"Everything is a Freaking DNS problem" dnsdist to the rescue

Hello!

I am Dominic

Working for cyon in Basel

Head of Software Engineering

@drdol

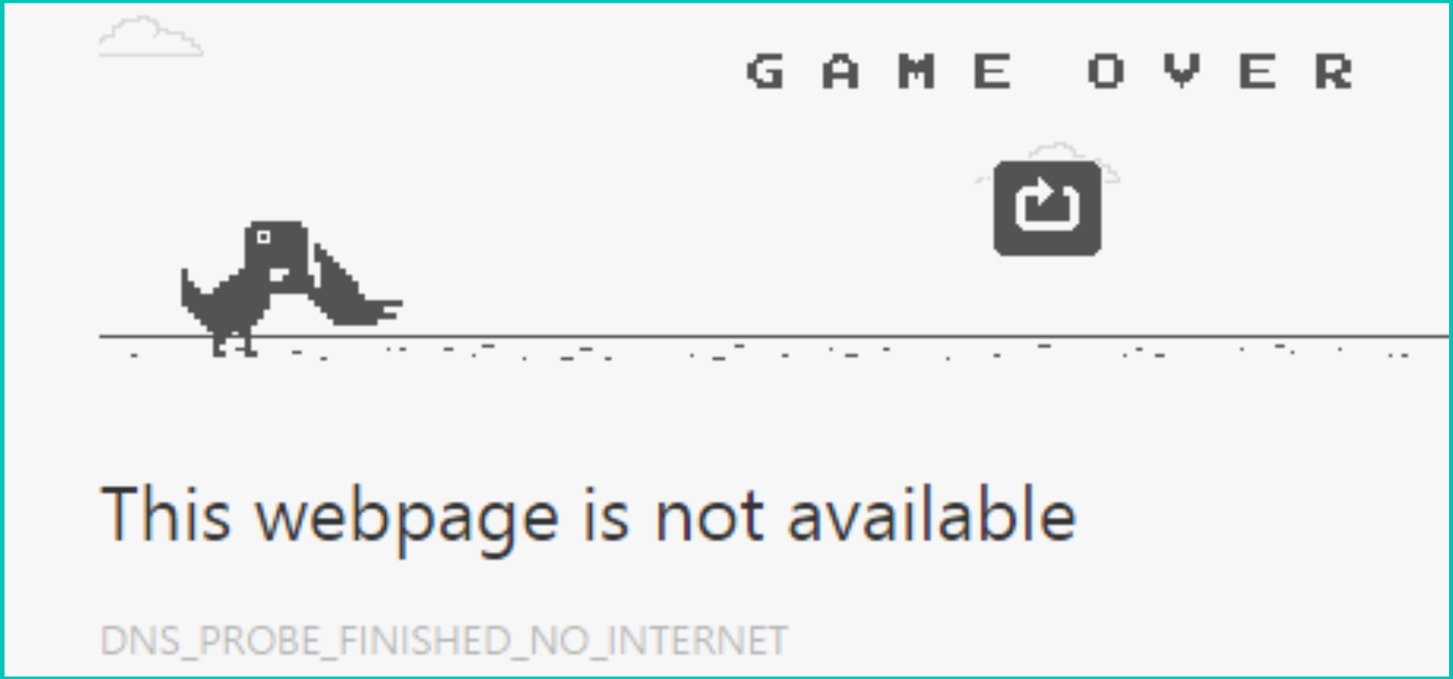

This webpage is not available

DNS_PROBE_FINISHED_NO_INTERNET

Preconditions

- Dependent on vendor
- Old software
- Software has bugs
  - NS delegation
  - Max size of TXT record
- No caching
- Minimal fault tolerance
- No stats or insights

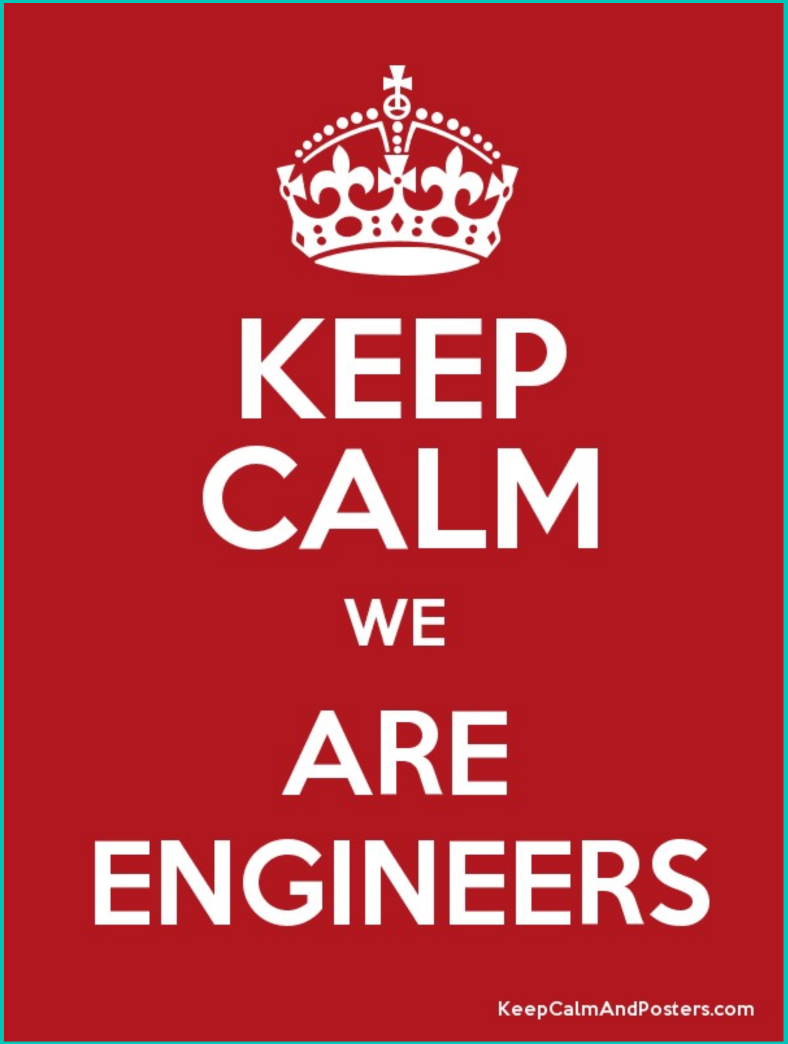
KEEP
CALM
WE
ARE
ENGINEERS
KeepCalmAndPosters.com

