it
           :\n:                                                        :\n: Get
started: https://learn.srlinux.dev                       :\n: Container:   https://go.
srlinux.dev/container-image          :\n: Docs:       https://doc.srlinux.dev/25-3
                 :\n: Rel. notes:  https://doc.srlinux.dev/rn25-3-1                :\n
: YANG:        https://yang.srlinux.dev/v25.3.1                :\n: Discord:     https
://go.srlinux.dev/discord                  :\n: Contact:     https://go.srlinux.dev/co
ntact-sales            :\n.............................................................
....\n\n'
2025-06-11 09:30:05,625 - INFO - Authentication (publickey) failed.
2025-06-11 09:30:05,687 - INFO - Authentication (password) successful!
2025-06-11 09:30:07,110 - INFO - Command executed successfully on node2.
2025-06-11 09:30:07,110 - INFO - Query processing complete. Sending response.
2025-06-11 09:30:07,110 - INFO - 127.0.0.1 - - [11/Jun/2025 09:30:07] "POST /query HTT
P/1.1" 200 -
```

"formatted_cli_commands.txt" selected (11.7 kB)

Result

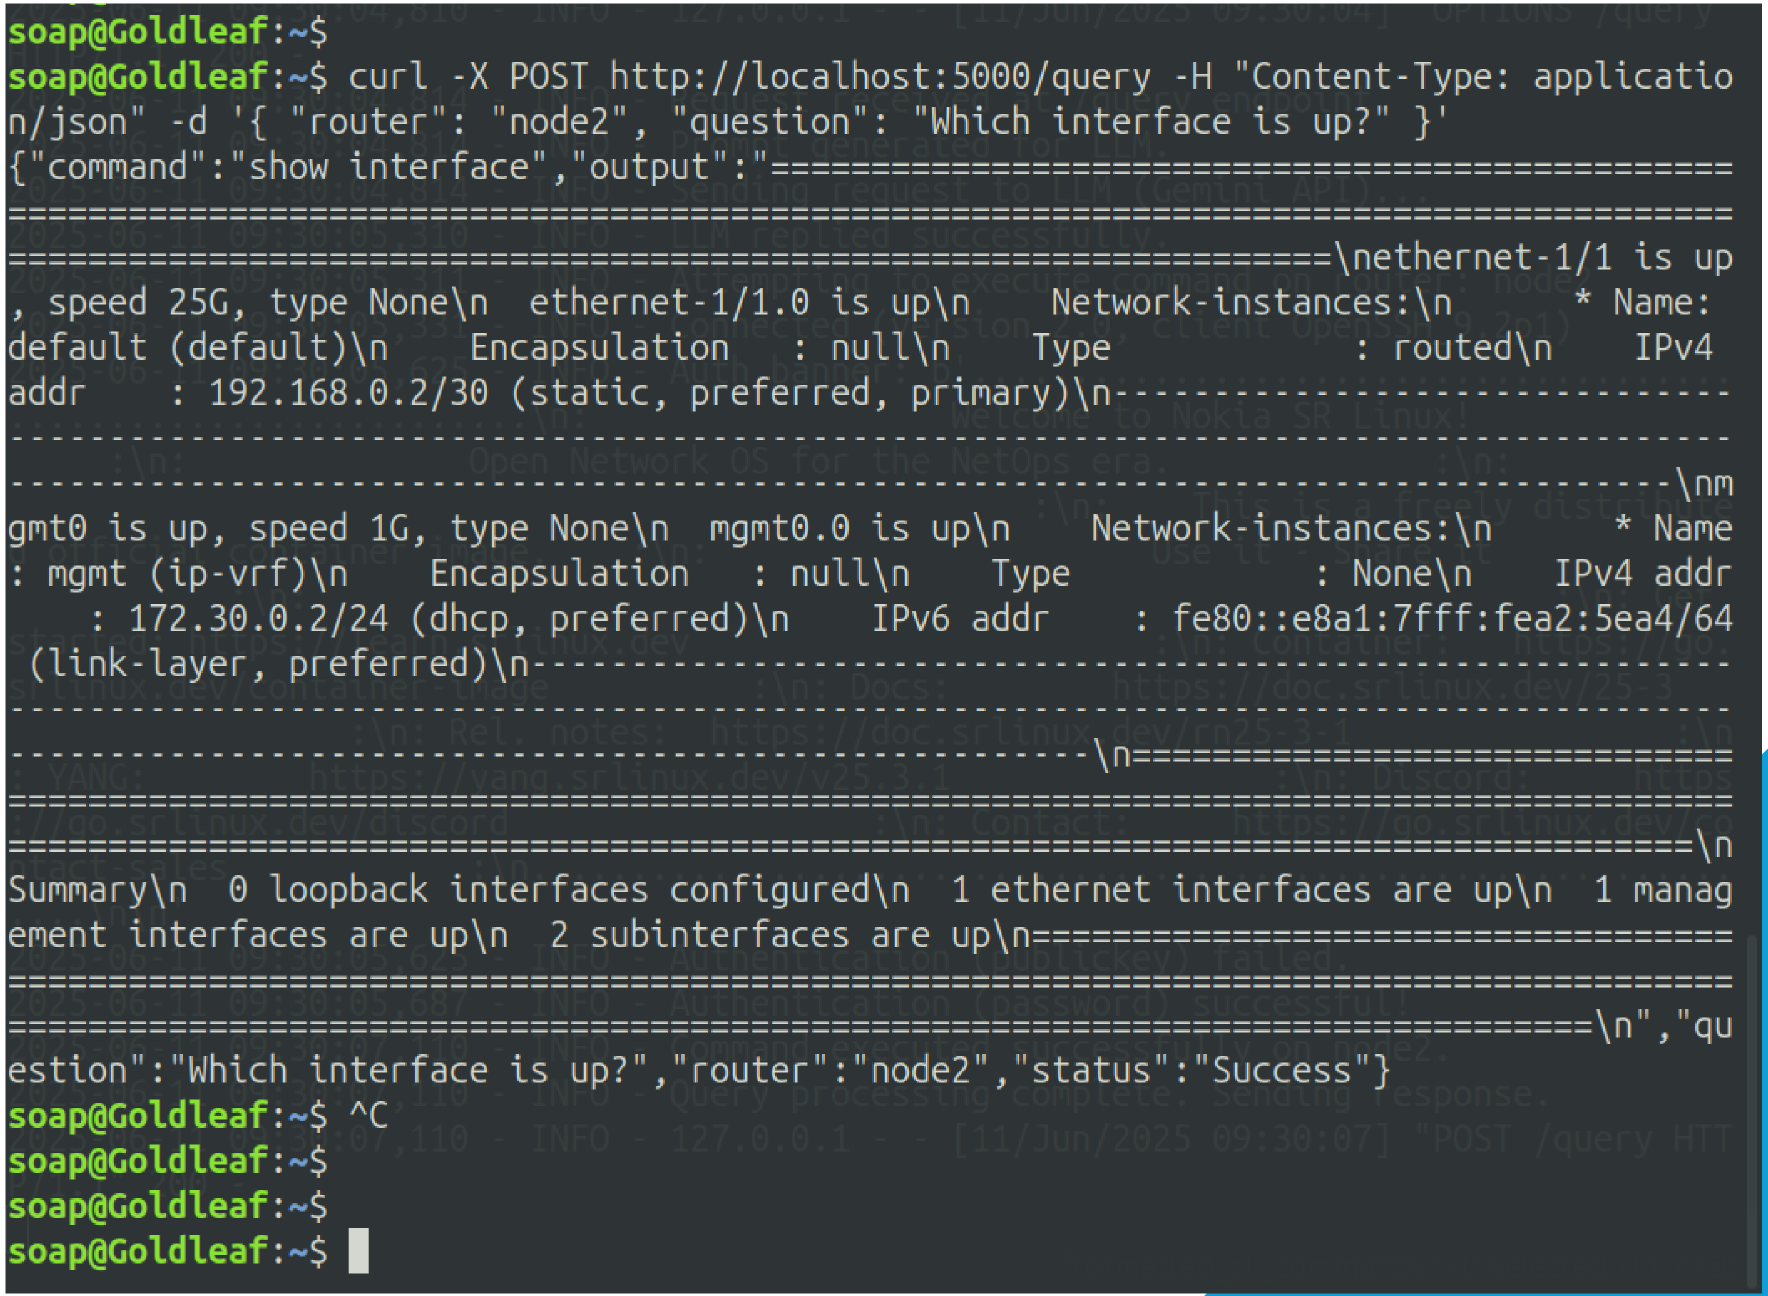

```
soap@Goldleaf:~$
soap@Goldleaf:~$ curl -X POST http://localhost:5000/query -H "Content-Type: applicatio
n/json" -d '{ "router": "node2", "question": "Which interface is up?" }'
{"command":"show interface","output":"=========================================
===============================================================================
=================================================================\nethernet-1/1 is up
, speed 25G, type None\n  ethernet-1/1.0 is up\n    Network-instances:\n      * Name:
default (default)\n    Encapsulation   : null\n    Type              : routed\n    IPv4
addr    : 192.168.0.2/30 (static, preferred, primary)\n--------------------------------
-------------------------------------------------------------------------------
----------------------------------------------------------------------------\nm
