



Welcome to the era of LLMs in <Networking>

Javier Antich

javier@anetta.ai

<https://www.linkedin.com/in/javier-antich-romaguera/>

ESNOG – Oct 26th 2023

Goals for this session

- Understand what a Language Model is.
- Understand the implications and applications of Large Language Models in the Networking Industry.
- Review some practical use cases.
- Introduce new architectures: LLM-powered autonomous agents.
- What's in it for you?



Setting the Industry context



Network
infrastructures are
more complex than
ever

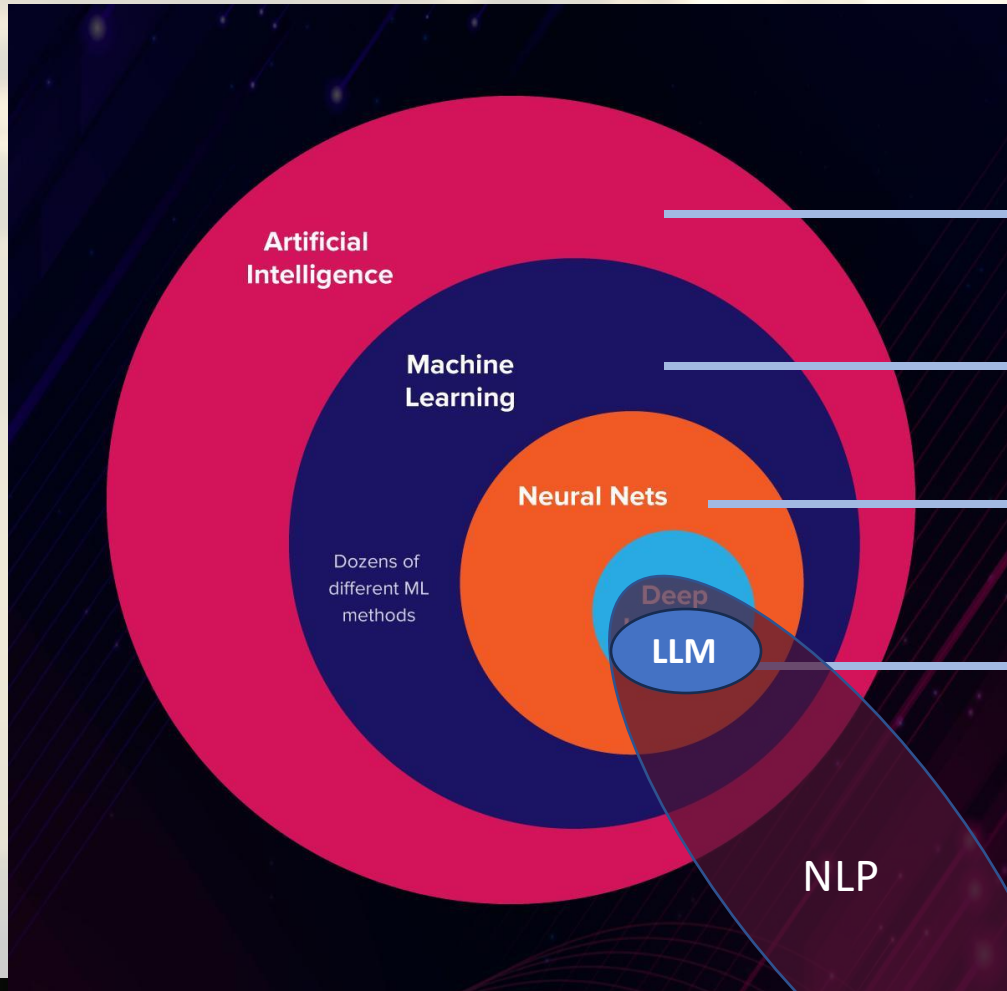


Qualified resources
are more scarce
than ever



Telecom
infrastructures are
more mission
critical than ever

setting the context and terminology



This is about making machines imitate human behavior

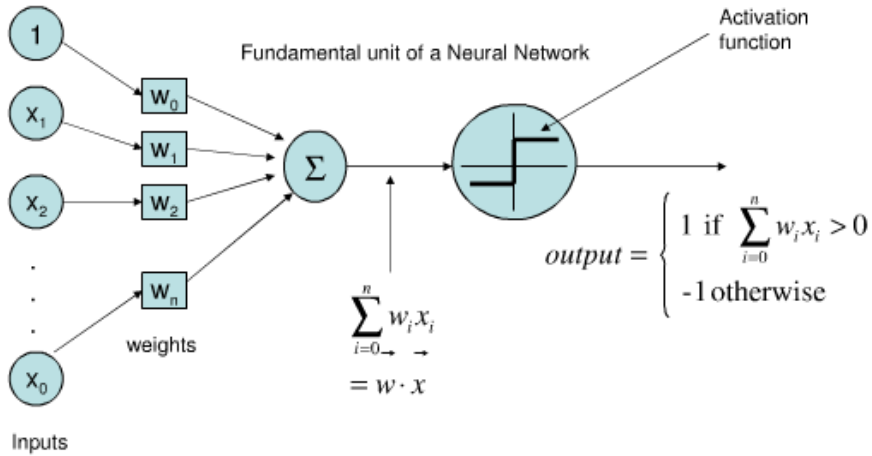
This is about doing it with algorithms that learn from data

This is about using Neural Networks

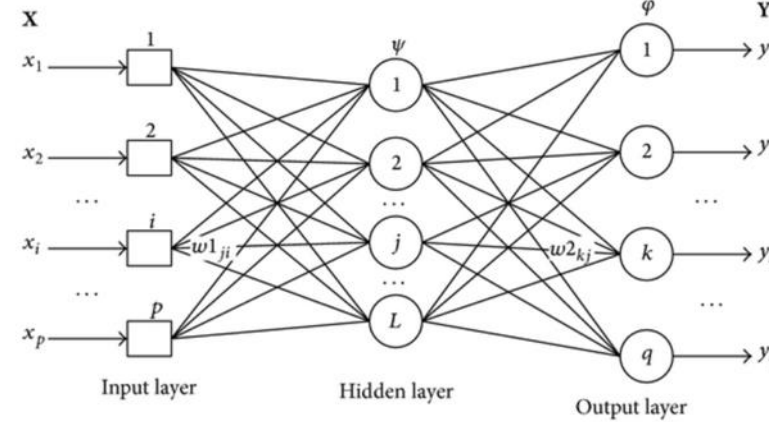
This is about using Neural Networks with many layers

