



DC – AI – Routing – Switching – Security – Wireless – Automation

LLM Connector: a Conversational AI Assistant for Your Network

An HPE Juniper Networking Techpost

Julian Lucek

2025-06-27

HPE



Contents

1	LLM Connector: a Conversational AI Assistant for Your Network	2
1.1	Introduction	2
1.2	How LLM Connector Works	2
1.3	Example queries	4
1.3.1	Troubleshooting example: NTP	5
1.3.2	Troubleshooting example: BGP	8
1.3.3	Requesting Graphs from LLM Connector	11
1.3.4	Queries about VPN Services	12
1.3.5	Queries about Active Assurance	13
1.3.6	Multi-lingual capability	15
1.4	Conclusion	16
1.5	Useful links	16
1.6	Glossary	16
1.7	Acknowledgments	17

1 LLM Connector: a Conversational AI Assistant for Your Network

Julian Lucek - 06/27/2025

There has been a lot of interest recently in Large Language Models (LLMs). One of the major applications of LLMs is conversational AI that enables natural language interactions between people and chatbots. In this article we'll talk about LLM Connector, which is a chatbot within Routing Director that leverages LLMs.

This article has been prepared and co-written by Julian, Vijay Gadde, Harsha Lakshmikanth, and Nilesch Simaria.

1.1 Introduction

Until now, network operators have been able to observe their network via the CLI of individual routers, or via the graphical user interface of Routing Director. LLM Connector provides an additional way of interacting with the network, and is particularly suited to ad-hoc queries and troubleshooting that would otherwise involve logging onto multiple routers, executing show commands, and trawling through the outputs to find the relevant information.

Note: LLM Connector is the new name for Ask Paragon and Routing Director is the new name for Paragon Automation

1.2 How LLM Connector Works


image below shows how LLM Connector works. LLM Connector is an application that resides within Routing Director. Once connected to the LLM, LLM Connector prepares it to answer the user's queries in the appropriate way by issuing a prompt. The prompt tells the LLM that it is a helpful and efficient Routing Director assistant whose primary role is to assist with network monitoring tasks using the Routing Director platform.

The user types their questions about the network into the LLM Connector chat-window. The question is passed verbatim to the LLM and is interpreted by the LLM. The LLM Connector app exposes functions to the LLM that the LLM can invoke in order to gather information pertinent to the user's question. These functions include:

- Retrieve the router inventory
- Issue show-commands via RPCs on one or more routers
- Retrieve KPIs and metrics that Routing Director has ingested, analysed, and stored in its databases
- Retrieve results from Routing Director Active Assurance tests and monitors
- Retrieve alarm and alert information from Routing Director.

Having digested the user's question, the LLM works out which functions it needs to call in order to answer it successfully. Having received the responses to the function-calls from the LLM Connector app, the LLM formulates a response in a human-friendly form. It passes this to the LLM Connector app, which sends it to the user's chat-window.

In general, answering a given question may involve several rounds of interaction behind the scenes between the LLM Connector app and the LLM, to enable the LLM to gather all the information it needs, but the user is unaware that this is happening.



In Version 2.4, LLM Connector supports OpenAI and Llama LLMs. Interestingly, the “stock” version of the LLM is used - it has had no additional training about networking beyond the general content of the internet that it has been exposed to during its training. Sometimes, when using an LLM to support a specialist role, Retrieval-Augmented Generation (RAG) is used to provide access to domain-specific information that the LLM would not have had access to during its training. However, in our case we found that no RAG was needed, due to the large amount of networking-related information on the internet that LLMs are exposed to as a matter of course during their training. As a result, as we’ll see shortly, LLM Connector does a remarkably good job of formulating useful replies to queries about the network without requiring RAG.

The LLM instance is provided by your organisation, rather than by Juniper. The chances are that your organisation already has an LLM in place for other purposes. This makes it easier to have the LLM instance conform to your company’s policies. For example, it could be one that your organisation hosts on-prem, or could be one that your organisation purchases as a cloud service, for example Azure OpenAI.

