

Challenges of self-hosting services

Tomas Hlavacek, CSNOG 2024

Is Cloud Exit a trend?



LATEST NEWS ARTIFICIAL INTELLIGENCE BIG DATA AND ANALYTICS

Home > Cloud

Cloud IT Management

Is Cloud Computing Worth the Cost?

If daunting cloud costs make you consider simply using your data center, realize that the cloud isn't just about cost

By **David Linthicum** - April 13, 2022



eWEEK content and product reviews are independent. We may make money from our partners. [Learn More.](#)

IT costs are complex and traditional IT, better know



CLOUD

Why your cloud computing costs are so high – and what you can do about them

 **SPECIAL REPORT: WHAT'S NEXT IN CLOUD BY PAUL GILLIN**

Small mistakes in the cloud can have big consequences.

Cloud Exit is indeed a reality



DAVID HEINEMEIER HANSSON

December 19, 2023

The Big Cloud Exit FAQ

Just over a year ago, we announced [our intention to leave the cloud](#). We then shared [our complete \\$3.2 million cloud budget for 2022](#), and the fact that we were going to build our own tooling rather than pay for [overpriced enterprise service contracts](#). The mission was set!

Články Zprávičky Fórum Podpořte Root Školení Galerie Kalendář Root do mailu RSS

Ad

Root.cz » Server » Odpovědi na nejčastější dotazy k našemu odchodu z předraženého cloudu

Odpovědi na nejčastější dotazy k našemu odchodu z předraženého cloudu

REDAKCE | 4. 1. 2024 | Doba čtení: 10 minut

143 NOVÝCH NÁZORŮ



Cloud Exit highs and lows

- **A: “We have 90% of instances in only four AWS regions”**
- **B: “Let’s build four PoPs / data centers in similar locations. What can go wrong?”**
- **A: “We had VPC peerings and transit gateways to link the micro-services.”**
- **B: “No problem, we just build fully-meshed IPsec tunnels between the on-prem sites. We could use a few OpnSense VMs for that. It will be fine...”**

Apps may not like it...

first-incident-responder-team



[Redacted] Dec 24th at 2:39 PM

Networking issues in a couple of DCs?

first-incident-responder-team



[Redacted] Dec 15th, 2023 at 8:34 AM

Sounds like "networking" issue to me because I can reproduce it inconsistently (edited)

e[Redacted]



[Redacted] Dec 14th, 2023 at 6:47 AM

Assuming this means it's a networking issue, @**[Redacted]** do you think there is anything we can do?

What can go wrong? Part 1

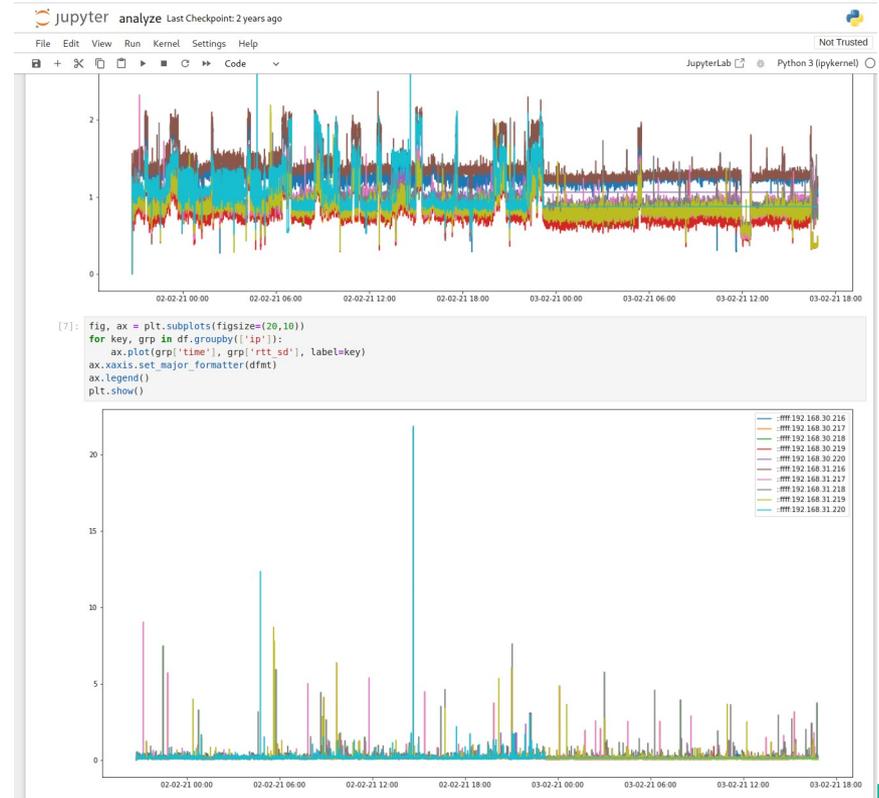
- **Proxmox cluster split to multiple DCs**
- **Proxmox uses Corosync for “membership management”**
- **Aggressive timeouts (1s) in Totem protocol**