Neural networks. the ultimate function approximation

perceptron



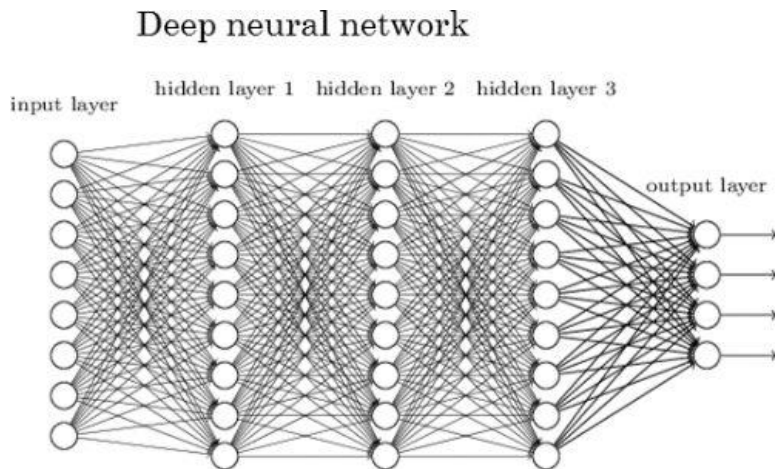
multilayer perceptron



some data points

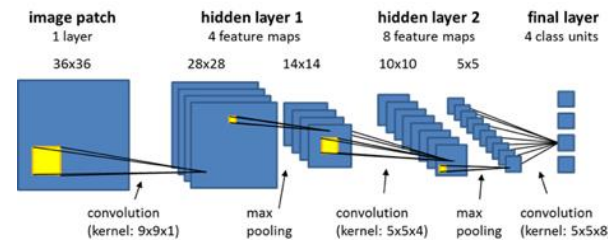
- Require A LOT of data
- May take a long time to train.
- Consume A LOT of resources.
- Largest model OpenAI GPT-4
 - 1.76 trillion parameters.
- Cost of training:
 - >\$100M

deep neural network

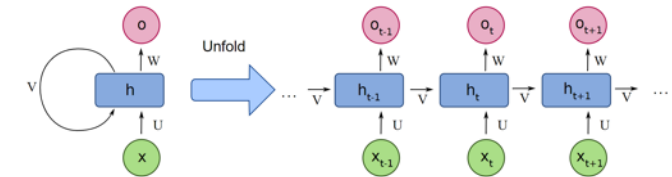


other types of neural networks

Convolutional Neural Networks (CNN)



Recurrent Neural Networks (RNN)



Autoencoders, Attention RNN, Graph NN, GAN, **Transformers**, ...

The background of the slide is a soft, out-of-focus photograph. It depicts a landscape during the 'golden hour' of a sunset or sunrise. The sky is filled with warm, yellow and orange light, with some light-colored clouds. In the foreground, there are silhouettes of rolling hills or mountains. On the right side, a single, tall tree with a curved trunk and a dense canopy of leaves stands out against the bright sky.

Language Models

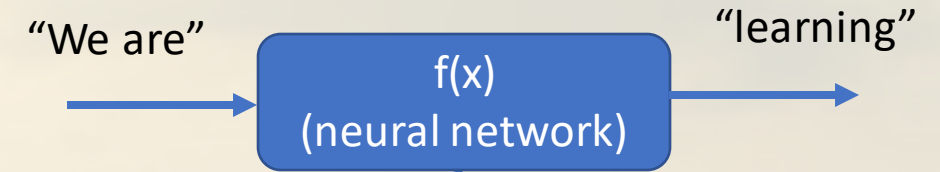
What is auto-regression?

Given some type of data, build a model that can predict the next values based on the previous values. Used widely in Time series forecasting (ARIMA, etc.). The model “learns” how the data behaves.

some data sequence

0100001000111101010001110001011010101100101010110...
2515542393659230919549603818103959683901939584684...
Hello World, how are you doing. We are learning about LLMs...

Training data	Target
Hello World	how
World, how	are
how are	you
are you	doing
...	...

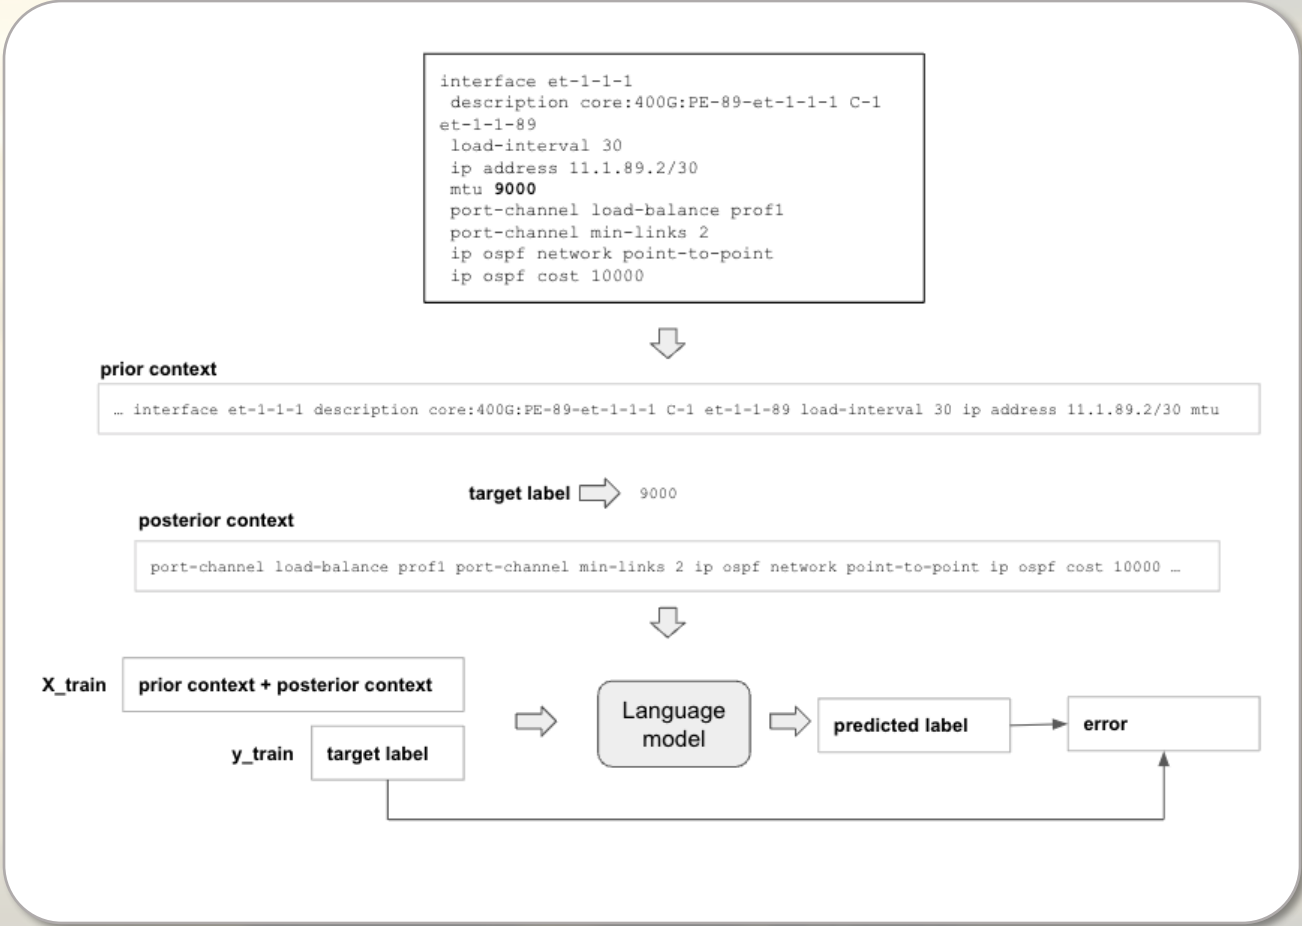
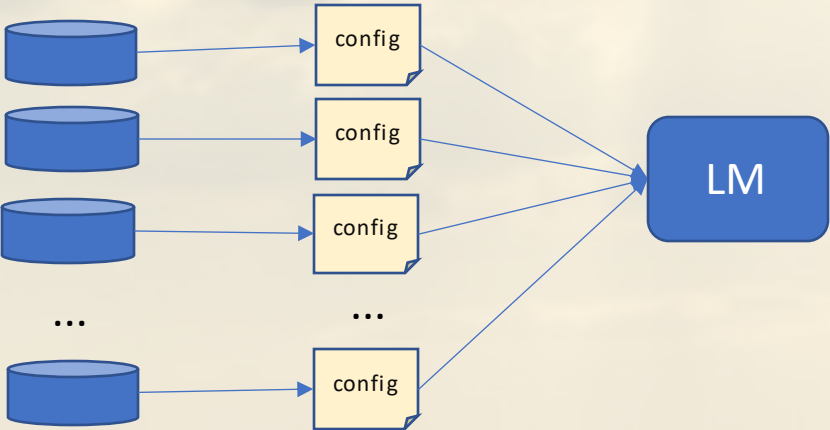