As you know, LLMs are not perfect and can sometimes hallucinate or give responses that don’t answer the user’s request very well. For this reason, while LLM Connector can answer your questions about the network, we have made sure, at least for now, that it can’t do anything that is potentially disruptive such as changing router configurations, shutting down ports or rebooting routers!

image shows the configuration of an LLM Connector profile on LLM Connector. You can create multiple profiles, each pointing to a different LLM, and switch between them, for example when experimenting with different LLM models to see which one suits you best, taking into account cost and performance.

LLM Connector Configuration ?

X

Name *

GPT4o-model

Provider * ?

Azure Open AI

▼

Model * ?

gpt-4o

API Key ?

Existing API Key

🔑

Base URL * ?

https://yoururl.openai.azure.com/

✓

API Version * ?

2025-01-01-preview

Active ?

☒

Delete

Save

Cancel

1.3 Example queries

Have done that, LLM Connector is ready to use. You might be wondering what types of query LLM Connector can deal with, so let's ask it directly...

What types of query can you answer about my network?

10:02



I can assist with numerous aspects of your network monitoring and management, including:

- Device Monitoring:** Provide information about onboarded devices, including their health, alerts, location, and model specifications.
- Insights:** Deliver insights into specific network metrics such as CPU usage, memory utilization, protocol statuses (like BGP, LLDP, STP), and device-specific metrics.
- VPN Management:** Fetch metrics for VPNs, understand service types, measure performance, and ensure connectivity health.
- TSDB Metrics:** Extract stored metrics and analyze historical data for trends and patterns.
- Service Management:** Manage and diagnose the service instances associated with customer networks.
- Alerts Handling:** Identify and analyze alerts raised within the network.
- Operational Commands:** Execute Junos RPC commands in operational mode for detailed diagnostic information.
- Plot Data and Visualization:** Provide visual data representation, like graphs or plots, for specified metrics upon request.

Feel free to ask any question or task related to your network, and I'll guide or assist as needed.

As you can see in its reply, LLM Connector supports a wide variety of queries related to monitoring and managing the live network.

LLM Connector is good at consolidating information from multiple locations in an easily digestible form. In the example below, LLM Connector has retrieved some of the information from Routing Director and other information from the individual routers in much less time than it would have taken the user to gather all of

the information manually.

Give me a table showing the hostname of each router, the model type, the software version and the system uptime

10:24



Here's a summarized table combining the requested details about each router, including hostname, model, software version, and system uptime:

Hostname	Model Type	Software Version	System Uptime
mx204-24	MX204	23.2R2.21	2 Days 10 Hours
mx204-93	MX204	23.2R2.21	4 Days 5 Hours
mx204-19	MX204	23.2R2.21	7 Days 1 Hour
ptx10001-4	PTX10001-36MR	23.4R1-S2.6-EVO	10 Days 3 Hours
mx10003-2	MX10003	23.4R1-S2.4	20 Days 6 Hours
mx10003-1	MX10003	23.4R1-S2.4	25 Days 12 Hours

This overview combines extracted data from onboarded routers, focusing primarily on active devices within the network. Let me know if you need further breakdowns or analysis!

Response Time: 9.4s

☆☆☆☆☆ 10:24

In response to the query below, LLM Connector does a good job of working out which router has the newest software, even though some of the versions deployed are quite similar and the release nomenclature involves a mixture of numerals, letters and punctuation.

Which router has the newest software. Explain your reasoning.

10:25



Running:

- execute_junos_rpc
- get_device_list

The router with the newest software version is "ptx10001-4" with the model PTX10001-36MR and JunOS version "23.4R1-S2.6-EVO". This determination is based on the version number, with the latest iteration labeled "23.4R" compared to other devices using an older "23.2R" series software. Higher revision number ("R1-S2.6") also indicates more recent updates compared to other devices using "R1" or "S2.4" configurations.