- **Suspected network issues needs a debugging tool**
- **fping? hping3?**

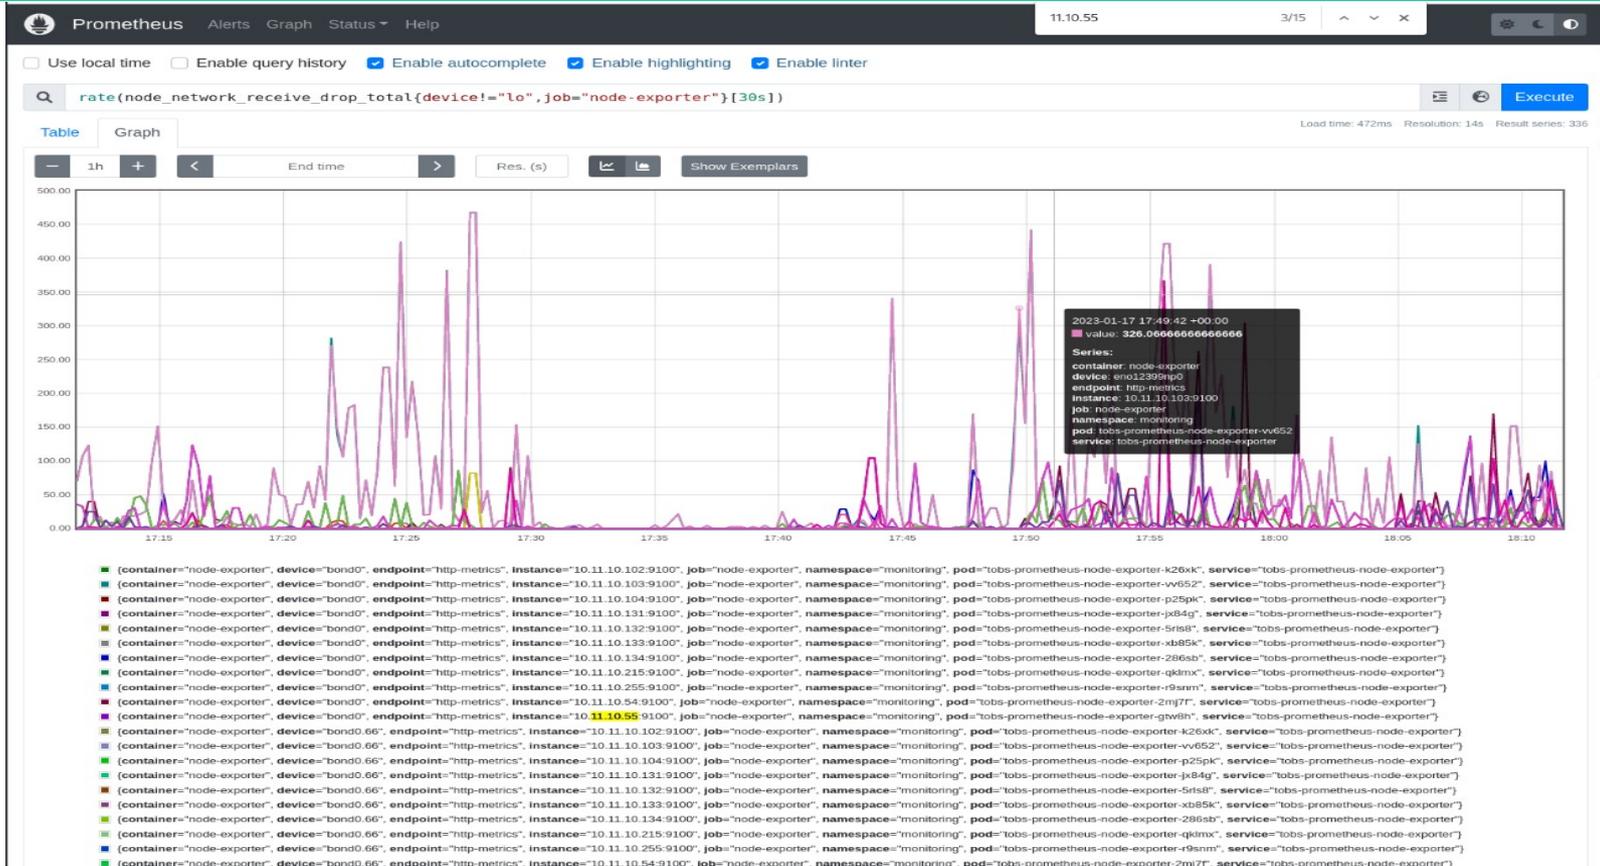
Part 1 happy end

- **fping = ICMP \neq TCP/UDP**
- **hping3 can do UDP**
- **Data science tools :-)**

- **<https://github.com/tmshlvck/udpsmoke>**



WCGW2: Fun with apps in K8s



Fun with apps in K8s – problem found?

```
root@[REDACTED]:/var/log# grep buffer syslog.1
Jan 17 16:06:50 zh1-dll15 kernel: [512287.300913] bnxt_en
0000:a1:00.1 ens2f1np1: RX buffer error d370004
Jan 17 16:06:50 zh1-dll15 kernel: [512287.300915] bnxt_en
0000:a1:00.1 ens2f1np1: RX buffer error d370004
Jan 17 16:06:50 zh1-dll15 kernel: [512287.300923] bnxt_en
0000:a1:00.1 ens2f1np1: RX buffer error d370004
Jan 17 16:06:50 zh1-dll15 kernel: [512287.300935] bnxt_en
0000:a1:00.1 ens2f1np1: RX buffer error d370004
...
```

Continuous (high-frequency) monitoring

- **gRPC based Telemetry**
- **SNMP :-)**
- **Icinga2**
- **SmokePing**
- **Prometheus Blackbox Exporter**



