

# nAnalyst

## Agentic Network Intelligence

# Networks Tell Everything

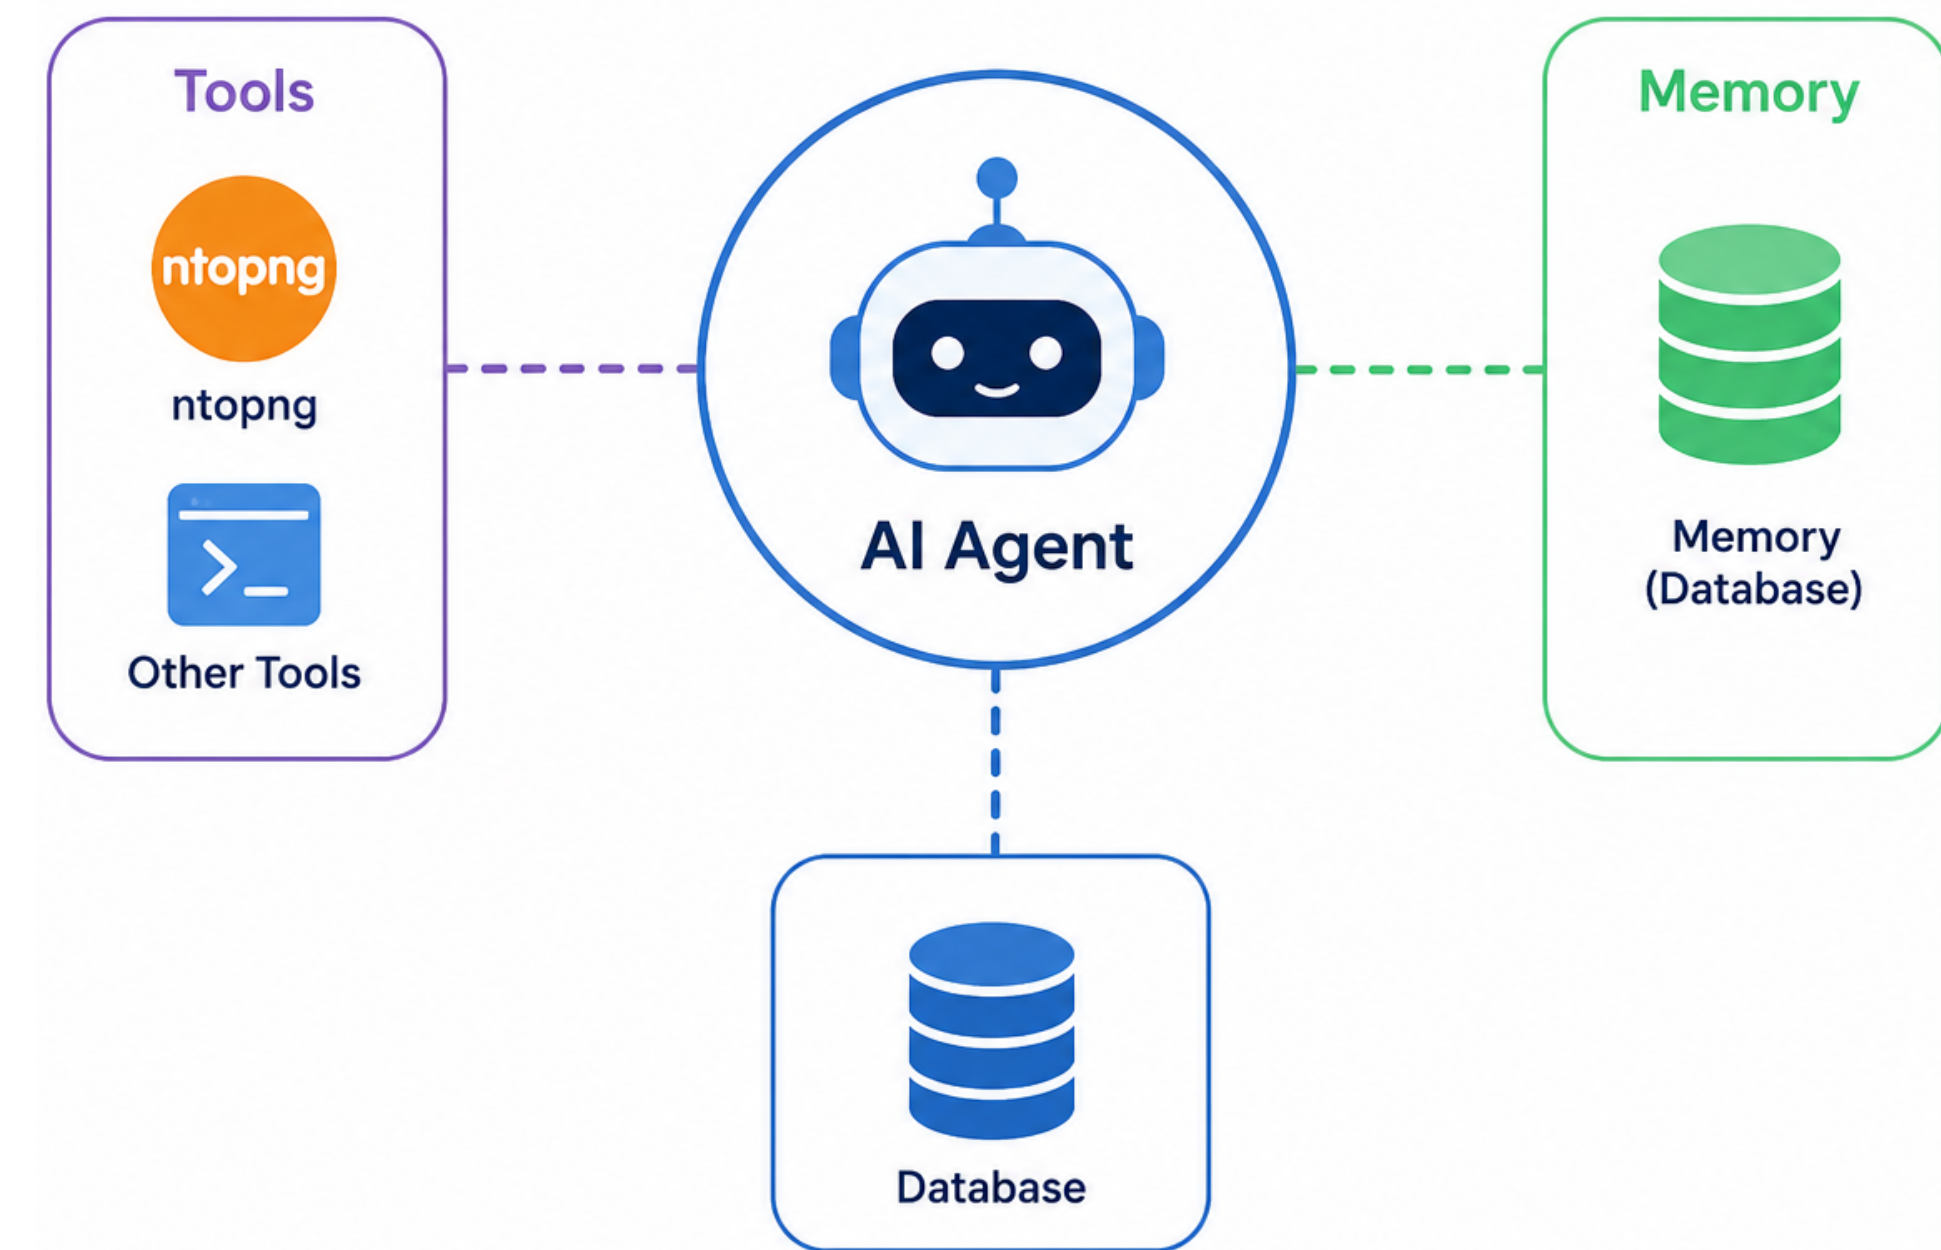
1. TBs of data per day
2. Thousands of alerts
3. Analysts spend hours investigating