1.3.1 Troubleshooting example: NTP

Let's see how LLM Connector can help with troubleshooting. First of all, let's ask for a list of critical alerts in the network. LLM Connector replies with a variety of alerts, including an issue with NTP on the router mx204-24.

List the critical alerts in the network

10:28



Running:

- get_alerts
- get_device_list

Below is a table summarizing the critical alerts in the network:

Hostname	Alert Description	Severity
mx204-24	NTP is not in sync	Critical
mx204-24	PEM1 is Offline	Critical
mx204-24	xe-0/1/7 link operator status is DOWN	Critical
mx204-93	xe-0/1/3 link operator status is DOWN	Critical
mx204-93	xe-0/1/3.0 status DOWN	Critical
ptx10001-4	et-0/2/9 link operator status is DOWN	Critical
ptx10001-4	et-0/1/6 link operator status is DOWN	Critical
ptx10001-4	et-0/2/11 link operator status is DOWN	Critical
mx10003-2	PEM3 is Offline	Critical
mx10003-2	PEM4 is Offline	Critical
mx10003-2	PEM5 is Offline	Critical
mx10003-2	BGP neighbor(14.15.0.2) session state is IDLE	Critical
mx10003-1	PEM3 is Offline	Critical
mx10003-1	PEM4 is Offline	Critical
mx10003-1	PEM5 is Offline	Critical

This table lists the devices with their critical alerts. Each alert warrants immediate attention to ensure network stability and security.

Let's now find out more about the NTP issue on mx204-24. As you can see below, LLM Connector gives some good advice about how to troubleshoot the issue:

Can you tell me more about the NTP not in sync alert on mx204-24?

10:36



Running:

- get_alerts
- get_device_list

Here are the details of the "NTP is not in sync" critical alert for the router **mx204-24**:

- Alert ID: **19a98a28-c4d7-472e-831d-423499578bb3**
- Severity: **Critical**
- Alert Category: **Hardware/System**
- Description: **NTP is not in sync**
- First reported: **04-11-2025**
- Last triggered: **04-28-2025**

Analysis:

1. **Source Issue:** NTP synchronization problems often occur due to network unavailability or misconfiguration within server accessibility or protocol setup.
2. **Configuration Verify:** Ensure NTP servers are correctly specified on the router.
3. **Connectivity:** Evaluate network accessibility and reachability of the configured NTP server.
4. **Offset Observation:** Inspect network time offset logs for potential discrepancies causing sync refusal.
5. **Hardware Specifics:** Examine hardware compatibility logs for any indicated inconsistencies.

Shall I assist you in diagnosing the issue further using detailed command-line investigations?

Let's take LLM Connector's advice about ensuring that the NTP server is correctly specified on router mx204-24. We'll start by asking what address is configured for the NTP server. LLM Connector in its reply gives the

address of the NTP server that appears in the configuration. Very usefully, it also volunteers the fact that the NTP configuration includes a deprecated “boot-server” statement, even though it wasn’t specifically asked to look out for this. LLM Connector also points out that this statement shouldn’t affect synchronization (so we know it is not the root cause of the NTP issue on mx204-24), but that it is good practice to remove that statement from the configuration.

What is the address of the NTP server configured on mx204-24?

19:38



Running:

- execute_junos_rpc
- get_device_list

The NTP server configured on **mx204-24** is at IP address **172.30.207.10**.

It is worth noting that the configuration includes a deprecated directive (`boot-server`), which also specifies the same address as the active NTP server. This should not impact synchronization but is recommended to be updated to align with newer standards in configuration practices.

Let’s now request LLM Connector to ping the address of the NTP server. As you can see from the output below, the pings timed out. It’s interesting to note that in its response, LLM Connector mentions the NTP server, even though we did not mention NTP explicitly in the request. This is because it retains context from previous queries, so “remembers” that the address 172.30.207.10 is related to the NTP server configuration on mx204-24:

Issue a ping request on mx204-24 to 172.30.207.10

19:40



Running:

- execute_junos_rpc

The ping results for the NTP server (172.30.207.10) from the router mx204-24 indicate the following:

- **Does not respond with requested packets. (Timeout):** There might be a connectivity issue or server blockage due to configurations or a network reachability malfunction.