function $f(x)$

When applied to text, an auto-regressive model “learns” how the natural language “works” -> **Language Model**

Small LM Use case: Configuration compliance verification*

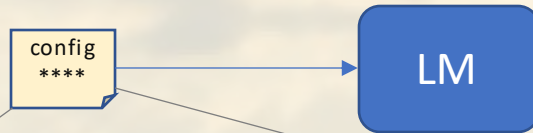
Model training:



* Use case #10 in the book

Small LM Use case: Configuration compliance verification*

Inference:



```
interface et-1-1-2
description core:400G:PE-100-et-1-1-2 C-2 et-1-1-100
load-interval 30
ip address 11.2.100.2/30
    mtu 9001
port-channel load-balance prof1
port-channel min-links 2
ip ospf network point-to-point
ip ospf cost 10000
!
```

```
interface et-1-1-2
description core:400G:PE-100-et-1-1-2 C-2 et-1-1-100
load-interval 30
ip address 11.2.100.2/30
    mtu 9001(expected: 9000)
port-channel load-balance prof1
port-channel min-links 2
ip ospf network point-to-point
ip ospf cost 10000
!
```

* Use case #10 in the book

Large Language Models (LLMs)

The emergence of Large Language Models

Language Model

- Self-supervised deep neural network (usually LSTM, Transformer)
- Trained to predict the most likely “next word” in a sentence.

How is this relevant/applicable to networking?



multiple sources of “natural language” in networks and clouds



- configuration files
- log messages
- customer support tickets
- alerts & incidents
- snmp traps
- labels
- manifests / workflows
- REGEX

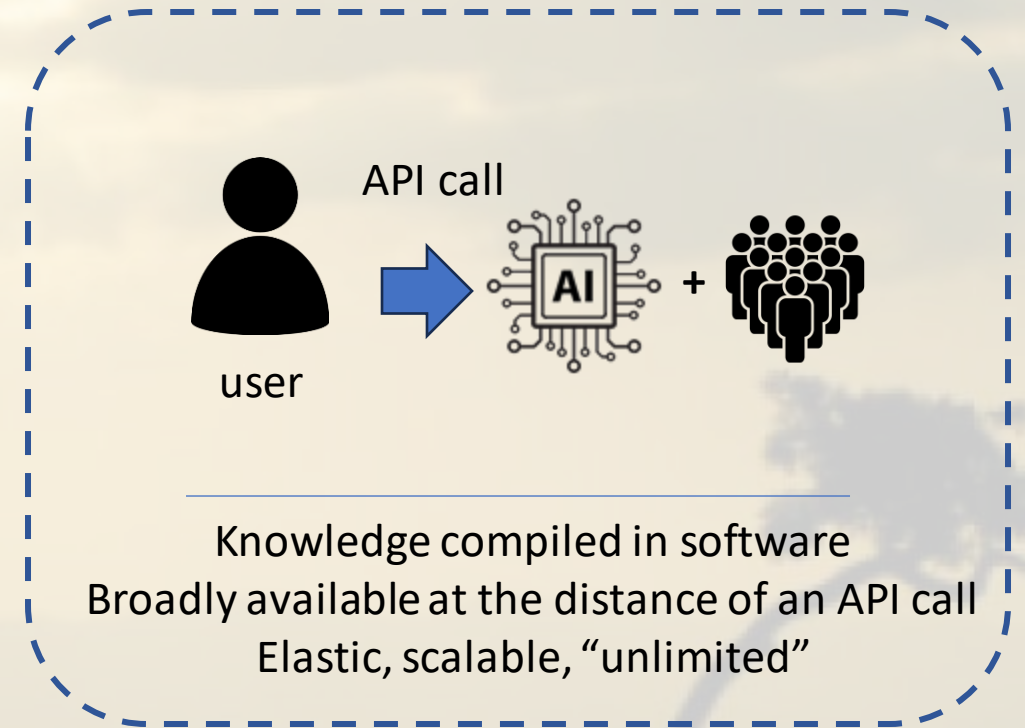
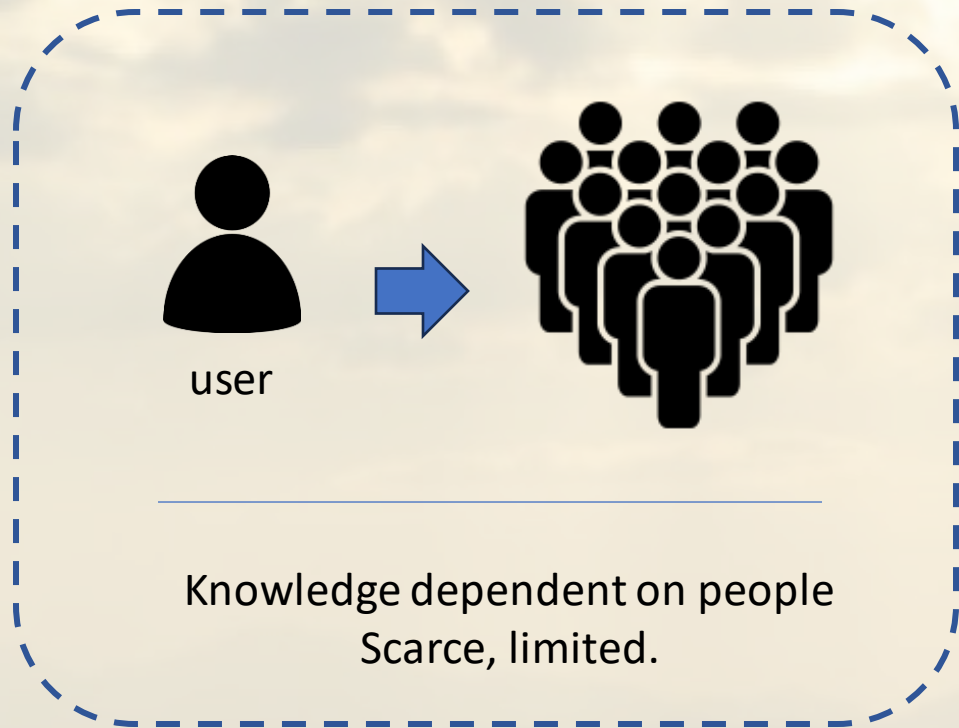
Large Language Model PaLM (BARD), GPT-4, ..