Most tools answer **what**. nAnalyst answers **why**

# What Is An Agent?

A program that perceives its environment, makes decisions, and takes actions to reach a specific goal autonomously.

- **Tools:** connections that let the agent act in the real world. (Cmd, call APIs...)
- **Memory:** storage for agent findings, identity
- **React Loop:** reasoning action loop to reach the goal
- **MCP:** standard protocol to let AI agents discover and use tools in a universal way



# What is nAnalyst?

<b>Today</b>	<b>With nAnalyst</b>
Dashboards	Explanations
Alerts	Prioritized incidents
Queries	Grounded Answers
Manual investigation	Automated reasoning

- Agentic layer that lives in ntopng
- Not a chatbot wrapper
- Purpose built in ntopng

# From Data to Answers

nAnalyst is an autonomous network intelligence layer that explains what is happening in your network and why

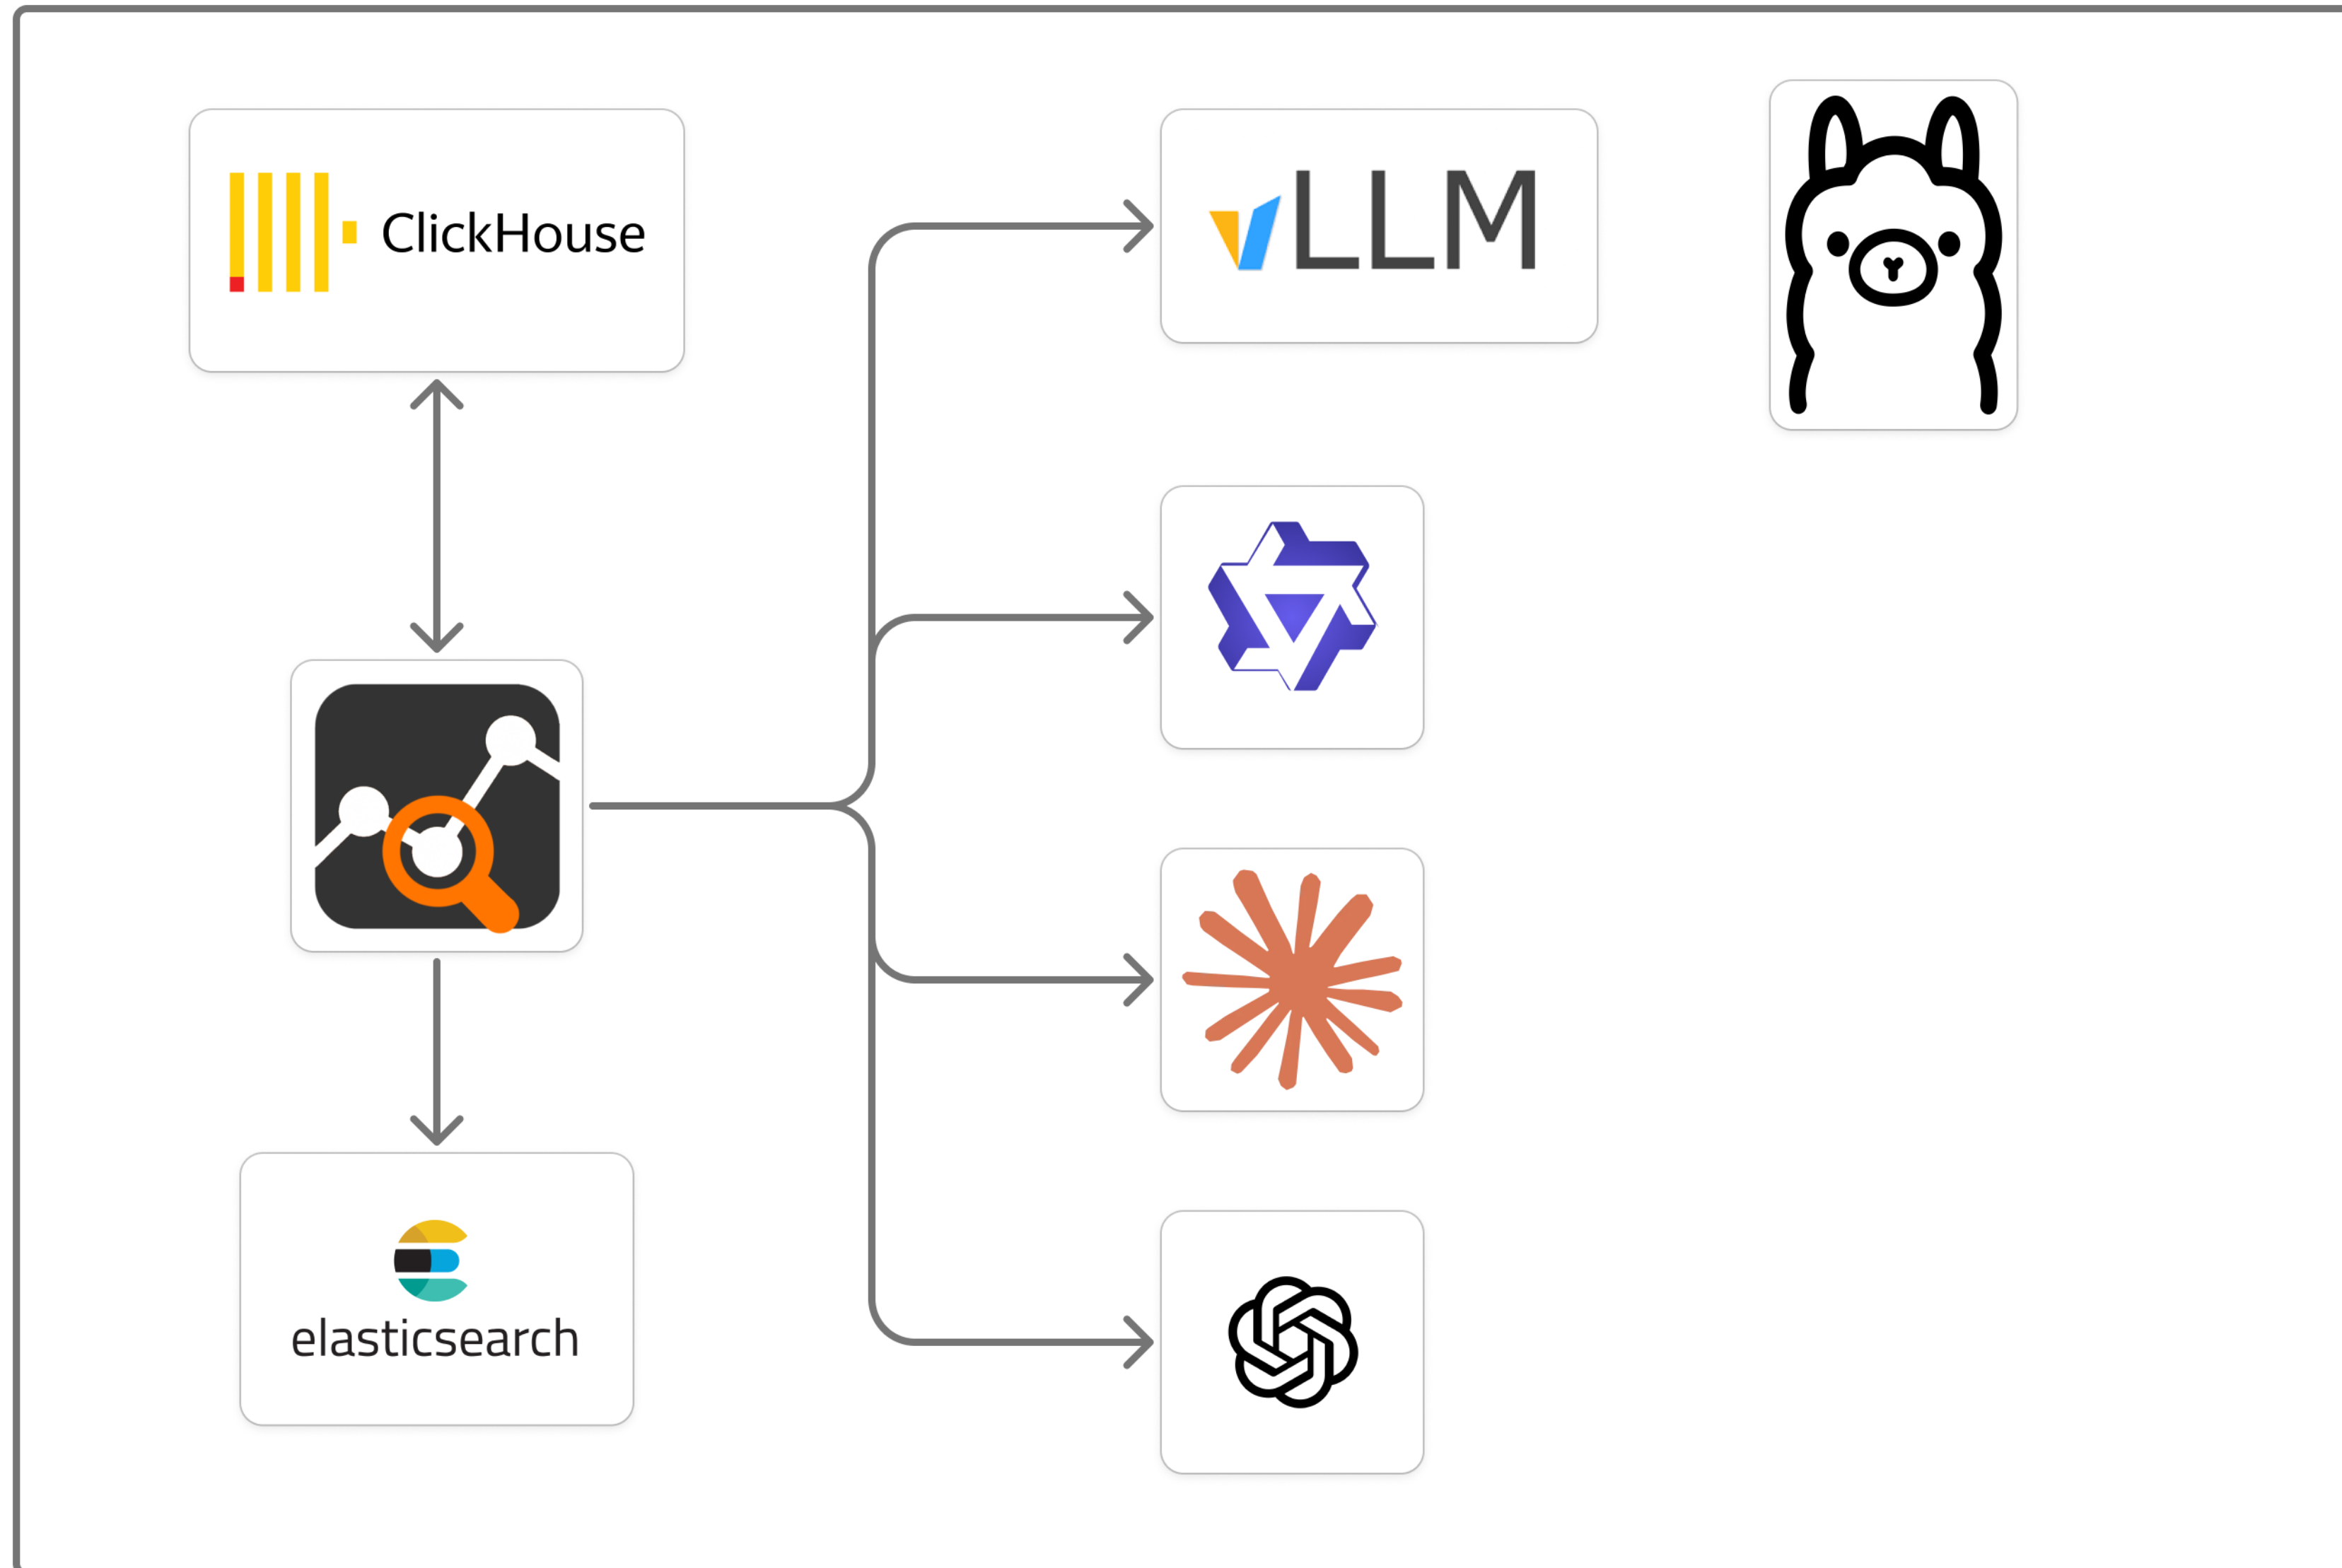
1. User asks a question
2. An agentic loop is started:
  1. Gather data
  2. Additional evidence is collected from flows, hosts, asset inventory, SNMP or alerts
  3. Evidence is visualized and persisted across sessions.