Given that NTP is not working on mx204-24, but is working fine on the other routers, let’s get LLM Connector to show the NTP server address configured on each router. As we can see from the response below, the address is different on mx204-24 to the other routers, so the address on mx204-24 must be wrong!

give me a table showing the NTP server address configured on each router

10:46



Running:

- get_device_list
- execute_junos_rpc

Device Name	Configured NTP Server
mx204-24	172.30.207.10
mx204-93	172.30.240.1
mx204-19	172.30.240.1
ptx10001-4	172.30.240.1
mx10003-2	172.30.240.1
mx10003-1	172.30.240.1

Let's now go onto the CLI on mx204-24 to change the address of the NTP server to the correct one, bearing in mind that LLM Connector is not permitted to change router configurations. We'll also remove the deprecated NTP boot-server configuration, as advised by LLM Connector:

```

1 jlucek@mx204-24> show ntp associations
2
3 remote refid auth st t when poll reach delay offset jitter rootdelay rootdisp
4 =====
5 172.30.207.10 .INIT. - 16 u - 1024 0 0.000 +0.000 0.000 0.000 0.000
6
7 jlucek@mx204-24> show configuration system ntp
8 boot-server 172.30.207.10; ## Warning: 'boot-server' is deprecated
9 server 172.30.207.10;
10
11 jlucek@mx204-24# delete system ntp
12 jlucek@mx204-24# set system ntp server 172.30.240.1
13 jlucek@mx204-24# commit and-quit

```

Now we can have LLM Connector check the NTP status again on mx204-24. As you can see below, NTP is now up and running on that router!

now check the NTP status on mx204-24

11:00



Running:

- execute_junos_rpc
- get_device_list

The NTP status for **mx204-24** indicates the following:

- **NTP is synchronized:** Indicates sync_ntp, which shows synchronization is active.
- **Stratum:** 5 (indicating intermediate level hierarchy in NTP topology).
- **Root Delay:** 139.899 ms
- **Root Dispersion:** 196.647 ms

1.3.2 Troubleshooting example: BGP

Let's now look at another troubleshooting example, this time involving BGP. First of all, we'll ask if there are any issues with BGP. The LLM knows from its general training on the content of the internet that "Established" is the desired BGP state. As you can see from the output below, most of the BGP sessions are in that state. However,

as you can see in the mx10003-2 section, one of the BGP sessions is in the Idle state. This is highlighted again by LLM Connector in the summary section as something that needs attention. LLM Connector correctly notes that mx204-19 and ptx10001-4 don't have BGP configured (they are P-routers).

Look at the BGP session status on each router and highlight any issues

11:18



Running:

- insights_on
- execute_junos_rpc
- get_device_list

BGP Session Status and Identified Issues Across Network Devices

Router: **mx204-24**

- Peer Count: 3
- Active Peers: All peers are established.
- Issues: No notable issues detected.

Router: **mx204-93**

- Peer Count: 3
- Active Peers: All peers are established.
- Issues: No notable issues detected.

Router: **mx204-19**

- Status: BGP is not currently active on this device.

Router: **ptx10001-4**

- Status: BGP is not currently active on this device.

Router: **mx10003-2**

- Peer Count: 4 (VPN service configured).
- Active Peers: 3 peers are established.
- Issues: One peer is in idle state, requiring investigation.

Router: **mx10003-1**

- Peer Count: 4
- Active Peers: All peers are established.
- Issues: Although no critical issues were highlighted, monitoring for flaps might be useful.

Summary

Most routers are functioning without critical issues in their BGP sessions except for **mx204-19** and **ptx10001-4**, which do not have BGP activated. Specific focus is needed on **mx10003-2** for one idle peer to ensure full peer connectivity. If more detailed metrics are required, feel free to request further diagnosis.

Let's investigate further by asking which BGP session is idle on mx10003-2:

Which BGP peer is idle on mx10003-2?