PORUCZ

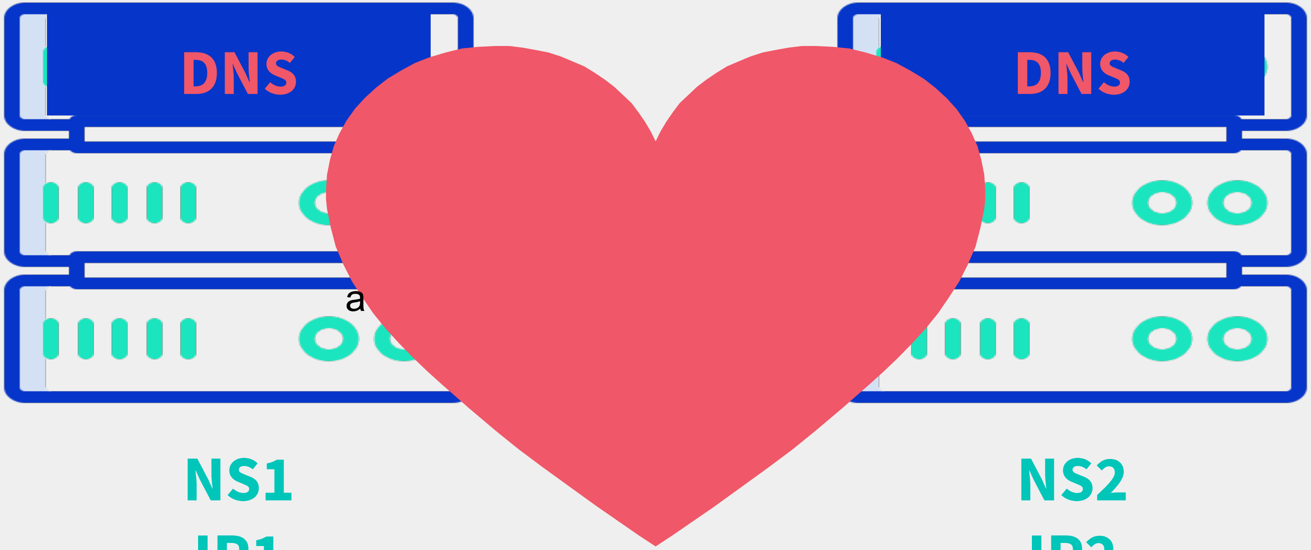

NS1
IP1
DNS 0.0.0.0:53

NS2
IP2
DNS 0.0.0.0:53

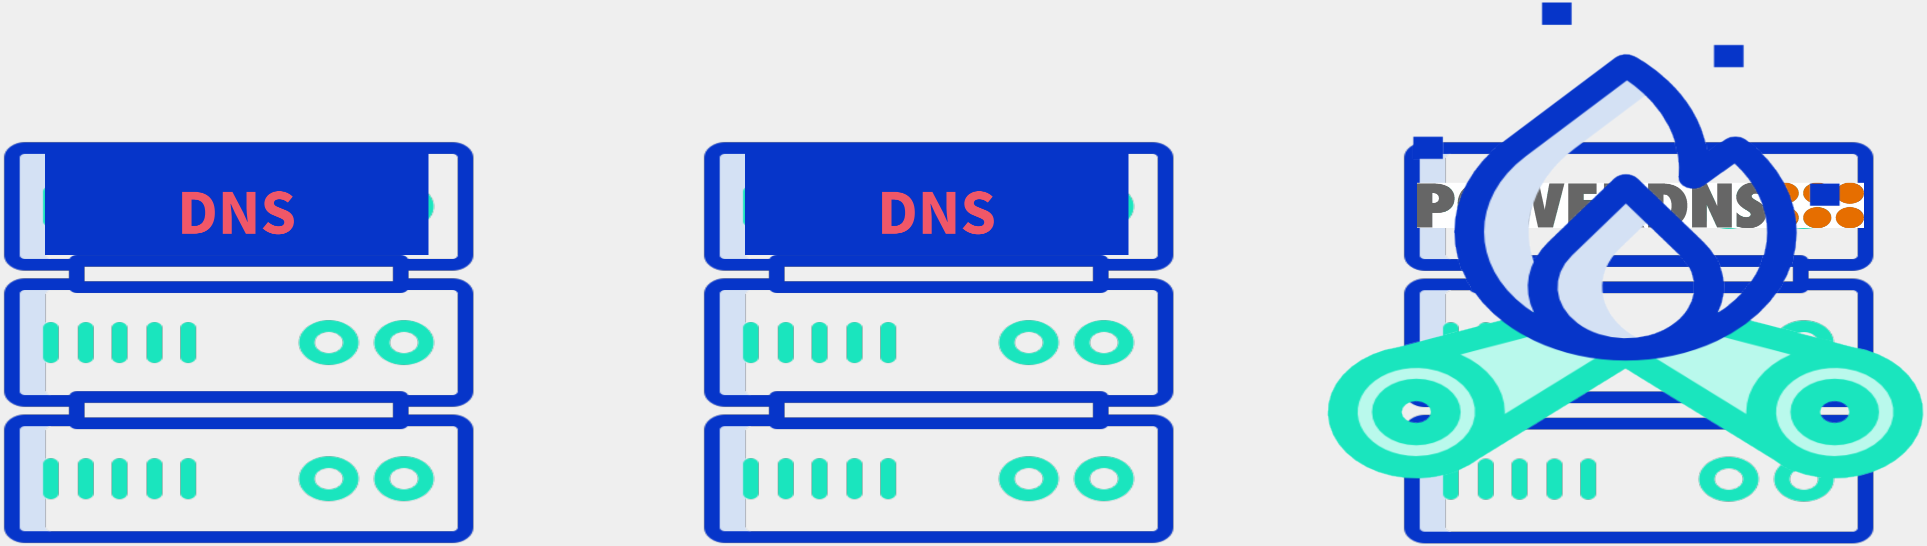

NS1
IP1
DNS 0.0.0.0:53

NS2
IP2
DNS 0.0.0.0:53

NS3
-
ALT 0.0.0.0:53

© Maura Teal 2015

dnsdist to the rescue

“

dnsdist *is a highly DNS-, DoS- and abuse-aware* ***loadbalancer****. Its goal in life is to* ***route traffic*** *to the best server, delivering top* ***performance*** *to legitimate users while shunting or* ***blocking*** *abusive traffic.*