# Inference Servers

An inference server:

- Loads an LLM into GPU VRAM / system RAM
- Handles HTTP(S) requests from applications and AI agents
- Execute model inference: generate text, reasoning, tool calls...
- Returns responses through REST API
- Supports concurrent users and agents

# Deployment



# LLM Setup

## 1. Fully supports anthropic and openAI format

1. Pay per use APIs
2. AWS Bedrock compatible
3. Local LLM servers: llama-cpp, vllm, sglang

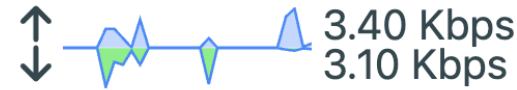
## 2. Easy Setup

1. Set LLM API key, endpoint and model name
2. No key leaves your machine, all is stored on the ntopng instance

# LLM Setup

- Dashboard
- Monitoring
- Alerts
- Flows
- nAnalyst
- Hosts
- Maps
- Interface
- Policies
- Settings
- Developer
- Help

enp1s0



2 1 12 6 10 25 ntop

Search, Notifications, User Profile

## Runtime Preferences

Search Preferences

Active Scan

Active Monitoring

Alerts

Applications

Assets

LLM Providers

Logging

Message Broker

Misc

Names

Notifications

Network Discovery

Reports

Telemetry

Timeseries

### Local LLM

#### API URL

Base URL for the local LLM API endpoint (e.g. http://localhost:11434)

https://aitest.interhost.it/v1/chat/complet

#### API Token

Authentication token for the local LLM API, if required

.....

#### Model Name

Local model name to use

Qwen3.6-27B

### Anthropic

#### API URL

Base URL for the Anthropic API endpoint

https://api.anthropic.com/v1/messages

#### API Token

API key for the Anthropic API



# Chat Interface

The screenshot displays the nAnalist chat interface. On the left is a dark sidebar with navigation icons and labels: Dashboard, Monitoring, Alerts, Flows, nAnalist (highlighted), Hosts, Maps, Interface, Policies, Settings, Developer, and Help. The top status bar shows the interface is connected to 'enp1s0' with a speed of 700.70 bps and 1.60 Kbps. It also features several status indicators: 2 alerts, 14 flows, 10 hosts, 13 (1) interfaces, 10 maps, 210 policies, and the ntop logo. The main chat area has a 'HISTORY' tab and a 'Local LLM Qwen3.6-27B' model selector. Below the history is a '+ New Chat' button and tabs for 'All', 'Live Flows', and 'Historical'. The chat area contains a list of previous questions, each with a gear icon and a star icon. The main chat area is currently empty, displaying a large orange speech bubble icon and the text 'Ask nAnalist a question'. Below this are four example questions in rounded rectangular boxes: 'What are the top 5 hosts by traffic in the last hour?', 'Which network protocols are generating the most traffic?', 'List the top talkers by bytes sent in the last 24 hours', and 'Show me DNS query volume and top queried domains'. At the bottom of the chat area is an input field with the placeholder text 'Ask a question' and an orange 'Investigate' button. A small disclaimer at the bottom left of the chat area reads: 'nAnalist can make mistakes. Always verify critical information independently'.

# Why is nAnalyst Different? [1/2]

- Leverages nDPI as high fidelity data source
- Works with anthropic, openAI and any openAI compatible inference server
- Privacy and explainability in mind, **no black box, interpretable results**
- Cost tracking and audit log ready (SOC2 style evidence)
- Domain specific tools - not a simple LLM wrapper

# Why is nAnalyst Different? [2/2]

Differentiator	nAnalyst	Generic Agent Assistant
Purpose built tools	25+ network specific tools	Generic web search
Evidence transparency	Full reasoning steps	Black box
Audit trail	Actions taken	None
Privacy	Fully on prem	Cloud only
Cost Tracking	Detailed Usage	N/A
Ntopng - nDPI native	Yes	No nDPI

# What does this mean for teams?

1. Investigation time from hours to seconds
2. Less alert fatigue
3. Less false positives and noise
4. Faster incident response

# Demo 1: Traffic Investigation

In seconds, users go from raw traffic to insights + next steps

1. **Question:** “Show me DNS query volume for each hour in the last day and top queried domains”
2. Agent collects data for DNS traffic, query to Clickhouse
3. Generates a line chart embedded in the chat
4. Reports a table of evidence and provides suggested next steps: “Investigate the two most contacted domains...”

# Demo 1: Traffic Investigation

Show me DNS query volume for each hour in the last day and top queried domains  
5/5/2026, 8:58:11 PM

### Hourly DNS Query Volume (Last 24h)

Time	Queries
May 04 18:00	6
May 04 19:00	1,124
May 04 20:00	1,124
May 04 21:00	600
May 05 00:00	650
May 05 03:00	650
May 05 06:00	600
May 05 09:00	650
May 05 12:00	1,196
May 05 15:00	650
May 05 18:00	700

DNS traffic over the last 24 hours shows steady volume between 600-800 queries/hour, with spikes at 19:00 (1,124) and 12:00 (1,196). The early hour of 18:00 yesterday was low (6 queries), likely due to partial data collection.

#### Top 10 Queried Domains

Domain	Queries
<a href="#">qubitowl.com</a>	9,016
<a href="#">encodeiq.ai</a>	5,442
<a href="#">grafana.com</a>	1,170
<a href="#">interhost.it</a>	402
<a href="#">ellio.tech</a>	384
<a href="#">microsoft.com</a>	360
<a href="#">azure.com</a>	354
<a href="#">github.com</a>	222
<a href="#">datadoghq.com</a>	192
<a href="#">npmjs.org</a>	162

[qubitowl.com](#) and [encodeiq.ai](#) dominate with 14,458 combined queries. Investigate whether these are expected services or potential beacons/domains of interest.

#### SUGGESTED NEXT STEPS

- Investigate [qubitowl.com](#) flows check source hosts and alert status for this high-volume domain.
- Investigate [encodeiq.ai](#) flows verify legitimacy and correlate with any active alerts.