11:28



Running:

- get_alerts
- get_device_list

BGP Peer Idle Status on Router mx10003-2

The BGP peer **14.15.0.2** is reported as idle on the router **mx10003-2**. This indicates a session issue that might require further investigation to establish connectivity properly.

We need to work out what to do next, so we'll seek advice from LLM Connector about why a BGP session might be in the Idle state. In its reply, LLM Connector gives several good reasons why this might be the case.

Give some reasons why the BGP session with peer 14.15.0.2 isn't established on mx10003-2

11:28



Potential Reasons for BGP Session Issues with Peer 14.15.0.2

The BGP session with the peer at IP 14.15.0.2 on router mx10003-2 might not be established due to several factors:

1. Configuration Mismatches:

- Incorrect peer IP address configuration.
- Different Autonomous System (AS) numbers configured between peers.

2. Network Connectivity Issues:

- Physical or logical link down status for the interface connected to the peer.
- Intermediate routing problems preventing reachability.

3. ACLs or Firewall Rules:

- Access Control Lists (ACLs) blocking the BGP traffic.
- Inadequate or restrictive firewall rules that prevent BGP establishment.

4. Authentication Issues:

- MD5 authentication set incorrect or not configured.

5. Outdated or Incorrect Software Versions:

- Disparities or incompatibilities in BGP implementation across devices.

6. Peer-Specific Configurations:

- Restrictions such as prefix limits or path filtering specifications.

Further investigation into the configuration and network connectivity of mx10003-2 is recommended to pinpoint the exact cause of this idle state. Should you require any targeted checks or logs, please specify.

Let's check if potential causes 1 or 2 in the list above might be the issue, by requesting a ping to the configured peer address. LLM Connector replies saying the peer address is unreachable:

Issue a ping from mx10003-2 to 14.15.0.2

11:30



Running:

- execute_junos_rpc
- get_device_list

The attempt to issue a ping to IP 14.15.0.2 from router mx10003-2 resulted in a timeout. This could indicate:

1. Connectivity Problems:

- The router might be unable to reach the specified IP address due to network issues.

2. Device Unreachability:

- Actual device unresponsiveness beyond the router itself.

Further troubleshooting steps include confirming the router's operational state and verifying any network policies or routes that could block such communication.

Now we'll tell LLM Connector to find out which interface on mx10003-2 is configured with the corresponding address, and then find out the status of that interface. As we can see from the output, the interface is administratively up but operationally down, so we have found the root cause of the BGP issue:

On router mx10003-2, find out which interface has the address 14.15.0.1. Then tell me the status of that interface.

11:35



Running:

- get_device_list
- execute_junos_rpc

The IP address 14.15.0.1 is assigned to the interface xe-0/0/0:1 on router mx10003-2.

The current status of this interface is:

- **Administrative Status:** Up
- **Operational Status:** Down

If you need further assistance or details, feel free to ask!

1.3.3 Requesting Graphs from LLM Connector

The example outputs from LLM Connector shown so far have been text-based. However, LLM Connector can also present information in the form of a graph. Let's see an example, in the context of fan-speed troubleshooting. First of all, let's find out about any critical alerts on one of the routers. As you can see, the fans are spinning above the alert threshold:

List the critical alerts on PTX10001-36MR 10:31



Running:

- get_alerts
- get_device_list

Critical Alerts on PTX10001-36MR (Device: ptx10001-4):

Hardware Alerts:

1. **Description:** Fan Tray 5 Fan 1 RPM percent 100% (21240 RPM) exceeds critical threshold 16992.
 - **Source:** Insights
 - **Severity:** Critical
2. **Description:** Fan Tray 3 Fan 2 RPM percent 98% (18720 RPM) exceeds critical threshold 15281.
 - **Source:** Insights
 - **Severity:** Critical
3. **Description:** Fan Tray 2 Fan 2 RPM percent 97% (18480 RPM) exceeds critical threshold 15241.
 - **Source:** Insights
 - **Severity:** Critical
4. **Description:** Fan Tray 5 Fan 2 RPM percent 98% (18720 RPM) exceeds critical threshold 15281.
 - **Source:** Insights
 - **Severity:** Critical
5. **Description:** Fan Tray 4 Fan 1 RPM percent 100% (21240 RPM) exceeds critical threshold 16992.
 - **Source:** Insights
 - **Severity:** Critical
6. **Description:** Fan Tray 0 Fan 1 RPM percent 100% (21360 RPM) exceeds critical threshold 17088.
 - **Source:** Insights
 - **Severity:** Critical
7. **Description:** Fan Tray 4 Fan 2 RPM percent 97% (18480 RPM) exceeds critical threshold 15241.
 - **Source:** Insights
 - **Severity:** Critical
8. **Description:** Fan Tray 2 Fan 1 RPM percent 100% (21240 RPM) exceeds critical threshold 16992.
 - **Source:** Insights
 - **Severity:** Critical
9. **Description:** Fan Tray 1 Fan 2 RPM percent 97% (18600 RPM) exceeds critical threshold 15340.
 - **Source:** Insights
 - **Severity:** Critical
10. **Description:** Fan Tray 3 Fan 1 RPM percent 100% (21240 RPM) exceeds critical threshold 16992.
 - **Source:** Insights
 - **Severity:** Critical
11. **Description:** Fan Tray 0 Fan 2 RPM percent 96% (18240 RPM) exceeds critical threshold 15200.
 - **Source:** Insights
 - **Severity:** Critical
12. **Description:** Fan Tray 1 Fan 1 RPM percent 100% (21360 RPM) exceeds critical threshold 17088.
 - **Source:** Insights
 - **Severity:** Critical

In order to find out for how long this has been happening, let's request a graph of the fan-speed on that router:

Plot me a graph of the fan speed on that router

10:32




Running:

- get_device_list
- insights_on
- get_tsdb_metrics_data
- get_metrics_name_from_vm

06/06/2025 10:02 AM - 06/06/2025 10:32 AM

30m 1H 3H 1D 1W



Let's ask why the fan speed might have increased. As you can see LLM Connector gives some useful advice about why this might be the case.

Can you suggest reasons for the fan speed to increase?

10:37



Reasons for Increased Fan Speed in Devices

Fan speed increases can typically be traced back to the following potential factors:

1. Temperature Increase:

- High ambient temperature around the device can raise the internal temperature of the equipment, causing fans to speed up to prevent overheating.

2. Processor Load:

- Compute-intensive operations or spikes in utilization (e.g., processing of large data streams or complex computations) generate excess heat that triggers elevated cooling mechanisms.

3. Hardware Usage:

- Extended use of specific components like power supplies or interfaces (e.g., SFPs and transceivers) leads to heat generation, requiring fans to adapt.

4. Blocked Cooling Paths:

- Dust, debris, or physical obstructions around the device can reduce the efficiency of heat dissipation, prompting the increased operation of fans.

5. Environmental Conditions:

- Poor airflow around the device's installation location or inadequate cooling infrastructure directly influences fan operation.

6. Hardware Degradation:

- Wear and degradation over time in fans or cooling units can cause them to work harder to maintain adequate cooling levels.

7. Configuration Updates:

- Certain firmware or software modifications can cause adjustments in dynamic fan policies.

8. Operational Changes:

- The introduction of new network workloads or hardware upgrades may impact thermal conditions.

Diagnosis Suggestions:

To accurately determine and mitigate the increased fan speeds, consider:

- Ensuring regular maintenance and cleaning.
- Monitoring thermal metrics and component operational metrics.
- Validating device operational loads align within expected norms.
- Reviewing environmental and hardware setup to optimize cooling efficiency.

If issues persist, proactive diagnostics and detailed analysis including comprehensive metrics review can be advisable.