Quickstart

- Add repository 'repo.powerdns.com'
- $ apt-get/yum install dnsdist
- $ sudo dnsdist -l 1.1.1.1 8.8.8.8 9.9.9.9
- $ dig swinog.ch @127.0.0.1 +short

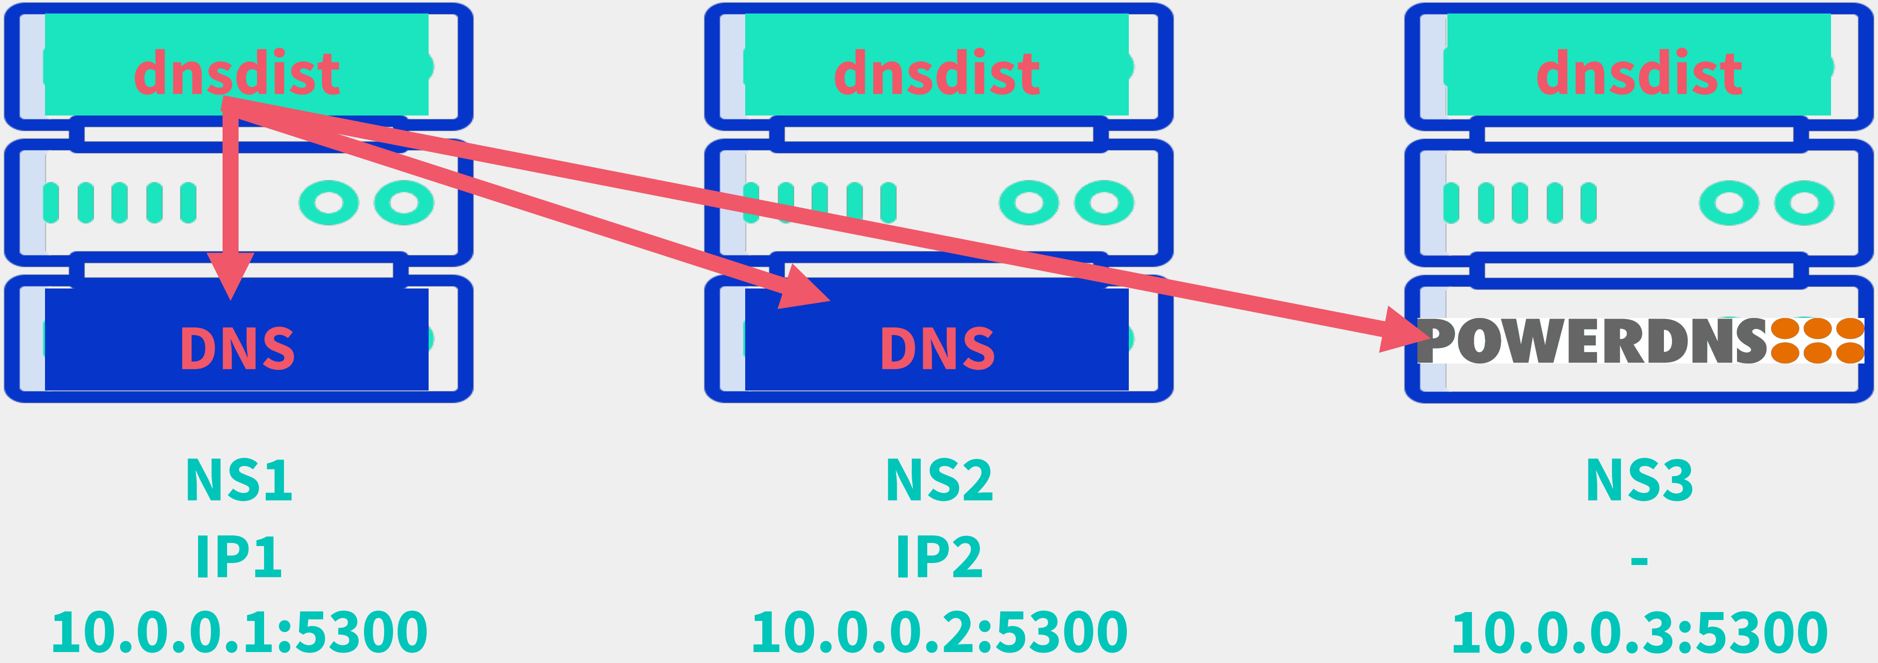
dnsdist
DNS
dnsdist
DNS
dnsdist
POWERDNS
NS1
IP1
10.0.0.1:5300
NS2
IP2
10.0.0.2:5300
NS3
-
10.0.0.3:5300

```
setLocal('0.0.0.0:53')
newServer({address='10.100.0.1:5300',order=1})
newServer({address='10.100.0.2:5300',order=2})
newServer({address='10.100.0.3:5300',order=3})
setServerPolicy(firstAvailable)
```

Batteries included

Console

Commandline console for live traffic inspection, controlling and configuring server