- Multi-task, zero-shot, one-shot, few-shot classification, summarization, generation

The background of the slide is a soft-focus landscape. The top half shows a bright, hazy sky with light clouds, suggesting a sunrise or sunset. The bottom half shows a silhouette of a tree on the right side, with its branches extending towards the center. The overall color palette is warm and muted, with shades of yellow, orange, and light blue.

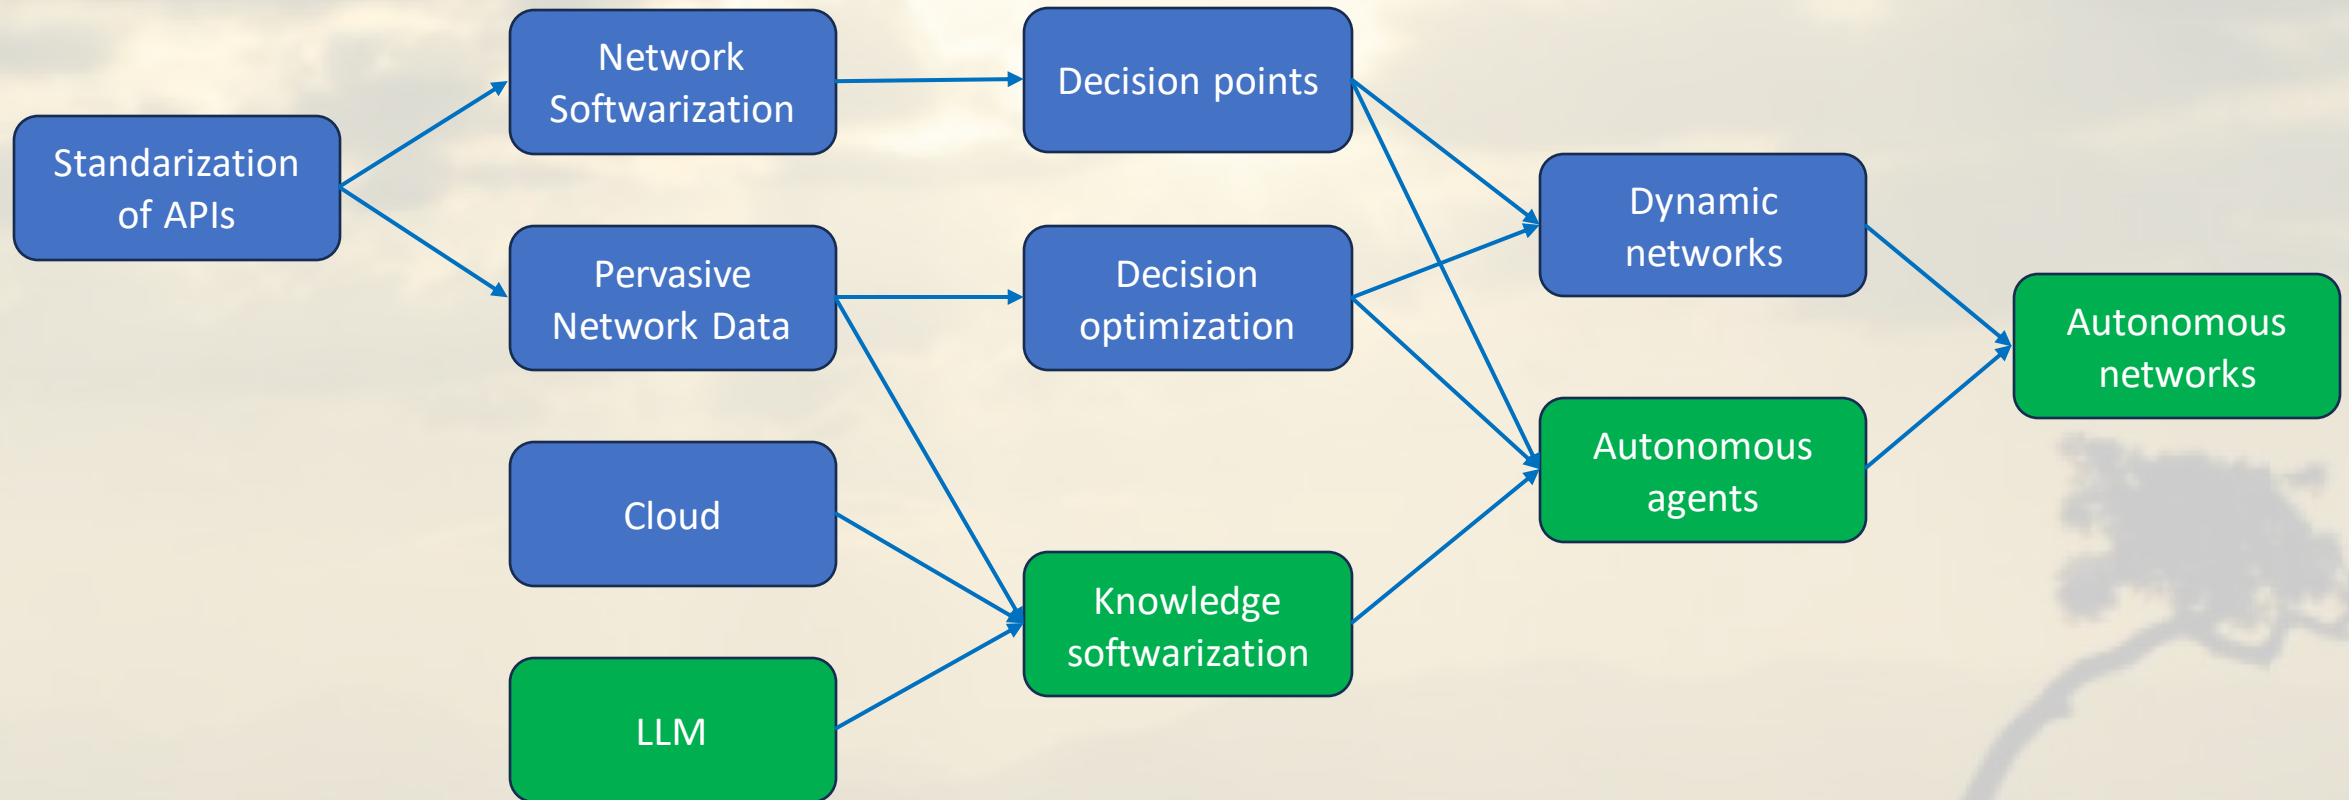
The softwarization of Knowledge

The big disruption



The LLMs are huge “knowledge compilers”

Network softwarization and intelligence






LLMs & Networks. Practical use cases

LLM Use case: Log data extraction

Here is a Syslog message: *Feb 14 09:40:09.325: %LINK-3-UPDOWN: Interface GigabitEthernet0/1, changed state to up.
I need a **regular expression** to extract the **interface name**.