5/5/2026, 8:59:00 PM · 49s · 26 tok/s

Ask a question Investigate

⚠ nAnalyst can make mistakes. Always verify critical information independently



# Evidence Logs - No Black Box

**qubitowl.com** and **encodeiq.ai** dominate with 14,458 combined queries. Investigate whether these are expected services or potential beaconing/domains of interest.

## SUGGESTED NEXT STEPS

→ **Investigate qubitowl.com flows** check source hosts and alert status for this high-volume domain.

→ **Investigate encodeiq.ai flows** verify legitimacy and correlate with any active alerts.

5/5/2026, 8:59:00 PM · 49s · 26 tok/s

✓ Show Evidence

^ Hide Evidence

### Q Describe Table

The user wants two things: DNS query volume per hour (line chart) and top queried domains (pie chart). I need to inspect the flows table schema first to understand the available columns for DNS traffic filtering and time aggregation.

### Q Chart

I have the schema. DNS traffic is identified by L7\_PROTO=5. I need two visualizations: 1) Hourly DNS query volume (line chart), 2) Top queried domains (pie chart). Per the rules, I can only show ONE chart per response. I will show the line chart for hourly volume first, and then query the top domains separately to display as a table.

```
SELECT formatDateTime(toStartOfHour(FIRST_SEEN), '%Y-%m-%d %H:%i') AS x, count() AS y FROM flows WHERE L7_PROTO = 5 AND F
```

### Q Query

I have the hourly DNS volume data. Now I need to get the top queried domains to display as a table. I will use the DOMAIN\_NAME column from the flows table, filtering for DNS (L7\_PROTO=5) in the last 24 hours.

```
SELECT DOMAIN_NAME AS domain, count() AS queries FROM flows WHERE L7_PROTO = 5 AND FIRST_SEEN >= now() - INTERVAL 24 HOUF
```

# Demo 2: Active Monitoring Automation

## AI doesn't just analyze - it acts

1. I want to monitor if there is connectivity to https://qubitowl.com
2. **Tool:** add active monitoring script
3. A new monitoring hook is added, monitoring HTTPS connectivity every minute. (ICMP, HTTP(S), speediest, throughput tests available)

The screenshot displays a chat interface with the following elements:

- User Query (orange bubble):** "I want to monitor if there is connectivity towards qubitowl.com, https, latency more than 100ms, how can i do it?" (10:39 AM)
- AI Response (white bubble):** "Active monitoring for HTTPS latency to qubitowl.com has been configured. An alert will trigger every minute if the response time exceeds 100 ms." Includes a button "Open Active Monitoring" and metadata "10:39 AM · 13s · 58.5 tok/s".
- Footer:** "Ask a question" input field, "Investigate" button, and a disclaimer: "nAnalyst can make mistakes. Always verify critical information independently".

# Demo 3: Policy Generation

**Idea -> detection -> alerts in seconds**

1. No SSH for host 192.168.2.38 during business hours
2. Agent generates an SQL query that is executed periodically (1 or 5 min, 1h, daily)
3. Generate an alert if the policy is violated
4. Agent executes the SQL query, provides interpretability and registers the alert
5. Complex query such as: traffic 2x hourly baseline and similar are possible

# Demo 3: Policy Generation

- Dashboard
- Monitoring
- Alerts
- Flows
- nAnalyst
- Hosts
- Maps
- Interface
- Policies

enp1s0 2 6 11 (1) 5 10 70 ntop

### Generate Policy from Natural Language

**DESCRIBE YOUR POLICY**  
Tell the AI what traffic to monitor and what constitutes a violation. Be specific about hosts, ports, protocols, time windows, or traffic volumes.

**LLM PROVIDERS**

**FREQUENCY**      Search:

Acti...	Periodicity	Status	Name	Description
	hourly	Active	192.168.2.38: no SSH 09:00-18:00	Detects any SSH (L7_PROTO=92) traffic initiated by or received by host 192.168.2.38 between 09:00 and 18:00 UTC.
	hourly	Active	192.168.2.38: no SSH 01:00-09:00	Detects any SSH (L7_PROTO=92) traffic initiated by or received by host 192.168.2.38 between 01:00 and 09:00 UTC.

# Review Policy

 Generate Policy from Natural Language

[← New](#)

**Policy Generated**  
Review the generated SQL query and check frequency.

**POLICY NAME**

192.168.2.38: SSH during business hours

**RUN EVERY**

Hourly

**ALERT SCORE**

Error (100)

**DESCRIPTION**

Detects any SSH (L7\_PROTO=92) traffic initiated by host 192.168.2.38 between 09:00 and 18:00 UTC, assuming standard business hours.

**ALERT MESSAGE**

Host 192.168.2.38 violated policy: SSH during business hours

**SQL QUERY**

NON EMPTY RESULT IS A VIOLATION

high

```
SELECT if(IPV4_SRC_ADDR != 0, IPv4NumToString(IPV4_SRC_ADDR), IPv6NumToString(IPV6_SRC_ADDR)) AS src_ip, if(IPV4_DST_ADDR != 0, IPv4NumToString(IPV4_DST_ADDR), IPv6NumToString(IPV6_DST_ADDR)) AS dst_ip, L7_PROTO AS l7_proto, count() AS query_count, formatDateTime(max(FIRST_SEEN), '%Y-%m-%d %H:%i:%S', 'UTC') AS last_seen FROM flows WHERE FIRST_SEEN >= now() - INTERVAL 1 HOUR AND IPV4_SRC_ADDR = IPv4StringToNum('192.168.2.38') AND L7_PROTO = 92 AND toHour(FIRST_SEEN) BETWEEN 9 AND 18 GROUP BY src_ip, dst_ip, l7_proto LIMIT 100
```

**EXPLANATION**

This policy monitors outbound SSH traffic (L7\_PROTO=92) initiated by host 192.168.2.38. It triggers an alert if any SSH flows are detected between 09:00 and 18:00 UTC, which are defined here as business hours. Periodicity is hourly to balance detection latency with noise reduction. Criticality is high because unauthorized SSH access during business hours may indicate compromised credentials or policy violation. Score 100 (error) reflects significant risk requiring prompt investigation.