Performance

C++ 2011, Lua(JIT), lots of experience from running/developing DNS software

Statistics

View and ship stats/telemetry data to you preferred tool

Policies

Built-in policies for load balancing

Caching

Effective packet cache per pool

Rule/action engine

Drop, delay, redirect traffic based on rules

Policies/Pools/Rules/Actions

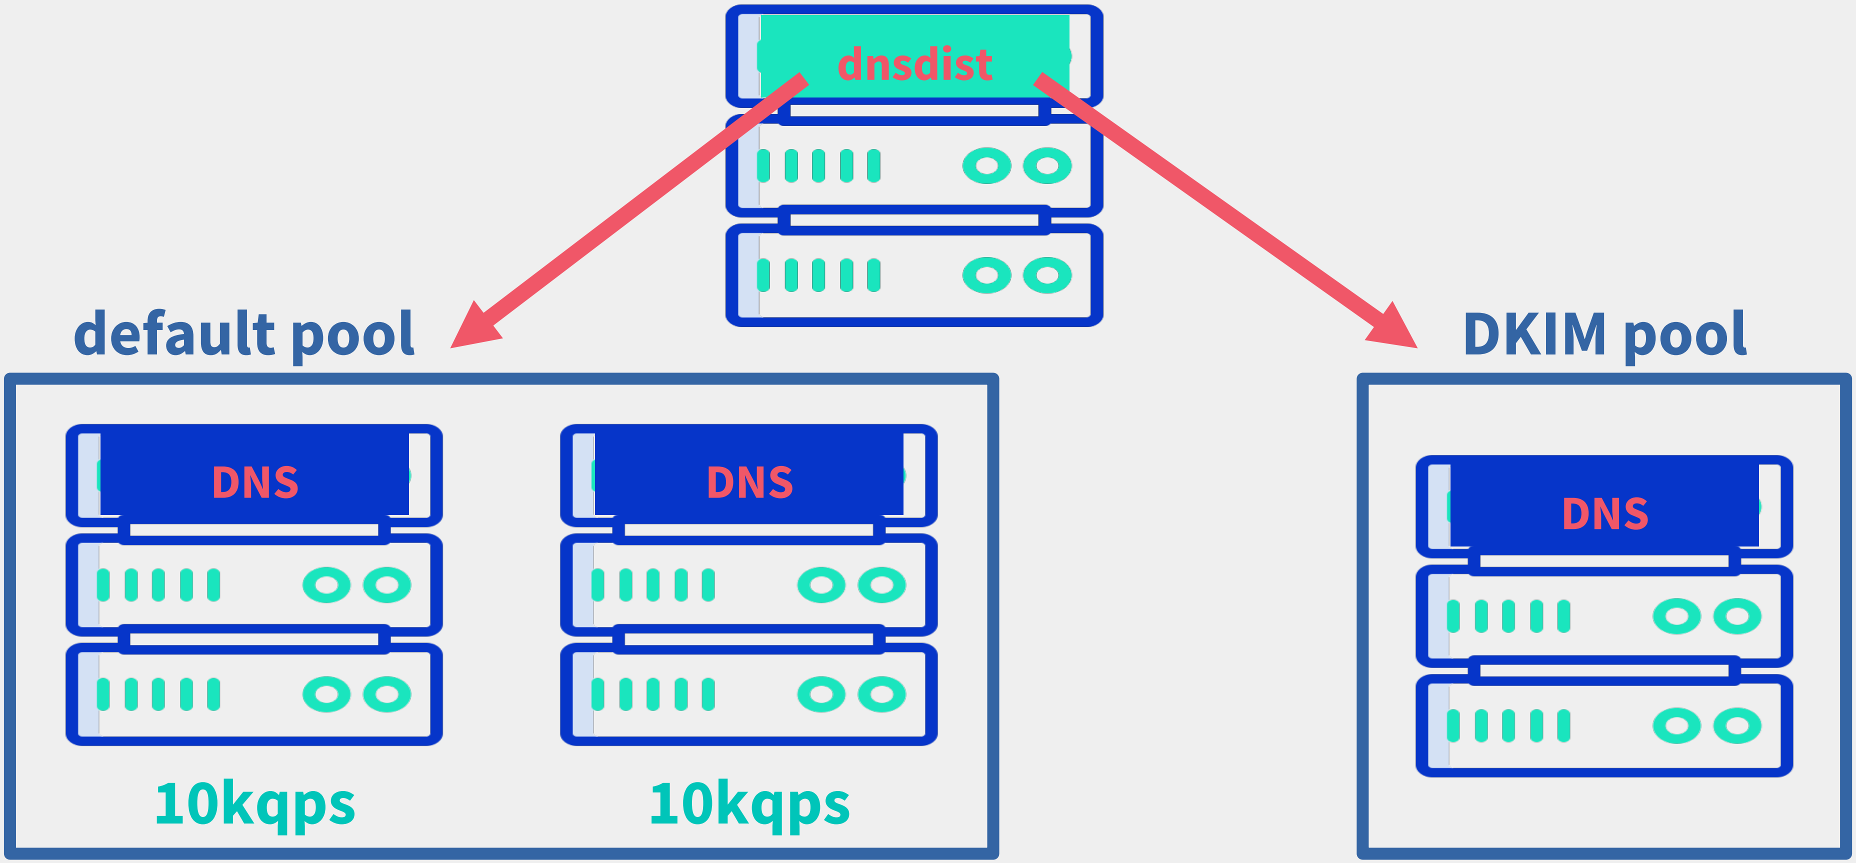

```
addAction(
  AndRule(
    {
      QTypeRule(dnsdist.TXT),
      RegexRule("\\._domainkey\\.")
    }
  ),
  PoolAction('dkim')
)
```

Stats and insights

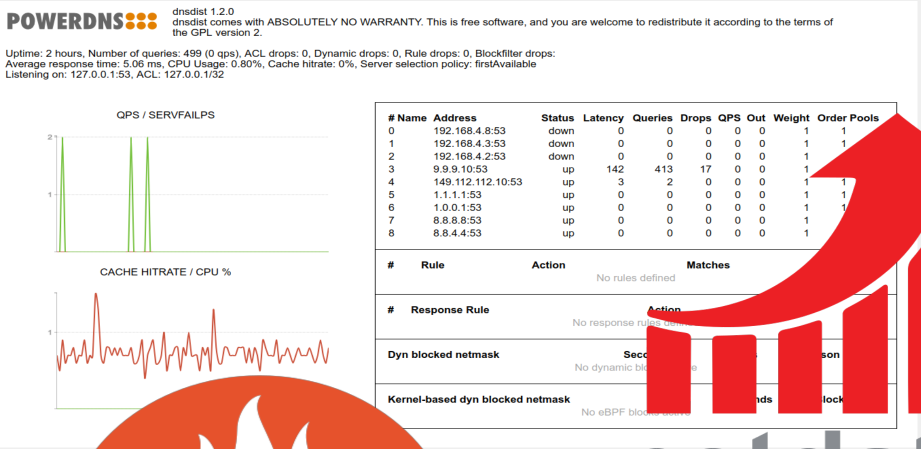

dnsdist 1.2.0
dnsdist comes with ABSOLUTELY NO WARRANTY. This is free software, and you are welcome to redistribute it according to the terms of the GPL version 2.