Sure, here is a regular expression to extract the interface name from the Syslog message:

```
^%LINK-3-UPDOWN: Interface (.*) , changed state to up$
```


This regular expression will match the following pattern:



```
^%LINK-3-UPDOWN: Interface (.*) , changed state to up$
```

where *()* is a group that matches any character except a newline. The interface name will be captured in the group *()*.

Here is an example of how to use this regular expression to extract the interface name from the Syslog message:

```
>>> import re
>>> syslog_message = """
*Feb 14 09:40:09.325: %LINK-3-UPDOWN: Interface GigabitEthernet0/1, changed state to up.
"""
>>> interface_name = re.search(r'^%LINK-3-UPDOWN: Interface (.*) , changed state to up$', syslog_message).group(1)
>>> interface_name
'GigabitEthernet0/1'
```



You can use the following regular expression to extract the interface name from the given Syslog message:  

```
SCSS Copy code
(?<=Interface\s)(\S+)
```

This regex pattern looks for a word that comes right after the keyword "Interface" followed by a whitespace. The interface name is then captured in a capturing group.

Here's a breakdown of the regex pattern:

- * `(?<=Interface\s)` is a positive lookbehind that asserts that the pattern is preceded by the word "Interface" followed by a whitespace.
- * `(\S+)` is a capturing group that matches one or more non-whitespace characters. This represents the interface name.

LLM Use case: Incident summarization

Events Detail:

Event #1:Interface in issue: {'device_name': 'PE-13', 'interface_ip_address': '12.13.2.1/30', 'interface_name': 'et-1-2-2', 'interface_role': 'edge', 'interface_type': '10G', 'link_id': '12.13.2', 'peer_device_name': 'E-13-2', 'peer_interface': 'et-1-1-1', 'peer_interface_ip_address': '12.13.2.2/30', 'unique_id': 'PE-13-et-1-2-2', 'alert': 'Interface in issue', 'metric': 'Interface in issue'}

Event #2:Interface out issue: {'device_name': 'E-13-2', 'interface_ip_address': '12.13.2.2/30', 'interface_name': 'et-1-1-1', 'interface_role': 'edge', 'interface_type': '10G', 'link_id': '12.13.2', 'peer_device_name': 'PE-13', 'peer_interface': 'et-1-2-2', 'peer_interface_ip_address': '12.13.2.1/30', 'unique_id': 'E-13-2-et-1-1-1', 'alert': 'Interface out issue', 'metric': 'Interface out issue'}

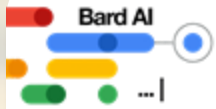
Event #3:Interface in issue: {'device_name': 'PE-13', 'interface_ip_address': '12.13.2.1/30', 'interface_name': 'et-1-2-2', 'interface_role': 'edge', 'interface_type': '10G', 'link_id': '12.13.2', 'peer_device_name': 'E-13-2', 'peer_interface': 'et-1-1-1', 'peer_interface_ip_address': '12.13.2.2/30', 'unique_id': 'PE-13-et-1-2-2', 'alert': 'Interface in issue', 'metric': 'Interface in issue'}

Event #4:Interface out issue: {'device_name': 'E-13-2', 'interface_ip_address': '12.13.2.2/30', 'interface_name': 'et-1-1-1', 'interface_role': 'edge', 'interface_type': '10G', 'link_id': '12.13.2', 'peer_device_name': 'PE-13', 'peer_interface': 'et-1-2-2', 'peer_interface_ip_address': '12.13.2.1/30', 'unique_id': 'E-13-2-et-1-1-1', 'alert': 'Interface out issue', 'metric': 'Interface out issue'}

Event #5:Interface in issue: {'device_name': 'PE-13', 'interface_ip_address': '12.13.2.1/30', 'interface_name': 'et-1-2-2', 'interface_role': 'edge', 'interface_type': '10G', 'link_id': '12.13.2', 'peer_device_name': 'E-13-2', 'peer_interface': 'et-1-1-1', 'peer_interface_ip_address': '12.13.2.2/30', 'unique_id': 'PE-13-et-1-2-2', 'alert': 'Interface in issue', 'metric': 'Interface in issue'}

I need a 3 line summary:

LLM Use case: Incident summarization



There is an interface issue between PE-13 and E-13-2. The issue is recurring and has been happening for 5 events.



Three-line summary:



1. Five interface issue events occurred between devices PE-13 and E-13-2.
2. The involved interfaces are et-1-2-2 (PE-13) and et-1-1-1 (E-13-2) with IP addresses 12.13.2.1/30 and 12.13.2.2/30, respectively.
3. Issues are reported as either "Interface in issue" or "Interface out issue" for both devices.

LLM Use case: Incident summarization (2)

Events Detail:

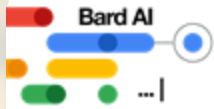
Event #1: BGP session down: {'device_name': 'PE-13', 'interface_ip_address': '12.13.2.1/30', 'interface_name': 'et-1-2-2', 'interface_role': 'edge', 'interface_type': '10G', 'link_id': '12.13.2', 'peer_device_name': 'E-13-2', 'peer_interface': 'et-1-1-1', 'peer_interface_ip_address': '12.13.2.2/30', 'unique_id': 'PE-13-et-1-2-2', 'alert': 'BGP session down', 'metric': 'BGP state', 'timestamp': 6201},

Event #2 : Interface down: {'device_name': 'PE-13', 'interface_ip_address': '12.13.2.1/30', 'interface_name': 'et-1-2-2', 'interface_role': 'edge', 'interface_type': '10G', 'link_id': '12.13.2', 'peer_device_name': 'E-13-2', 'peer_interface': 'et-1-1-1', 'peer_interface_ip_address': '12.13.2.2/30', 'unique_id': 'PE-13-et-1-2-2', 'alert': 'Interface down', 'metric': 'Interface in issue', 'timestamp': 5601},

Event #3: OSPF down issue: {'device_name': 'PE-13', 'interface_ip_address': '12.13.2.1/30', 'interface_name': 'et-1-2-2', 'interface_role': 'edge', 'interface_type': '10G', 'link_id': '12.13.2', 'peer_device_name': 'E-13-2', 'peer_interface': 'et-1-2-2', 'peer_interface_ip_address': '12.13.2.1/30', 'unique_id': 'E-13-2-et-1-1-1', 'alert': 'OSPF down', 'metric': 'OSPF state', 'timestamp': 5606}.