Meet Telephant in the room

- **Continuous high-frequency mesh probing**
- **Prometheus exporter**
- **YAML traces → Publish to web and share URL only**



Continuous high-frequency monitoring



What can go wrong – part 3: MTU issues

```
2: th@pes:~/Sync/projects/telephant
```

Peer	Tx	Rx	RTT (ms)	Lost	0o0	Crpt	PingRx
krtek6 2001:1ab0:7e1e:150::6	319 5.0/s	319 (100.0 %) 4.9/s	33.4 32.3 N/A	0	0	0	0
slon6 2a02:2a8:1:20::49c6	319 5.0/s	319 (100.0 %) 5.0/s	17.6 17.1 N/A	0	0	0	0
krtek4 62.109.150.6	319 5.0/s	319 (100.0 %) 5.0/s	24.5 24.9 N/A	0	0	0	0
slon4 93.188.73.198	319 5.0/s	319 (100.0 %) 5.0/s	18.1 17.6 N/A	0	0	0	0

```
0:01:04 | Tx 1412B | --- | q:Quit u:up d:down c:reConfig r:send Report l: Logs +/-: TxLen+/-
```

- UDP ping
- Easy to extend
- Extensive traces

- Rename: Telephant
- Prometheus exporter
- YAML traces upload

<https://github.com/tmshlvck/telephant>

What can go wrong – part 3: MTU issues

```
2: th@pes:~/Sync/projects/telephant
```