Uptime: 2 hours, Number of queries: 499 (0 qps), ACL drops: 0, Dynamic drops: 0, Rule drops: 0, Blockfilter drops:
Average response time: 5.06 ms, CPU Usage: 0.80%, Cache hitrate: 0%, Server selection policy: firstAvailable
Listening on: 127.0.0.1:53, ACL: 127.0.0.1/32

#	Name	Address	Status	Latency	Queries	Drops	QPS	Out	Weight	Order	Pools
0		192.168.4.8:53	down	0	0	0	0	0	1	1	
1		192.168.4.3:53	down	0	0	0	0	0	1	1	
2		192.168.4.2:53	down	0	0	0	0	0	1		
3		9.9.9.10:53	up	142	413	17	0	0	1		
4		149.112.112.10:53	up	3	2	0	0	0	1	1	
5		1.1.1.1:53	up	0	0	0	0	0	1	1	
6		1.0.0.1:53	up	0	0	0	0	0	1	1	
7		8.8.8.8:53	up	0	0	0	0	0	1		
8		8.8.4.4:53	up	0	0	0	0	0	1		

#	Rule	Action	Matches

No rules defined

#	Response Rule	Action

No response rules defi

Dyn blocked netmask	Seco	son

No dynamic blo e

Kernel-based dyn blocked netmask	nds	lock

No eBPF blocks active

SNMP
Simple Network Management Protocol

netdata

graphite

metronome.powerdns.com

Graphing as a Service

```
carbonServer('37.252.122.50', 'pick-something')
```