1.3.4 Queries about VPN Services

LLM Connector can retrieve information from Routing Director about VPNs that it has orchestrated. Let's find out which VPNs are present on the network. LLM Connector responds by providing a table showing the name

of each VPN, the associated customer and the VPN type:

Give a list of VPNs in the network, and the customer associated with each

14:59



Running:

- get_service_instances
- get_vpns
- get_customers

VPNs and Associated Customers

VPN Name	Associated Customer	Design ID
ACME-resilient	ACME	I3vpn
ACME-picked-intf	ACME	I3vpn
ACME-ce-selection	ACME	I3vpn
Widget-pick-port	Widget	I3vpn

These are the VPNs configured in the network along with their associated customers.

Now, let's ask for details about which PEs are involved in the ACME-resilient VPN. LLM Connector helpfully points out that the interface on mx204-93 is down, and that static routing is used in all the VRFs:

which routers are involved with the ACME-resilient VPN?

15:00



Running:

- get_vpn_metrics
- get_customers
- get_vpn_health
- get_service_instances
- get_device_list
- get_vpns

Routers Associated with ACME-resilient VPN

Following are the routers involved in the ACME-resilient VPN:

Router Name	Interface Name	Interface Status
MX204-93	xe-0/1/3	DOWN
MX10003-1	xe-0/0/2:3	UP
MX10003-2	xe-0/0/1:1	UP

Each router uses static protocols for the network instance.

1.3.5 Queries about Active Assurance

So far in this blog, the examples shown have been from Version 2.4. In version 2.5, LLM Connector can also retrieve Active Assurance performance measurements from Routing Director. Let's first find out what type of performance measurements the system is capable of. The output below shows that a wide variety of measurements are supported.

list plugins

11/8



Running:

active_assurance_list_plugin_schema

The following are the available plugins:

Ping: Tests the reachability and response time of hosts.

RPM CISCO_TWAMP: Router Performance Measurement that configures CISCO_TWAMP probe service and collects results from the service running on devices.

RPM UDP: Router Performance Measurement that configures UDP probe service and collects results from the service running on devices.

RPM Y.1731 SLM: Router Performance Measurement that collects Y.1731 Synthetic Loss Measurement (SLM) results from the service running on devices.

TWAMP reflector: Runs a TWAMP reflector tool on a Test Agent interface.

IPTV MPEG: Receives one or several IPTV channels, measuring MPEG TS parameters according to ETSI TR 101290.

RPM Y.1731 DM: Router Performance Measurement that collects Y.1731 Delay Measurement (DM) results from the service running on devices.

TWAMP/TWAMP Light: Measures two-way loss and delay using the Two-way Active Measurement Protocol (TWAMP).

UDP: Flexible hub-and-spoke UDP traffic generation. Configurable rate, packet size, priority marking, and direction.

Netflix Speedtest: Tests the bandwidth and latency performance towards Netflix servers following fast.com Speedtest algorithm.

TCP: Flexible hub-and-spoke TCP traffic generation. Configurable rate, priority marking, and direction.

RPM RFC2544 INET: Router Performance Measurement that configures RFC2544 INET probe service and collects results from the service running on devices.

RPM HTTP: Router Performance Measurement that configures HTTP probe service and collects results from the service running on routers.

RPM RFC2544 CCC: Router Performance Measurement that configures RFC2544 CCC probe service and collects results from the service running on devices.

DNS: Tests the reachability and response time of DNS servers.

OTT - HLS: Measures user experience for adaptive video streaming using HLS (HTTP Live Streaming).

RPM TWAMP: Router Performance Measurement that configures TWAMP probe service and collects results from the service running on devices.

RPM TCP: Router Performance Measurement that configures TCP probe service and collects results from the service running on devices.

HTTP: Tests the reachability and response time of web servers.