I need a summary of what happened:

LLM Use case: Incident summarization (2)



Sure. Here is a summary of what happened:

- Event #1 and Event #3: BGP session down and Interface down.
- Event #2: OSPF down issue.

It seems that there is an issue with the BGP and OSPF routing protocols on PE-13. This is causing the interface to go down and the BGP session to drop. This is a critical issue that needs to be addressed immediately.

I would recommend that you contact your network administrator or vendor to investigate this issue further.



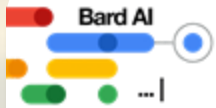
Summary:



1. BGP Session Down (Timestamp: 6201): The BGP session between devices PE-13 and E-13-2 went down, impacting routing information exchange and potentially causing network reachability issues.
2. Interface Down (Timestamp: 5601): The et-1-2-2 interface on device PE-13 went down, affecting the connection between devices PE-13 and E-13-2. This could be due to hardware, software, or physical connectivity issues.
3. OSPF Down Issue (Timestamp: 5606): The OSPF protocol failed to maintain adjacency between devices PE-13 and E-13-2, disrupting routing information exchange and affecting network performance. This issue was likely triggered by the interface going down on device PE-13.

LLM Use case: Root Cause Analysis

I need a root cause analysis of the incident above:



Here is a possible root cause analysis:

- **BGP configuration error:** The BGP configuration on PE-13 may be incorrect, which is causing the BGP session to drop.
- **Hardware failure:** There may be a hardware failure on PE-13, which is causing the interface to go down.
- **Software bug:** There may be a software bug on PE-13, which is causing the BGP session to drop.

To further investigate this issue, you can check the BGP configuration on PE-13, check the hardware on PE-13, and check the software on PE-13. You can also contact your network administrator or vendor for assistance.

If the BGP configuration error is the root cause, you can correct the error by following the steps in the BGP configuration guide for your network. If the hardware failure is the root cause, you can replace the faulty hardware. If the software bug is the root cause, you can upgrade the software to the latest version.



Root Cause Analysis:

The sequence of events suggests that the primary cause of this incident was the "Interface Down" event between devices PE-13 and E-13-2.

1. Interface Down (Timestamp: 5601): The et-1-2-2 interface on device PE-13 went down, resulting in a loss of connectivity between devices PE-13 and E-13-2. The cause of the interface going down could be attributed to a hardware or software issue, a physical cable problem, or misconfiguration. A thorough investigation of logs, hardware diagnostics, and physical inspection of the cabling and connectors should be conducted to identify the root cause of the interface failure.


2. OSPF Down Issue (Timestamp: 5606): The OSPF down issue occurred shortly after the interface went down, which implies that it was triggered by the loss of connectivity between the two devices. OSPF relies on an active interface to maintain adjacencies and exchange routing information. With the interface down, OSPF could not maintain its adjacency, resulting in the OSPF down issue.

3. BGP Session Down (Timestamp: 6201): The BGP session down event was also likely caused by the interface going down between devices PE-13 and E-13-2. Since BGP requires an active interface to exchange routing information, the loss of connectivity due to the interface failure would have disrupted the BGP session as well.

In conclusion, the root cause of this incident appears to be the interface failure on device PE-13. Resolving the underlying issue with the et-1-2-2 interface should restore network connectivity, subsequently resolving the OSPF and BGP issues.

LLM use case: Create a configuration for Cloud infrastructure

I need a Terraform template for 3 virtual machines in a Google VPC subnet:



Sure, here is a Terraform template for 3 virtual machines in a Google VPC subnet:

```
provider "google" {
  project = var.project_id
  region = var.region
}


resource "google_compute_network" "default" {
  name = "default"
}

resource "google_compute_subnet" "default" {
  network = google_compute_network.default.name
  name = "default"
  ip_cidr_block = "10.0.0.0/16"
  region = var.region
}

resource "google_compute_instance" "vm1" {
  name = "vm1"
  machine_type = "n1-standard-1"
  zone = "us-central1-a"
  network_interface {
    network = google_compute_network.default.name
    access_config {
      // Ephemeral IP address
    }
  }
  boot_disk {
    initialize_params {
      image = "debian-cloud/debian-9"
    }
  }
  metadata = {
    "user-data" = <<EOF
#!/bin/bash

echo "Hello, world!"
EOF
  }
}

resource "google_compute_instance" "vm2" {
  name = "vm2"
  machine_type = "n1-standard-1"
  zone = "us-central1-b"
  network_interface {
    network = google_compute_network.default.name
    access_config {
      // Ephemeral IP address
    }
  }
  boot_disk {
    initialize_params {
      image = "debian-cloud/debian-9"
    }
  }
}
```



Sure, here's a Terraform template that creates 3 virtual machines in a Google VPC subnet:

```
graph TD
  subgraph Terraform_Template [Terraform Template]
    direction TB
    T1["# Set the provider for Google Cloud  
provider \"google\" {  
  credentials = file(\"path/to/credentials.json\")  
  project = \"your-project-id\"  
  region = \"us-central1\"  
}  
# Create a VPC network  
resource \"google_compute_network\" \"vpc_network\" {  
  name = \"my-vpc-network\"  
}  
# Create a subnet within the VPC network  
resource \"google_compute_subnetwork\" \"vpc_subnet\" {  
  name = \"my-vpc-subnet\"  
  network = google_compute_network.vpc_network.self_link  
  ip_cidr_range = \"10.0.1.0/24\"  
}  
# Create a firewall rule to allow SSH traffic  
resource \"google_compute_firewall\" \"allow_ssh\" {  
  name = \"allow-ssh\"  
  network = google_compute_network.vpc_network.self_link  
  allow {  
    protocol = \"tcp\"  
    ports = [\"22\"]  
  }  
  source_ranges = [\"0.0.0.0/0\"]  
}"]
  end
```

LLM Use case: Configuration compliance verification

Given the following configuration:



```
set protocols bgp group internal-peers local-address 192.168.1.1  
set protocols bgp group internal-peers neighbor 192.168.1.2  
set protocols bgp group internal-peers neighbor 192.168.1.2 remote-as 65000  
set protocols bgp group internal-peers neighbor 192.168.1.2 next-hop-self  
set protocols bgp group internal-peers neighbor 192.168.1.2 update-source loopback0  
set protocols bgp group internal-peers neighbor 192.168.1.2 hold-time 200
```

Given the following template:

```
set protocols bgp group internal-peers local-address <IP address>  
set protocols bgp group internal-peers neighbor <IP address>  
set protocols bgp group internal-peers neighbor <IP address> remote-as 65000  
set protocols bgp group internal-peers neighbor <IP address> next-hop-self  
set protocols bgp group internal-peers neighbor <IP address> update-source loopback0  
set protocols bgp group internal-peers neighbor <IP address> hold-time 180
```

Is there any non-compliance? If so where?



The Junos router configuration and the template you provided match, with the exception of the "hold-time" value.  


In the configuration you provided, the hold-time is set to 200 seconds, whereas in the template it's set to 180 seconds. This is a minor difference and may be intentional, but I wanted to point it out nonetheless.