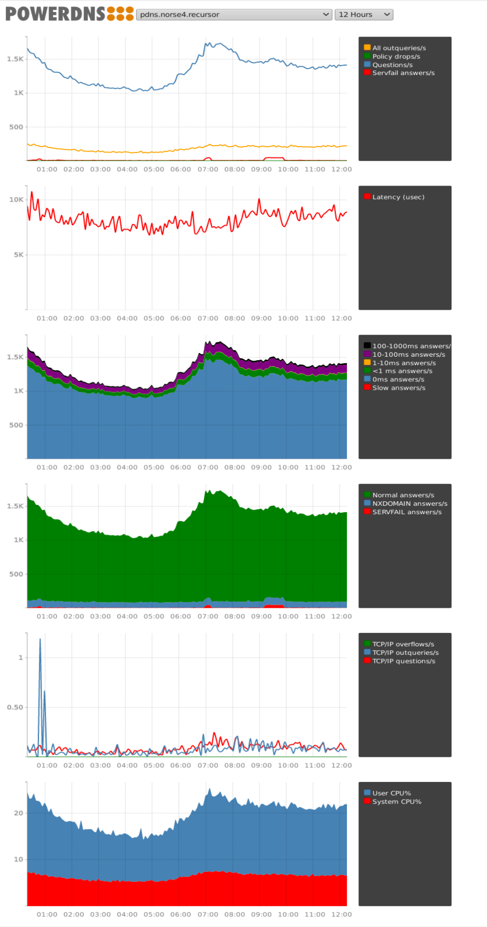

pdns.norse4.recursor

12 Hours

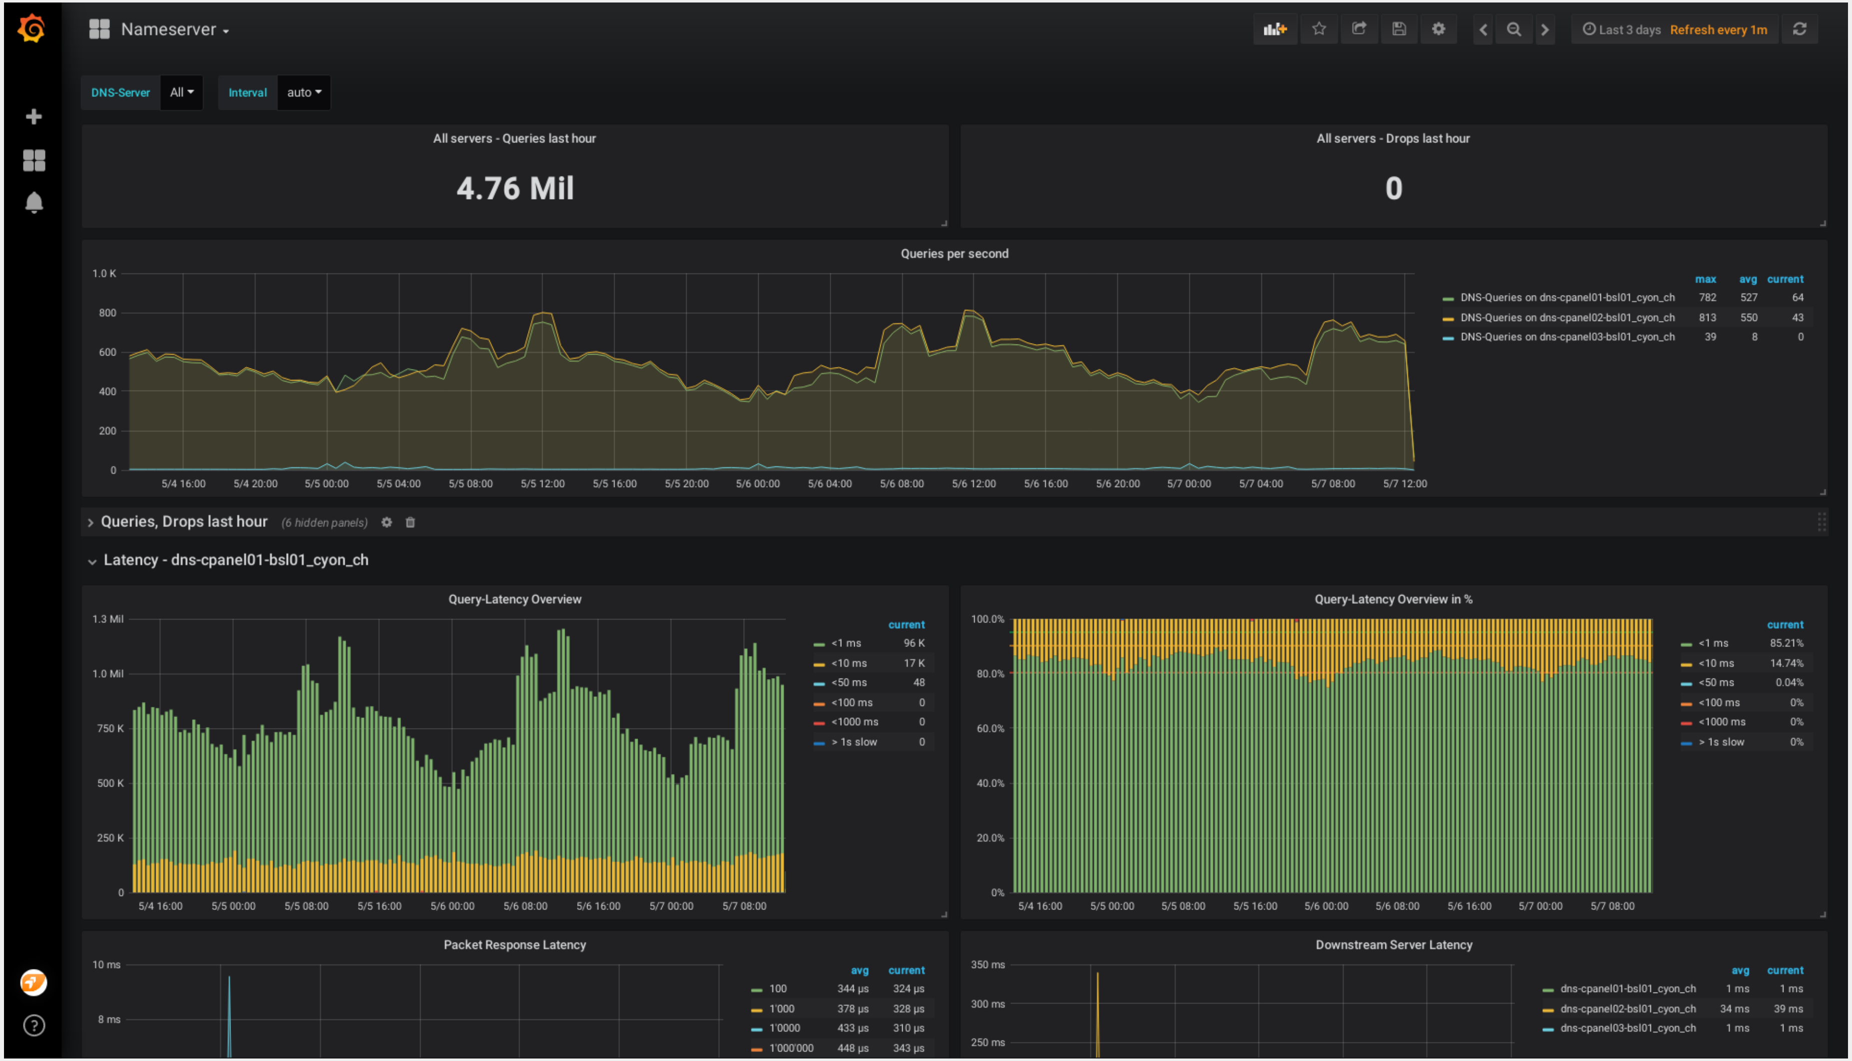

DNS-Server: All
Interval: auto

All servers - Queries last hour

4.76 Mil

All servers - Drops last hour

0

Queries per second

	max	avg	current
🟩 DNS-Queries on dns-cpanel01-bsl01_cyon_ch	782	527	64
🟨 DNS-Queries on dns-cpanel02-bsl01_cyon_ch	813	550	43
🟦 DNS-Queries on dns-cpanel03-bsl01_cyon_ch	39	8	0

› Queries, Drops last hour *(6 hidden panels)*

⌄ Latency - dns-cpanel01-bsl01_cyon_ch

Query-Latency Overview

	current
🟩 <1 ms	96 K
🟨 <10 ms	17 K
🟦 <50 ms	48
🟧 <100 ms	0
🟥 <1000 ms	0
🟦 > 1s slow	0

Query-Latency Overview in %

	current
🟩 <1 ms	85.21%
🟨 <10 ms	14.74%
🟦 <50 ms	0.04%
🟧 <100 ms	0%
🟥 <1000 ms	0%
🟦 > 1s slow	0%

Packet Response Latency

	avg	current
🟩 100	344 μs	324 μs
🟨 1'000	378 μs	328 μs
🟦 1'0000	433 μs	310 μs
🟧 1'000'000	448 μs	343 μs

Downstream Server Latency

	avg	current
🟩 dns-cpanel01-bsl01_cyon_ch	1 ms	1 ms
🟨 dns-cpanel02-bsl01_cyon_ch	34 ms	39 ms
🟦 dns-cpanel03-bsl01_cyon_ch	1 ms	1 ms