 Save Policy

# Alert Generation

Detect any policy violation, inbound and outbound to the specified host(s)

Alerts Explorer | All **Host** Interface Local Networks AS SNMP Flow MAC Address System Active Monitoring User

Alerts Last Day 04/05/2026 19:06 → 05/05/2026 19:06 Apply

Filters

Top Hosts Top Alerts Top Mitre Tactics 10

Actions	Date / Time	Score	Category	Alert	Host	Description
⋮	13:01:45	150	🛡️	AI Policy	192.168.2.153 <span>L</span> <span>🖥️</span>	Host 192.168.2.153 initiated 10 SSH flows to 192.168.2.38 during business hours (9am-6pm). Last seen: 2026-05-05 12:5
⋮	11:00:33	150	🛡️	AI Policy	develgd <span>L</span> <span>🖥️</span>	Host 192.168.2.38 initiated 2 SSH flows to 192.168.2.61 during business hours (9am-6pm). Last seen: 2026-05-05 10:27:

# Demo 4: Flow Explanation

**Scenario:** Click “Explain” on a suspicious flow

1. Live flows page -> Explain
2. Ask questions about a flow
3. Agent fetches flow context, has access to all other tools, DB
4. Agent generates analysis/answer

# Analysis Pipelines

**Use case:** Repeated investigations with different filters.  
Compliance analysis

**Solution:** Grafana style analysis pipelines

1. Describe an analysis pipeline, 1+ stages
2. LLM generates custom SQL queries with parametrized filter values
3. Editable filters

# Analysis Pipelines: Example

**NIS2 policy enforcement.** I have a nis2 policy to enforce

1. Return all IPs that did not do any NTP traffic in the last 24h
2. For hosts that instead did NTP traffic, i want to know the destination server, the server ASN and total volume

# Analysis Pipelines: Example

The screenshot displays the nAnalyst interface with a sidebar on the left containing navigation options: Dashboard, Monitoring, Alerts, Flows, nAnalyst (highlighted), Interface, Policies, Settings, Developer, and Help. The main content area is titled "Generate Playbook from Natural Language" and shows a "Playbook Generated" status with a green checkmark and a duration of 34.7s. Below this, a table lists the generated playbook:

NAME	DESCRIPTION
NTP Compliance Audit	Identify hosts violating NTP policy and analyze compliant hosts' NTP traffic.

An informational message states: "This playbook enforces NIS2 NTP policy requirements. Stage 1 identifies all internal hosts that have not initiated any NTP traffic in the last 24 hours (potential compliance violations). Stage 2 analyzes the hosts that \*did\* initiate NTP traffic, showing which NTP servers they contacted, the server's ASN, and the total volume of NTP data exchanged."

The "STAGES" section shows 2 stages. Stage 1, titled "s1 Hosts Without NTP Traffic", is expanded to show its description and SQL query:

1 s1 Hosts Without NTP Traffic 1/2

Finds all unique source IPs that had ANY traffic in the last 24h but did NOT initiate NTP (L7\_PROTO=9) requests.

```
SELECT
  IPv4NumToString(flows.IPV4_SRC_ADDR) AS host_ip,
  count() AS total_flows,
  formatReadableSize(sum(flows.TOTAL_BYTES)) AS total_bytes
FROM flows
WHERE flows.INTERFACE_ID = {{ifid}}
  AND flows.FIRST_SEEN >= toDateTime('{{epoch_start}}')
  AND flows.FIRST_SEEN <= toDateTime('{{epoch_end}}')
  AND flows.IPV4_SRC_ADDR != 0
  AND flows.IPV4_SRC_ADDR NOT IN (
    SELECT DISTINCT IPV4_SRC_ADDR
    FROM flows
    WHERE INTERFACE_ID = {{ifid}}
      AND FIRST_SEEN >= toDateTime('{{epoch_start}}')
      AND FIRST_SEEN <= toDateTime('{{epoch_end}}')
      AND L7_PROTO = 9
  )
GROUP BY flows.IPV4_SRC_ADDR
ORDER BY total_bytes DESC
```

# nAnalyst - MCP Server

**nAnalyst is also accessible as an MCP server running on ntopng**

- No additional installation required
- Add MCP credentials in Claude Code, cursor etc..
- Any MCP client can connect and interact with ntopng
- Same tools available as nAnalyst
- Strong security based on ntopng server

# nAnalyst - MCP Server - Alert Investigation

## Security Findings

### ● HIGH – Unidentified Raspberry Pi persistently hammering SSH

Source: 192.168.2.153 – MAC B8:27:EB:4D:44:C8 → Raspberry Pi Foundation

This device is making continuous SSH connections to .38:22 every 1–3 minutes for 24 hours straight, 24/7 with no gaps:

- 921 alerts generated – all as client, all targeting .38
- Contacts only 1 unique server (exclusively .38)
- SSH flows: 18 sessions, ~15 MB transferred
- Also hitting port 3000 on .38 (dev web server) with FLOW\_RISK=4128, SCORE=70
- Some SSH flows carry FLOW\_RISK=34359738368 (bit 35 – SSH weak/obsolete algorithm) with SCORE=10
- 3 behavioral score anomalies on the Pi itself
- Pi has no hostname, no label, no DNS activity – entirely unidentified
- attacker\_hits: 28 – flood flagged

**Assessment:** This pattern – an unidentified headless device making automated, high-frequency SSH connections to a single target, 24/7 – is a textbook persistent access mechanism signature. It could be legitimate (a CI/CD runner, backup agent, or monitoring probe), but it is **not documented, not labeled, and triggering 921 alerts**. This requires immediate investigation.

#### Action required:

1. Physically identify the Pi – check DHCP leases for .153
2. Determine what process on the Pi is initiating SSH (check ~/.ssh/authorized\_keys on .38 for unexpected keys)
3. If legitimate: label the device, whitelist, create an AI policy for expected behavior
4. If unauthorized: isolate .153 immediately

# What nAnalyst Is Evolving Into

1. “What changed?” Over time
2. Alerts analysis and prioritization (root cause analysis):
  1. Reduce alert fatigue
  2. Highlight relevant alerts worth investigating
3. Behavioral baselines

# Observability and Visibility

## 1. Usage Stats

1. Monitor historical usage in tokens (and \$) for each model and user
2. Breakdown agent usage: tool call, response generation etc