Now let's find out what are the five highest Round-Trip Times (RTTs) measured by the TWAMP monitors that are running in the network.

list 5 max rtt for all monitors with "rpm_twamp" plugin

18/18




Running:

active_assurance_list_plugin_schema
active_assurance_list_measurements_with_metrics

Below are the top 5 maximum round-trip times (RTT) observed for the monitors equipped with the "rpm_twamp" plugin:

RTT: 104.966 msMeasurement ID: 6e672b56-28a6-4aa8-98a0-975022a237e4
Observed Time: 2025-06-17T02:18:23Z[Details](#)**RTT: 102.874 ms**Measurement ID: 6e672b56-28a6-4aa8-98a0-975022a237e4
Observed Time: 2025-06-18T21:27:59Z[Details](#)**RTT: 98.943 ms**Measurement ID: 6e672b56-28a6-4aa8-98a0-975022a237e4
Observed Time: 2025-06-19T05:55:16Z[Details](#)**RTT: 97.702 ms**Measurement ID: 6e672b56-28a6-4aa8-98a0-975022a237e4
Observed Time: 2025-06-18T20:13:07Z[Details](#)**RTT: 95.930 ms**Measurement ID: 6e672b56-28a6-4aa8-98a0-975022a237e4
Observed Time: 2025-06-19T08:24:31Z[Details](#)

As can be seen from the output above, each of the five data-points has a link to more details about the measurement. Let's click on the first one. This takes us to the chart below, which is automatically centred on the data-point of interest, i.e. the peak delay value of 104.966 ms. We can also see some secondary peaks that occurred within the time-window of the chart.



1.3.6 Multi-lingual capability

As the LLMs that LLM Connector supports are trained on the general content of the internet, they have been exposed to many languages. This means that each user can converse with LLM Connector in their preferred language, independently of which languages other users on the same Routing Director deployment are using. In the example below, we ask in Polish for a list of VPNs in that are present on the network.

Podaj mi listę sieci VPN obecnych w sieci 09:47



Running:

- get_service_instances
- get_vpnns
- get_customers

Lista dostępnych sieci VPN w Twojej sieci to:

Nazwa VPN	Typ sieci
ACME-resilient	L3VPN
ACME-picked-intf	L3VPN
ACME-ce-selection	L3VPN
Widget-pick-port	L3VPN

Jeżeli potrzebne są bardziej szczegółowe informacje o konkretnej VPN, daj znać!

In the example below, we ask for details in Bulgarian about the alerts on mx204-24:

Кажете ми повече за mx204-24

13:22

Running:

- get_alerts
- get_device_list

Детайлна информация за MX204-24:

Устройство:

- **Модел:** MX204
- **MAC адрес:** 00c52c50a850

Аларми:

Критични:

- **Описание:** PEM1 е изключен.
 - Категория: Хардуер
 - Последно съобщено: 28 април, 2025
 - Източник: Insights
 - Повторно настъпване: 2895 пъти
- **Описание:** Връзката xe-0/1/7 е в статус DOWN.
 - Категория: Интерфейс
 - Последно съобщено: 28 април, 2025
 - Източник: Insights
 - Повторно настъпване: 2890 пъти

1.4 Conclusion

We have seen that LLM Connector provides a very user-friendly way of finding out what is happening in the network using natural language. This saves the user a large amount of time and effort compared to the traditional approaches to monitoring and troubleshooting. The outputs are provided in a natural language format, accompanied by easy to read charts and tables.

1.5 Useful links

- <https://www.juniper.net/gb/en/products/network-automation/paragon-automation.html>
- <https://www.juniper.net/documentation/product/us/en/juniper-paragon-automation/>

1.6 Glossary

- BGP: Border Gateway Protocol
- CLI: Command-Line Interface
- KPI: Key Performance Indicator
- LLM: Large Language Model

- NTP: Network Time Protocol
- PE: Provider Edge
- RAG: Retrieval Augmented Generation
- RPC: Remote Procedure Call
- VPN: Virtual Private Network

1.7 Acknowledgments

This article has been prepared and co-written by Julian, Vijay Gadde, Harsha Lakshmikanth, and Niles Simaria.



DC – AI – Routing – Switching – Security – Wireless – Automation

© Copyright 2026 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty.

Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Trademark acknowledgments, if needed. All third-party marks are the property of their respective owners.

HEWLETT PACKARD ENTERPRISE

HPE.com