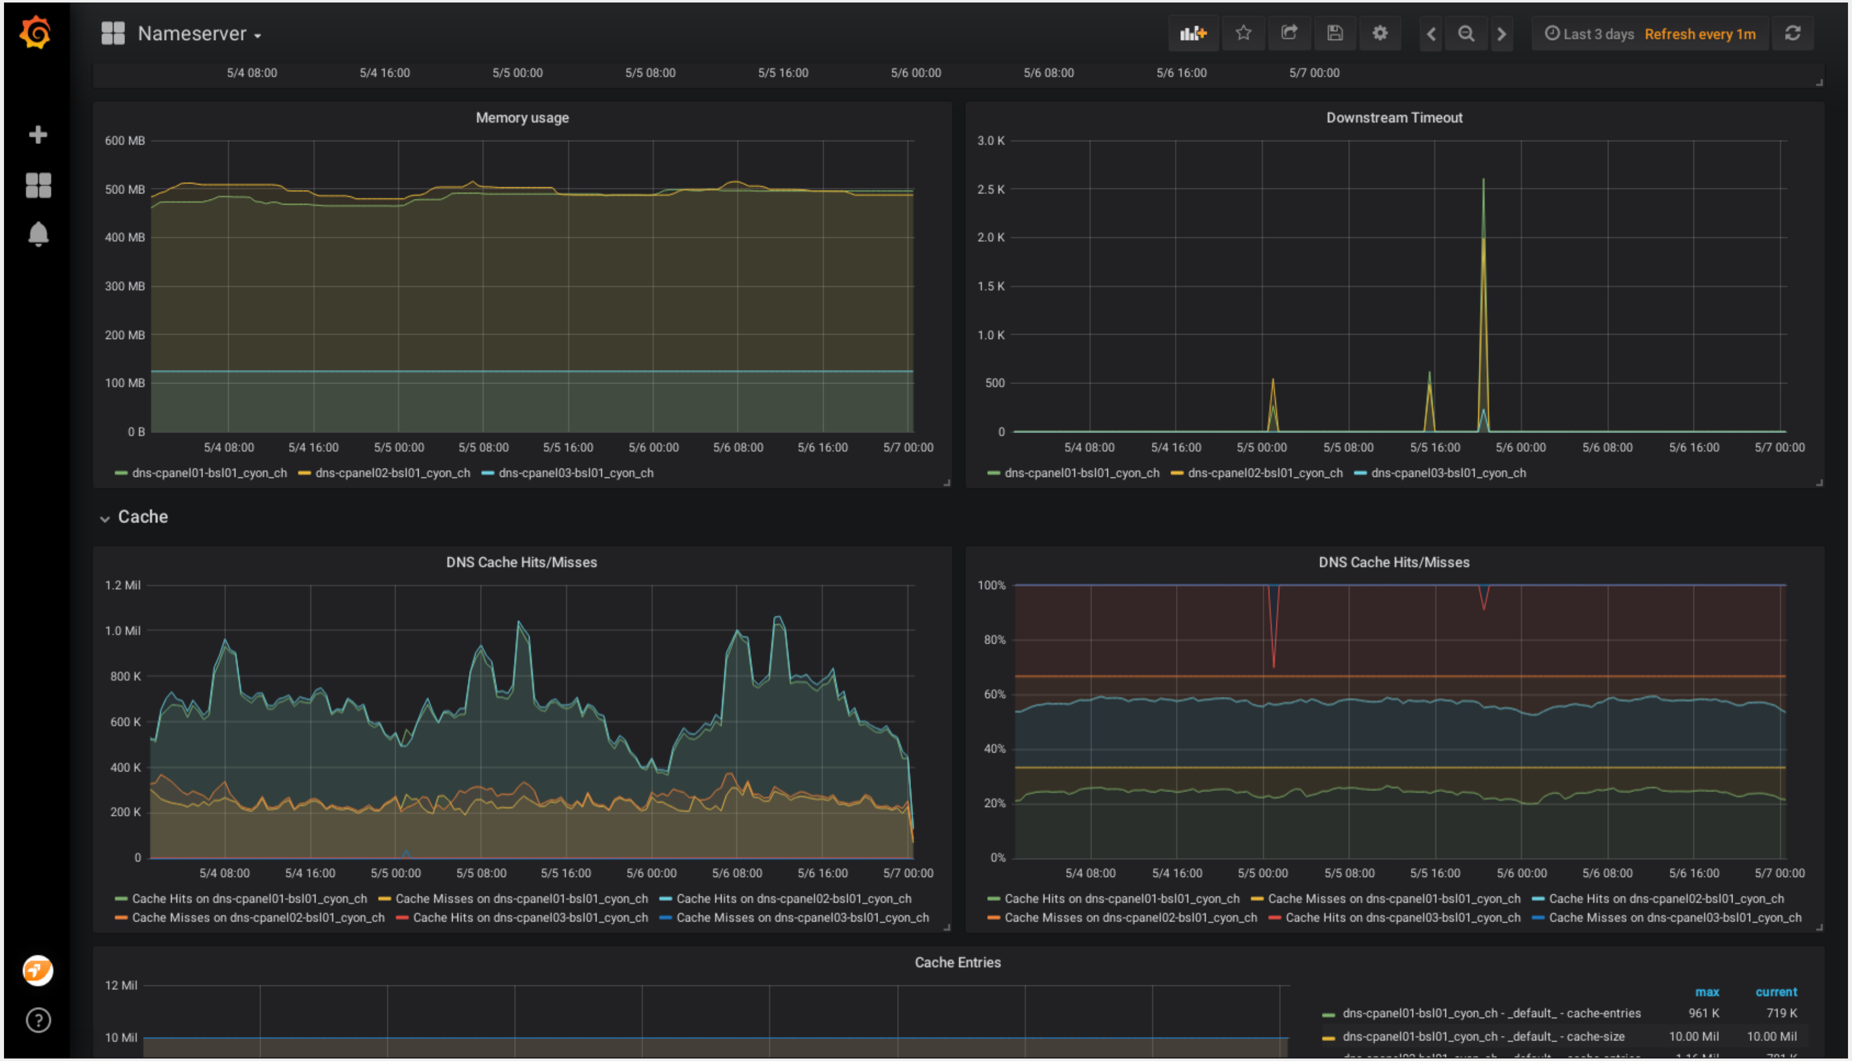

	max	current
dns-cpanel01-bsl01_cyon_ch - _default_ - cache-entries	961 K	719 K
dns-cpanel01-bsl01_cyon_ch - _default_ - cache-size	10.00 Mil	10.00 Mil