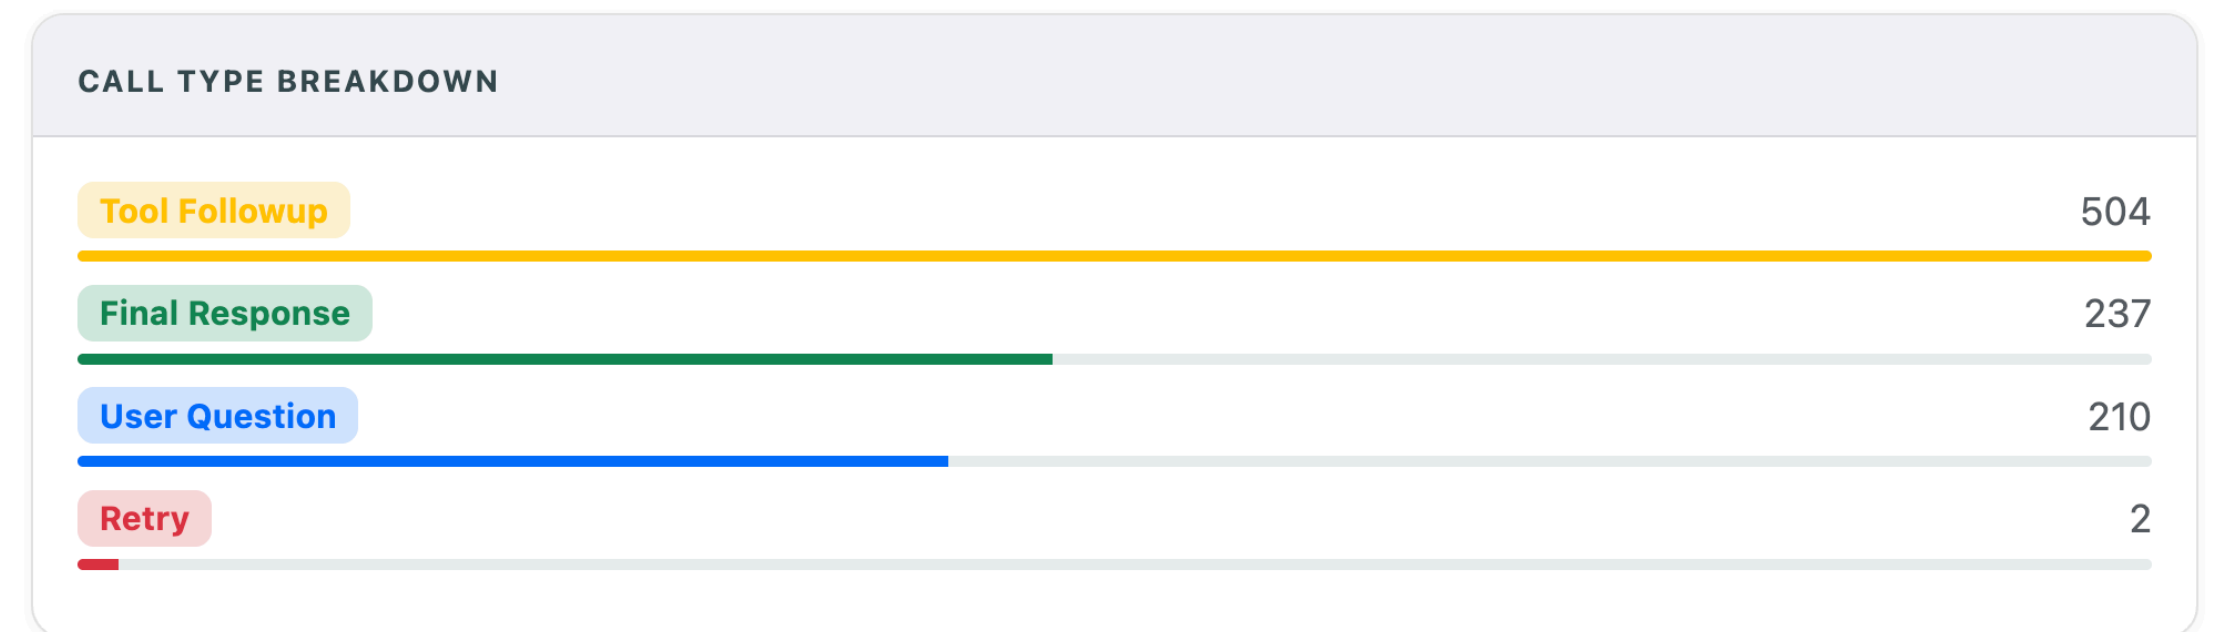
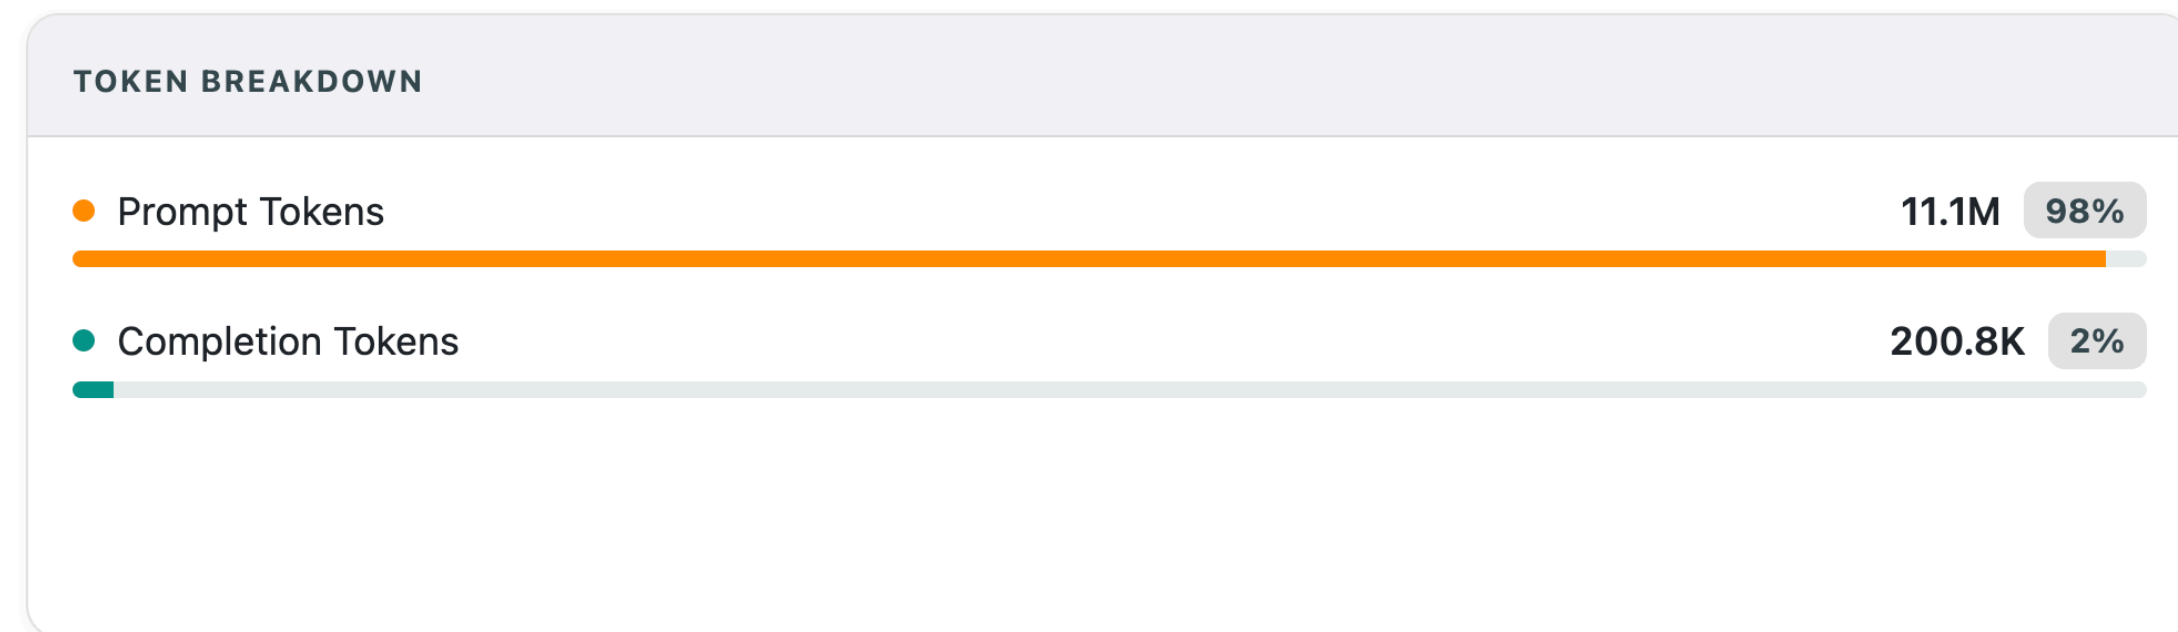
## 2. Action Audit