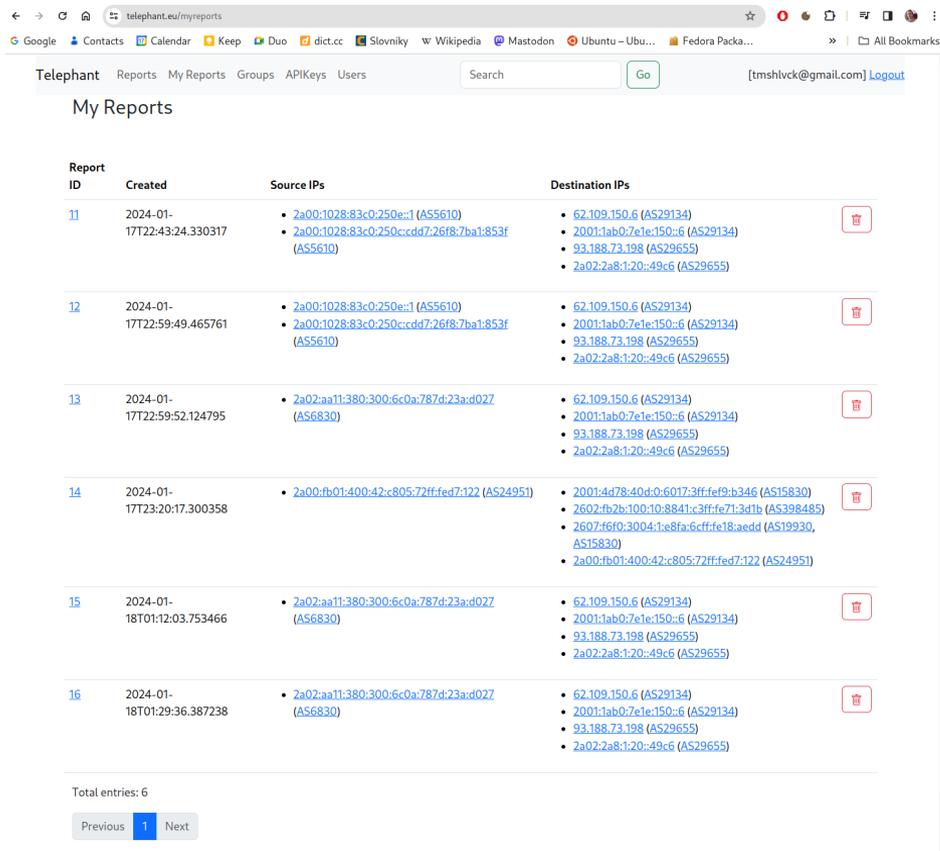
Peer	Tx	Rx	RTT (ms)	Lost	OoO	Crpt	PingRx
krtek6 2001:1ab0:7e1e:150::6	932 5.0/s	932 (100.0 %) 5.0/s	33.1 32.5 N/A	0	0	0	0
slon6 2a02:2a8:1:20::49c6	932 5.0/s	932 (100.0 %) 5.0/s	17.9 17.0 N/A	0	0	0	0
krtek4 62.109.150.6	932 5.0/s	417 (45.4 %) 0.4/s	25.8 27.3 N/A	492	0	0	0
slon4 93.188.73.198	932 5.0/s	393 (45.3 %) 0.1/s	18.4 19.0 N/A	514	0	0	0

```
0:03:07 | Tx 1413B | --- | q:Quit u:up d:down c:reConfig r:send Report l: Logs +/-: TxLen+/-
```

- DS-lite (CH)
- v4 works @ 1412B
- fails @ 1413B data
- + 20B Telephant header
- + 8B UDP header
- +20B IPv4 header
- +40B IPv6 header
- = 1501B