gmt0 is up, speed 1G, type None\n  mgmt0.0 is up\n    Network-instances:\n      * Name
: mgmt (ip-vrf)\n    Encapsulation   : null\n    Type              : None\n    IPv4 addr
    : 172.30.0.2/24 (dhcp, preferred)\n    IPv6 addr    : fe80::e8a1:7fff:fea2:5ea4/64
 (link-layer, preferred)\n-----------------------------------------------------
-------------------------------------------------------------------------------
-------------------------------------------------------\n=========================
===============================================================================
===============================================================================\n
Summary\n  0 loopback interfaces configured\n  1 ethernet interfaces are up\n  1 manag
ement interfaces are up\n  2 subinterfaces are up\n===================================
===============================================================================
==============================================================================\n","qu
estion":"Which interface is up?","router":"node2","status":"Success"}
soap@Goldleaf:~$ ^C
soap@Goldleaf:~$
soap@Goldleaf:~$
soap@Goldleaf:~$
```

What I learnt

- Vibe coding is viable, but only for prototypes.
- For prototype, locally hosted LLMs are very underpowered for this task. Unless you have a few top-of-the-line GPUs lying around.
- For production, it is a different story.
- Giving the LLM context is key, even in big ones.
- Open source and free(as in free beer) tools are ready to take on this challenge.

Future

- I do believe this can be a good tool for network engineers, especially in large scale.
- I want to work on it further, and I am asking help from community.
  - Do you see any major faults in the idea?
  - Do you see yourself using this kind of system?
  - Do you have any thoughts or concerns?
- Any feedback is most welcome.

Thank you for your time

What I have done till now will be available on:
https://github.com/jonathansaha/llm-router-interface