1. Monitor which user added an active monitoring script, created a policy
2. Blame the user or the agent for a given action
3. See action, who initiated (agent suggestion or user request)
4. Full visibility on the tool call content

# Usage Stats

**TIME RANGE** | 1h | 6h | 24h | 7d | **30d** |
 **PROVIDER** | All Providers |
 **MODEL** | All Models |
 **USER** | All Users |
 [Back To Chat](#) | [Add Model Cost](#) | [Refresh](#)

<b>953</b> Total LLM Calls	<b>11.3M</b> Total Tokens	<b>6.1s</b> Average Response Time	<b>153</b> Unique Chats
<b>\$1.04</b> Total Cost	<b>\$0.8485</b> Input Cost	<b>\$0.1870</b> Generation Cost	<b>\$0.5379</b> Tool Usage Cost



Usage By Model | Usage By User | Model Prices | 10 | [Grid](#) | [Refresh](#) | [Eye](#) | Search:

Provider	Model	Ca...	Prompt Tokens	Completion Tokens	Total Tokens	Avg Response	Max Response	Input Cost	Generation Cost	Total Cost
Anthropic	claude-haik...	106	848.5K	37.4K	<b>885.9K</b>	0ms	0ms	\$0.00	\$0.00	\$0.00
Anthropic	claude-opus...	10	138.0K	1.9K	<b>139.9K</b>	0ms	0ms	\$0.00	\$0.00	\$0.00
Local LLM	Qwen3.6-35B...	36	336.8K	5.6K	<b>342.4K</b>	1.1s	3.9s	\$0.00	\$0.00	\$0.00
Local LLM	gemma4-E4B	3	43.2K	485	<b>43.7K</b>	1.4s	2.3s	\$0.00	\$0.00	\$0.00

# Model Cost

1. Choose a model that was used or specify a custom model
2. Add price in input and output per 1M tokens
3. Monitor costs in the page from the previous slide

**Add Model Price**

**Provider** Anthropic

**Model**   
 ✓ — —   
 claude-haiku-4-5   
 claude-opus-4-5   
 Custom model name

**Input Price (1M Tokens)**   
 Cost per 1M input (prompt) tokens in USD   
 \$ 0 / 1M tokens

**Output Price (1M Tokens)**   
 Cost per 1M output (completion) tokens in USD   
 \$ 0 / 1M tokens

Save

# Audit Usage

enp1s0 2 1 11 7 11 39 ntop

User: All Triggered By: All Tool: All Reset Last Month: 05/04/2026 20:23 → 05/05/2026 20:23 Apply 10 Search:

Timestamp	User	Triggered By	Tool	Action	Content
Yesterday, 17:23:52	admin	LLM Agent	Add Active Monitor	Added active monitor: https://www.encodeiq.ai:443 (h...	{"measurement":"http","host":"https://www.encodeiq.ai:443","th...
Yesterday, 17:23:11	admin	LLM Agent	Add Active Monitor	Added active monitor: https://www.encodeiq.ai:443 (h...	{"measurement":"https","host":"https://www.encodeiq.ai:443","i...
Yesterday, 17:15:33	admin	Direct User Action	Create AI Policy	Created policy #22: 192.168.2.38: no SSH 01:00-09:00	{"description":"Detects any SSH (L7_PROTO=92) traffic initiated
Yesterday, 13:46:34	admin	Direct User Action	update_ai_policy	Updated policy #21: 192.168.2.38: no SSH 09:00-18:00	{"policy_id":"21","periodicity":"hourly","alert_description":"I...
Yesterday, 13:25:26	admin	Direct User Action	Create AI Policy	Created policy #21: 192.168.2.38: no SSH 09:00-18:00	{"custom_score":50,"periodicity":"hourly","name":"192.168.2.38
Yesterday, 10:37:38	admin	Direct User Action	Create AI Policy	Created policy #20: 192.168.2.38: >1MB TLS to qubito...	{"description":"Monitors TLS traffic from host 192.168.2.38 to
Fri, 06:56:13	admin	LLM Agent	Add Active Monitor	Added active monitor: 192.168.2.38 (icmp)	{"measurement":"icmp","host":"192.168.2.38","ifid":1,"threshol...
Thu, 07:38:47	admin	LLM Agent	Create AI Policy	Created policy #19: 192.168.2.38: >10MB to qubitowl.c...	{"custom_score":50,"periodicity":"5min","alert_description":"H...
Thu, 07:32:49	admin	LLM Agent	Add Active Monitor	Added active monitor: https://qubitowl.com:443 (http)	{"measurement":"http","host":"https://qubitowl.com:443","thresl...

Showing page 1 of 1: total 9 rows 1

ntopng Enterprise XXL v.6.7.260423 (Ubuntu 24.04.4 LTS) | © 1998-26 - ntop 20:24:07 +0000 UTC | Uptime: 1 Day, 09:50:12



# Timeline

1. **Feedback collection:** starting today
2. **Open Beta:** June 22nd
3. **Release:** Q3 26

**Requirements:** ntopng Enterprise M and above  
+ clickhouse



**nAnalyst is not yet another dashboard.**  
It's the layer that finally explains your network.

For any questions or custom requests: [gabriele@ntop.org](mailto:gabriele@ntop.org)