Other than that, there is no non-compliance between the Junos router configuration and the template you provided. They both use the same configuration hierarchy and options.

Configuration translation

old Manually programmed rules

JUNIPER NETWORKS | Support Portal Home Knowledge Quick Links ▾

 Knowledge Base

[Back](#)

[Archive] IOS-to-JUNOS (I2J) Conversion Tool - Tool Fact Sheet

Article ID **KB15482** Created 2009-10-10 Last Updated 2020-03-10

Description

Note: The I2J tool has been decommissioned.

The Juniper Networks IOS-to-JUNOS (I2J) Conversion (I2J) Tool was a software configuration translator that converts Cisco IOS to Juniper Networks JUNOS Software full IOS Configuration files for Cisco 65xx, 75xx and 12xxx platforms into Juniper Networks JUNOS format).

Solution

The IOS-to-Junos conversion tool is no longer available. Please visit the Juniper Networks [TechLibrary](#) product documentation for assistance with configuration or the [CLI Explorer](#) tool for assistance with Junos statements and commands.

Modification History

2019-12-23: Archived, **The I2J tool has been decommissioned.**




new Deep learning models

Here is an example configuration for Junos:

```
set protocols bgp group internal-peers type internal
set protocols bgp group internal-peers local-address 192.168.1.1
set protocols bgp group internal-peers neighbor 192.168.1.2
set protocols bgp group internal-peers neighbor 192.168.1.2 remote-as 65000
set protocols bgp group internal-peers neighbor 192.168.1.2 next-hop-self
set protocols bgp group internal-peers neighbor 192.168.1.2 update-source loopback0
set protocols bgp group internal-peers neighbor 192.168.1.2 hold-time 180
```

I need a translation to Nokia SROS

BARD's answer:

 Here is a translation of the Junos configuration to Nokia SROS:

```
# BGP internal peers
router bgp 65000
  neighbor 192.168.1.2
    ip address 192.168.1.2
    bgp-peer local
    remote-as 65000
    next-hop-self
    update-source loopback0
    hold-time 180
```


Image recognition



Can you describe in detail what you see in the image?



Certainly! Here's a detailed description of the image:



1. The primary object in the image is a black networking device, identified as a TP-Link Fast Ethernet Switch, model TL-SF1008.
2. The switch is rectangular in shape and sits atop a beige or light brown surface.
3. The front side of the switch, which faces the viewer, has several ports and indicators:
 - A green "Power" LED light on the left-most side is illuminated.
 - Next to the "Power" LED, there are eight ethernet ports labeled from 1 to 8.
 - Ports 1 through 4 have white ethernet cables plugged into them. The cables extend outwards and to the left, and some run beneath the switch.
 - Ports 5 to 8 are empty.
4. A human hand is visible on the left side of the image, pointing or pressing something on the switch, near the "Power" LED.
5. The background shows a portion of a black or dark-colored object, possibly another device or equipment, but its exact nature is not clear from the image.



Can you tell me what ports are up and which ones are down, out of the ports that have a cable connected?



Based on the image:



1. Ports 1 and 2: Both have white ethernet cables connected, and the corresponding LED lights adjacent to the ports are illuminated, indicating that these ports are "up" or active.
2. Port 3: Has a white ethernet cable connected, but the corresponding LED light is not illuminated, indicating this port is "down" or inactive.
3. Port 4: Has a white ethernet cable connected, and its adjacent LED light is illuminated, indicating that this port is "up" or active.

In summary:

- Ports 1, 2, and 4 are "up" or active.
- Port 3 is "down" or inactive.



Can you re-evaluate which ports are up and which are down?



Certainly! Based on the image provided:



Ports 1, 2, and 3: All three have white ethernet cables connected, and the corresponding LED lights adjacent to the ports are illuminated, indicating that these ports are "up" or active.

Port 4: Has a white ethernet cable connected, but the corresponding LED light is not illuminated, indicating this port is "down" or inactive.

In summary:

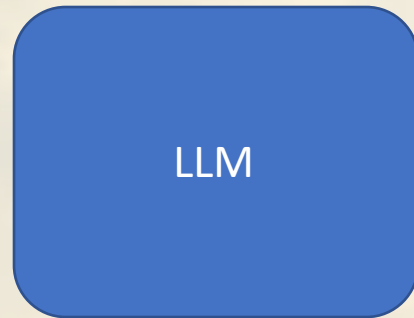
- Ports 1, 2, and 3 are "up" or active.



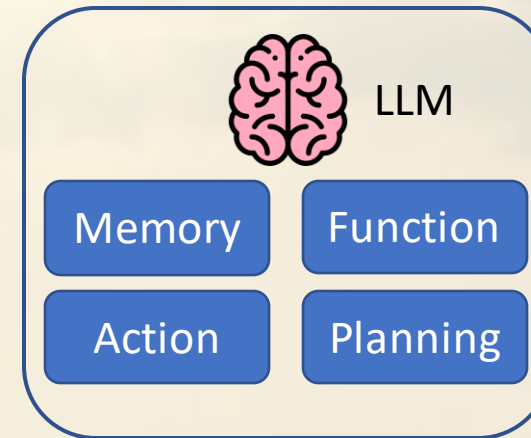
The next frontier: Agents

LLM-powered autonomous agents

From



To



Imperative / Prescriptive

Task execution:

- Generate text
- Summarize text
- Q&A
- ...

Declarative / Intent-based

LLM Powered agent

- Complex problem solver
- Autonomous execution

anetta.ai

autonomous **network** troubleshooting **agent**

where **AI-compiled knowledge** meets **networks**

Industry first autonomous intelligent agent for networks and clouds

conditional search

Anetta

How can I help you?

Start a new conversation

Anetta

Hi I'm Anetta, your automated network troubleshooting agent, how can I help you?. Please describe the problem you need me to troubleshoot or the information you need.

Your Answer

In the network some routers have vrf "test" configured. For those routers with vrf "test" configured, I need to know the next-hop for prefix 4.4.4.4/32 in the global routing table

Anetta

Thanks!. You need to get the following information from the network: User needs information about the next-hop for prefix 4.4.4.4/32 in the global routing table for routers with vrf 'test' configured in the network.

Anetta

The information requested is the following: {'vsrxng6': '20.20.21.2', 'vsrxng7': '20.20.22.2'}

Anetta

Hope this response was helpful.

Anetta

FYI, in this request you saved 213 seconds of your time.

configuration analysis

 Anetta

Hi I'm Anetta, your automated network troubleshooting agent, how can I help you?. Please describe the problem you need me to troubleshoot or the information you need.