<https://www.telephant.eu/report?showreport=16>

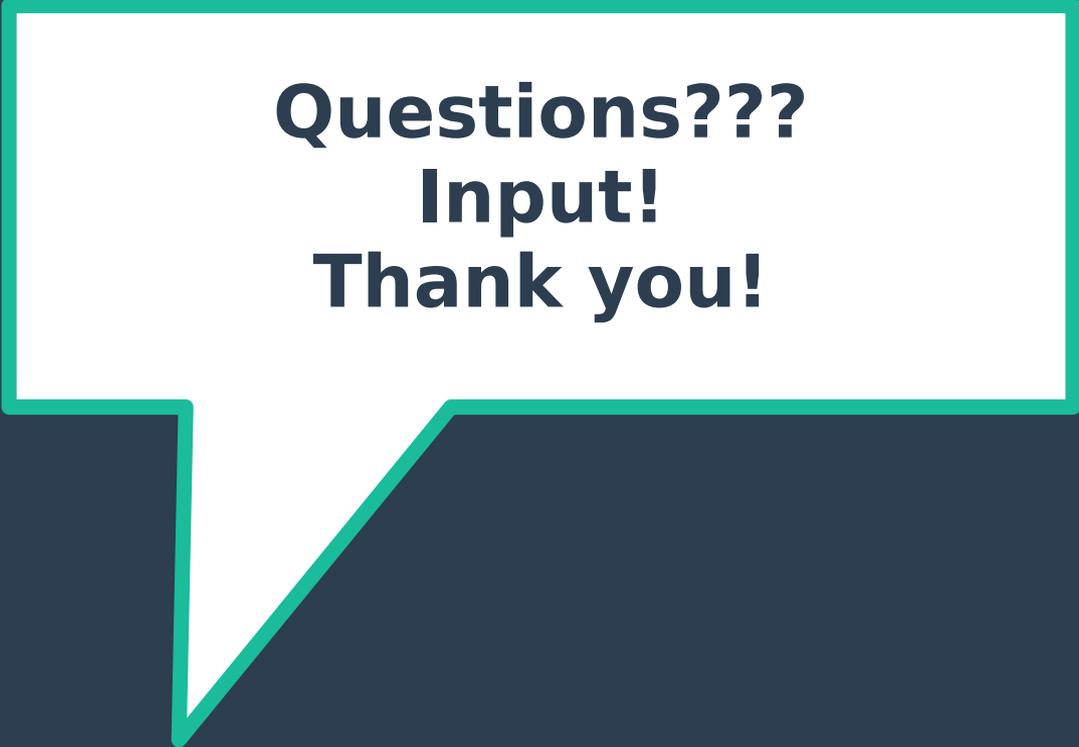
Telephant Web



The screenshot shows the Telephant Web interface. At the top, there is a navigation bar with links for Reports, My Reports, Groups, APIKeys, and Users. A search bar and a user profile link ([tmshlvck@gmail.com] Logout) are also present. The main content area is titled "My Reports" and displays a table of report entries. Each entry includes a Report ID, a Created timestamp, Source IPs, Destination IPs, and a delete icon. The table shows 6 entries, with the first one selected. Below the table, there is a pagination control showing "Total entries: 6" and "Previous 1 Next".

Report ID	Created	Source IPs	Destination IPs	
11	2024-01-17T22:43:24.330317	<ul style="list-style-type: none">2a00:1028:83c0:250e::1 (AS5610)2a00:1028:83c0:250c:cd47:26f8:7ba1:853f (AS5610)	<ul style="list-style-type: none">62.109.150.6 (AS29134)2001:1ab0:7e1e:150::6 (AS29134)93.188.73.198 (AS29655)2a02:2a8:1:20::49c6 (AS29655)	
12	2024-01-17T22:59:49.465761	<ul style="list-style-type: none">2a00:1028:83c0:250e::1 (AS5610)2a00:1028:83c0:250c:cd47:26f8:7ba1:853f (AS5610)	<ul style="list-style-type: none">62.109.150.6 (AS29134)2001:1ab0:7e1e:150::6 (AS29134)93.188.73.198 (AS29655)2a02:2a8:1:20::49c6 (AS29655)	
13	2024-01-17T22:59:52.124795	<ul style="list-style-type: none">2a02:aa11:380:300:6c0a:787d:23a:d027 (AS6830)	<ul style="list-style-type: none">62.109.150.6 (AS29134)2001:1ab0:7e1e:150::6 (AS29134)93.188.73.198 (AS29655)2a02:2a8:1:20::49c6 (AS29655)	
14	2024-01-17T23:20:17.300358	<ul style="list-style-type: none">2a00:fb01:400:42:c805:72ff:fed7:122 (AS24951)	<ul style="list-style-type: none">2001:d478:40d:0:6017:3ff:fe9:b346 (AS15830)2602:fb7b:100:10:8841:c3ff:fe71:3d1b (AS398485)2607:f60:3004:1e8fa:6c:fff:fe18:aedd (AS19930, AS15830)2a00:fb01:400:42:c805:72ff:fed7:122 (AS24951)	
15	2024-01-18T01:12:03.753466	<ul style="list-style-type: none">2a02:aa11:380:300:6c0a:787d:23a:d027 (AS6830)	<ul style="list-style-type: none">62.109.150.6 (AS29134)2001:1ab0:7e1e:150::6 (AS29134)93.188.73.198 (AS29655)2a02:2a8:1:20::49c6 (AS29655)	
16	2024-01-18T01:29:36.387238	<ul style="list-style-type: none">2a02:aa11:380:300:6c0a:787d:23a:d027 (AS6830)	<ul style="list-style-type: none">62.109.150.6 (AS29134)2001:1ab0:7e1e:150::6 (AS29134)93.188.73.198 (AS29655)2a02:2a8:1:20::49c6 (AS29655)	

- <https://telephant.eu>
- Report is static once uploaded
- Report has meta-data → searchable
- List of SRC and DST IPs included
- ping & traceroute to all targets
- “RRD”-style compressed stats
- 15s (24 hrs) + 5m (30 days) time series
- Aiming to provide a new debugging tool
- **Simplify/standardize issue reporting**



Questions???
Input!
Thank you!