Your Answer

I need to know the detailed differences between the configurations of routers vsrxng3 and vsrxng6

 Anetta

Thanks!. You need to get the following information from the network: The user is requesting detailed comparison of configurations between routers vsrxng3 and vsrxng6

 Anetta

The information requested is the following: Summary: The detailed differences between the configurations of routers 'vsrxng3' and 'vsrxng6' have been identified. They differ in host-name, root-authentication, login user, interfaces, and protocols. The 'vsrxng3' router has two interfaces ge-0/0/0 and ge-0/0/1, while 'vsrxng6' has ge-0/0/0 and ge-0/0/2. The 'vsrxng3' router uses OSPF and LDP protocols, while 'vsrxng6' uses OSPF, BGP, and LDP. The 'vsrxng6' router also has a routing-instance named 'test', which is not present in 'vsrxng3'. Details: The configurations of routers 'vsrxng3' and 'vsrxng6' have several differences. For instance, the host-name, root-authentication, login user, interfaces, and protocols are different. The 'vsrxng3' router has two interfaces ge-0/0/0 and ge-0/0/1, while 'vsrxng6' has ge-0/0/0 and ge-0/0/2. The 'vsrxng3' router uses OSPF and LDP protocols, while 'vsrxng6' uses OSPF, BGP, and LDP. The 'vsrxng6' router also has a routing-instance named 'test', which is not present in 'vsrxng3'.

 Anetta

Hope this response was helpful.

 Anetta

FYI, in this request you saved 365 seconds of your time.

troubleshoot a network problem

 Anetta

Hi I'm Anetta, your automated network troubleshooting agent, how can I help you?. Please describe the problem you need me to troubleshoot or the information you need.

Your Answer

Router vsrxng6 has a BGP session towards the loopback address (7.7.7.7) of router vsrxng7. However, this BGP session is not established. I need to know why.

 Anetta

Thanks!. You need to solve the following problem: The user is experiencing a problem with a BGP session between router vsrxng6 and router vsrxng7. The session is not established and the user needs to know the reason.

 Anetta

As step #1 I will focus on the following: Check the BGP configuration on both routers to ensure they are correctly configured to establish a BGP session. This includes checking the neighbor IP address, AS number, and other BGP parameters.

 Anetta

I found a root cause!

 Anetta

PROBLEM ROOT CAUSE: The BGP session is not established because the AS number for the BGP neighbor on router vsrxng6 is incorrectly configured as 65002 instead of 65001.

 Anetta

RECOMMENDED REMEDIATION ACTION: Correct the AS number for the BGP neighbor on router vsrxng6 to 65001.

©2023 Javier Antich



Network infrastructures are more complex than ever



Qualified resources are more scarce than ever

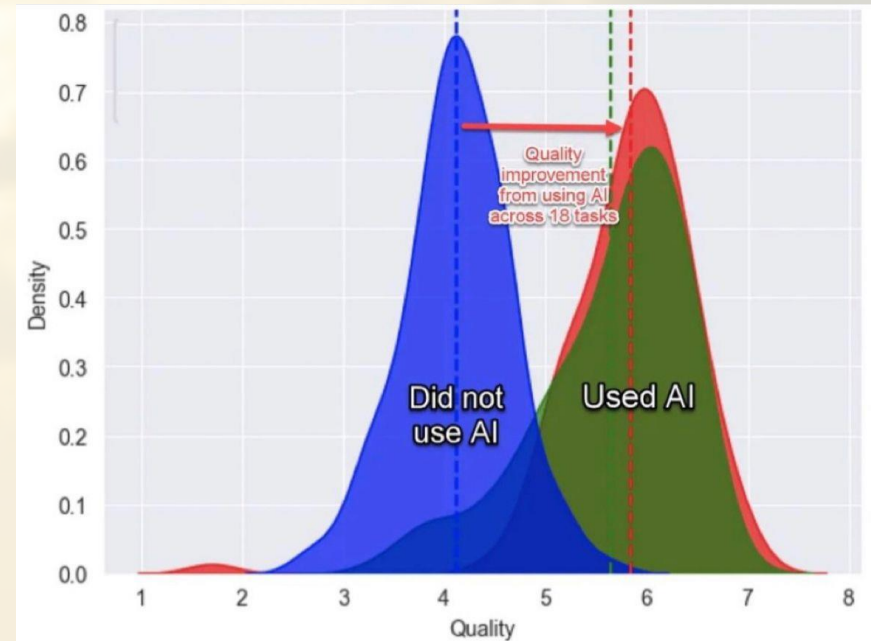
LLM-Powered Autonomous Agents



Telecom infrastructures are more mission critical than ever

Future trends for LLMs in Networks and Clouds

- More powerful LLMs (GPT-5, etc.)
- Longer context in LLMs
- Lower cost in cloud hosted LLMs.
- More domain specific LLMs (Anthropic to build Telco LLM with SK Telecom and DT)
- Powerful autonomous Agents: Anetta.ai
- More network software that embeds logic using LLMs.
- More pressure in your network, if you need to train and run your own LLMs

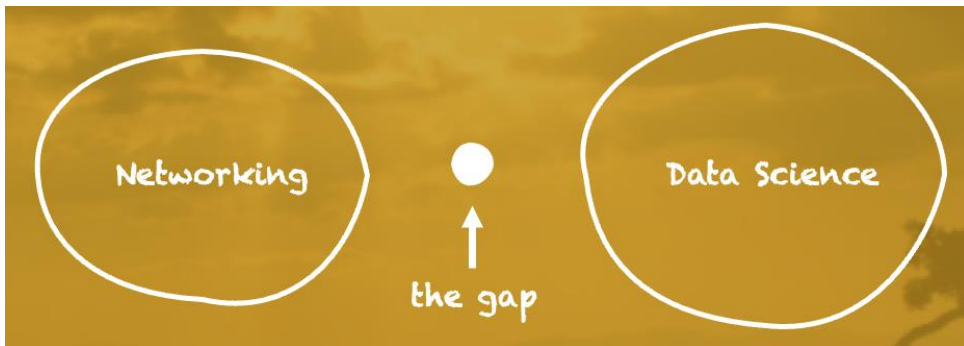


Distribution of output quality across all the tasks. The blue group did not use AI, the green and red groups used AI, the red group got some additional training on how to use AI.

What's in it for you?

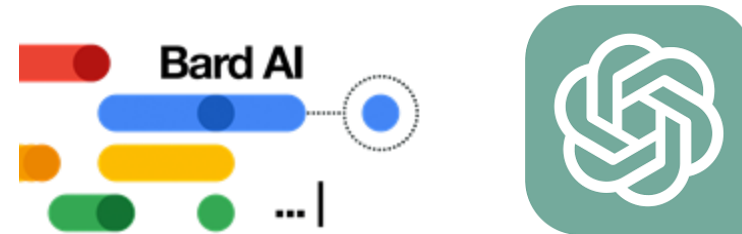
As a Network Engineer...

- Understand LLMs and leverage them where applicable.
- Identify opportunities.
- Learn & practice.



As an Automation Engineer...

- The new battlefield (Knowledge Plane) is a reality.
- It is already defining market leaderships (example: Mist)
- LLMs coming to disrupt the market, take advantage of them.



Thanks!

Machine Learning for Network and Cloud Engineers

Get ready for the next Era of Network Automation

Javier Antich
foreword by Ivan Pepelnjak



anetta  **.ai**

**Welcome to the future of
network and cloud operation**