Cache

```
getPool(""):setCache(newPacketCache(10000000))
```

Cache

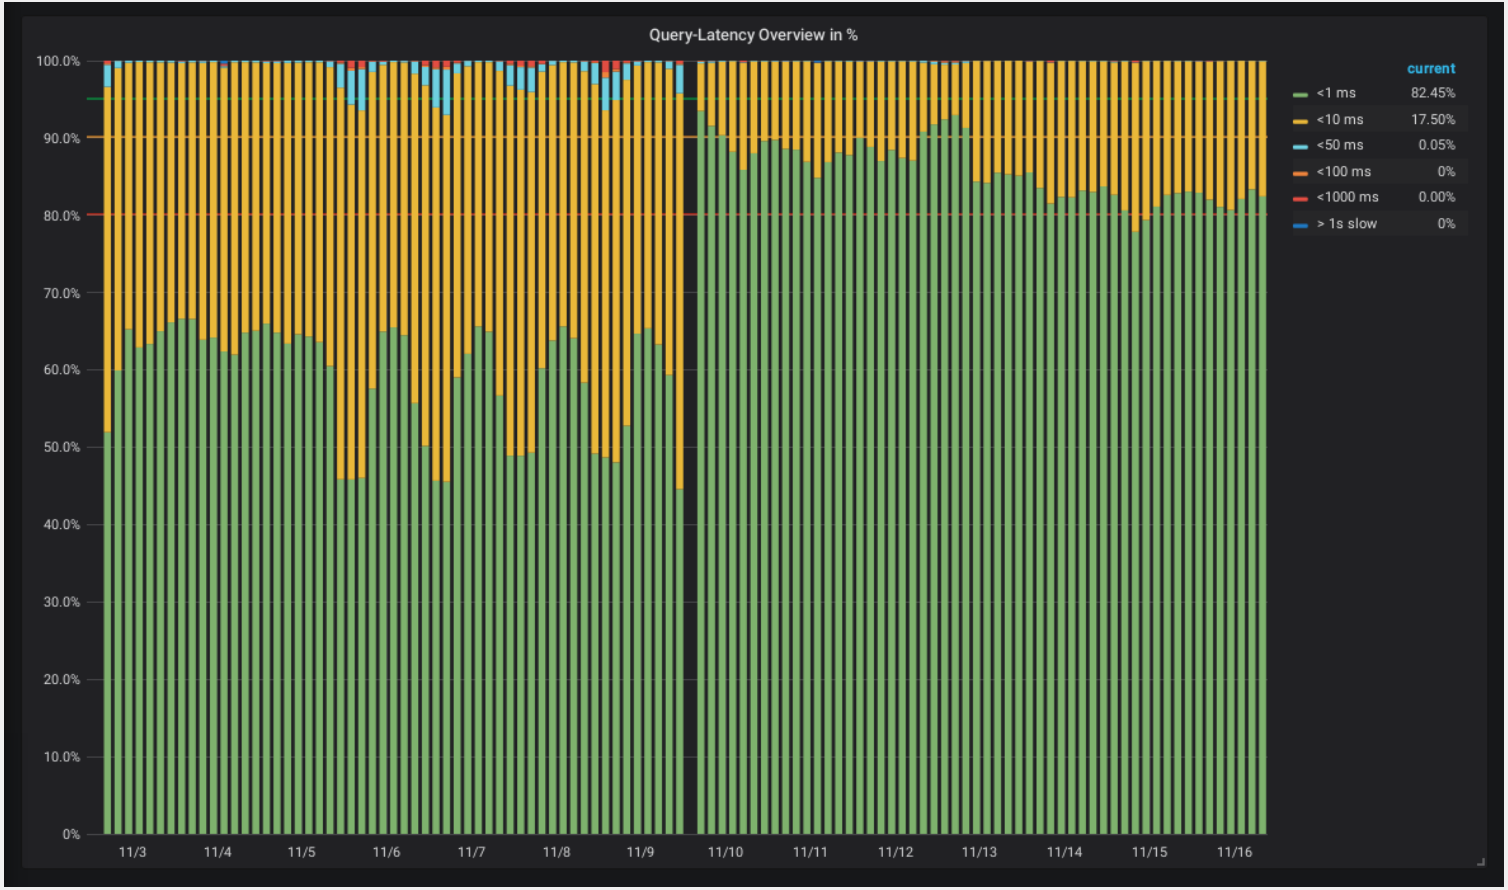
Query-Latency Overview in %
100.0%
90.0%
80.0%
70.0%
60.0%
50.0%
40.0%
30.0%
20.0%
10.0%
0%
11/3
11/4
11/5
11/6
11/7
11/8
11/9
11/10
11/11
11/12
11/13
11/14
11/15
11/16
current
<1 ms 82.45%
<10 ms 17.50%
<50 ms 0.05%
<100 ms 0%
<1000 ms 0.00%
> 1s slow 0%

Advanced features

- DOT and DNSCrypt support
- eBPF filtering
- Protobuf and dnstap logging
- …
- Go read the docs! https://dnsdist.org

~1500 q/s

Used to be more*

84% <1ms | 99% < 10ms

It's fast 👍

~70% cache hits

Absorbs traffic

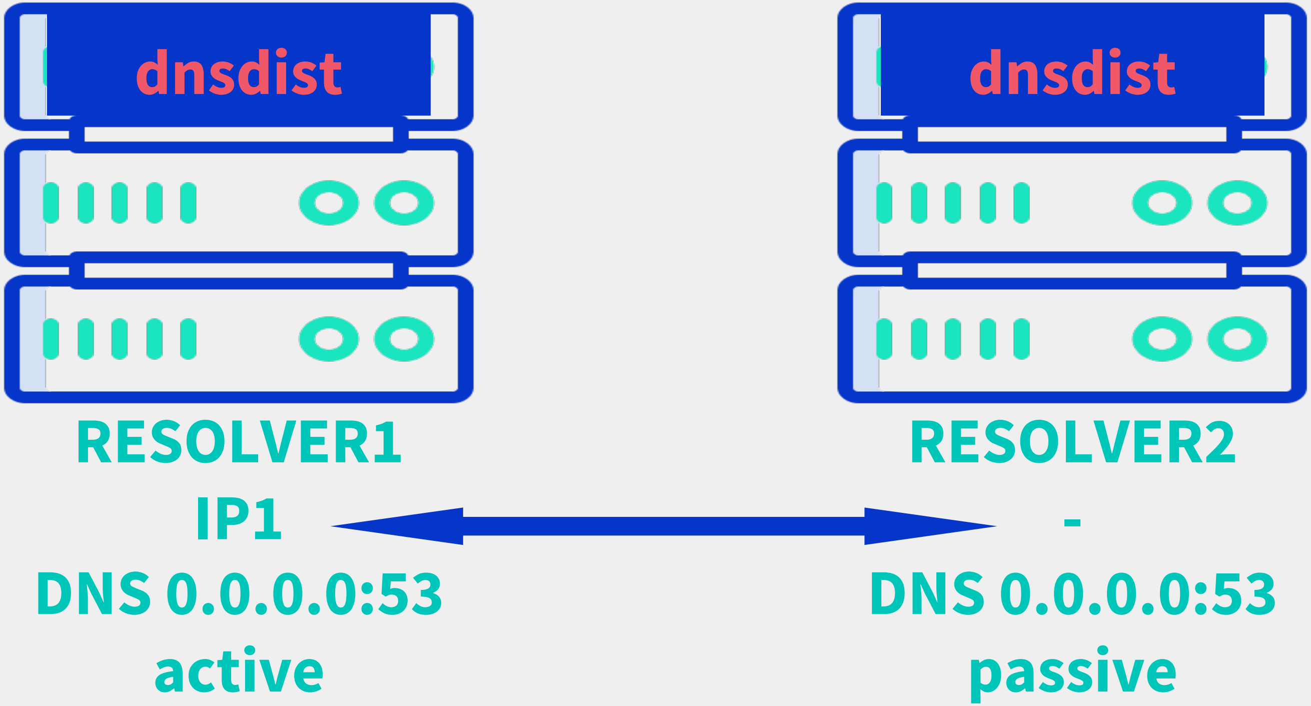

RESOLVER1
IP1
DNS 0.0.0.0:53
active

RESOLVER2
-
DNS 0.0.0.0:53
passive

Resolver

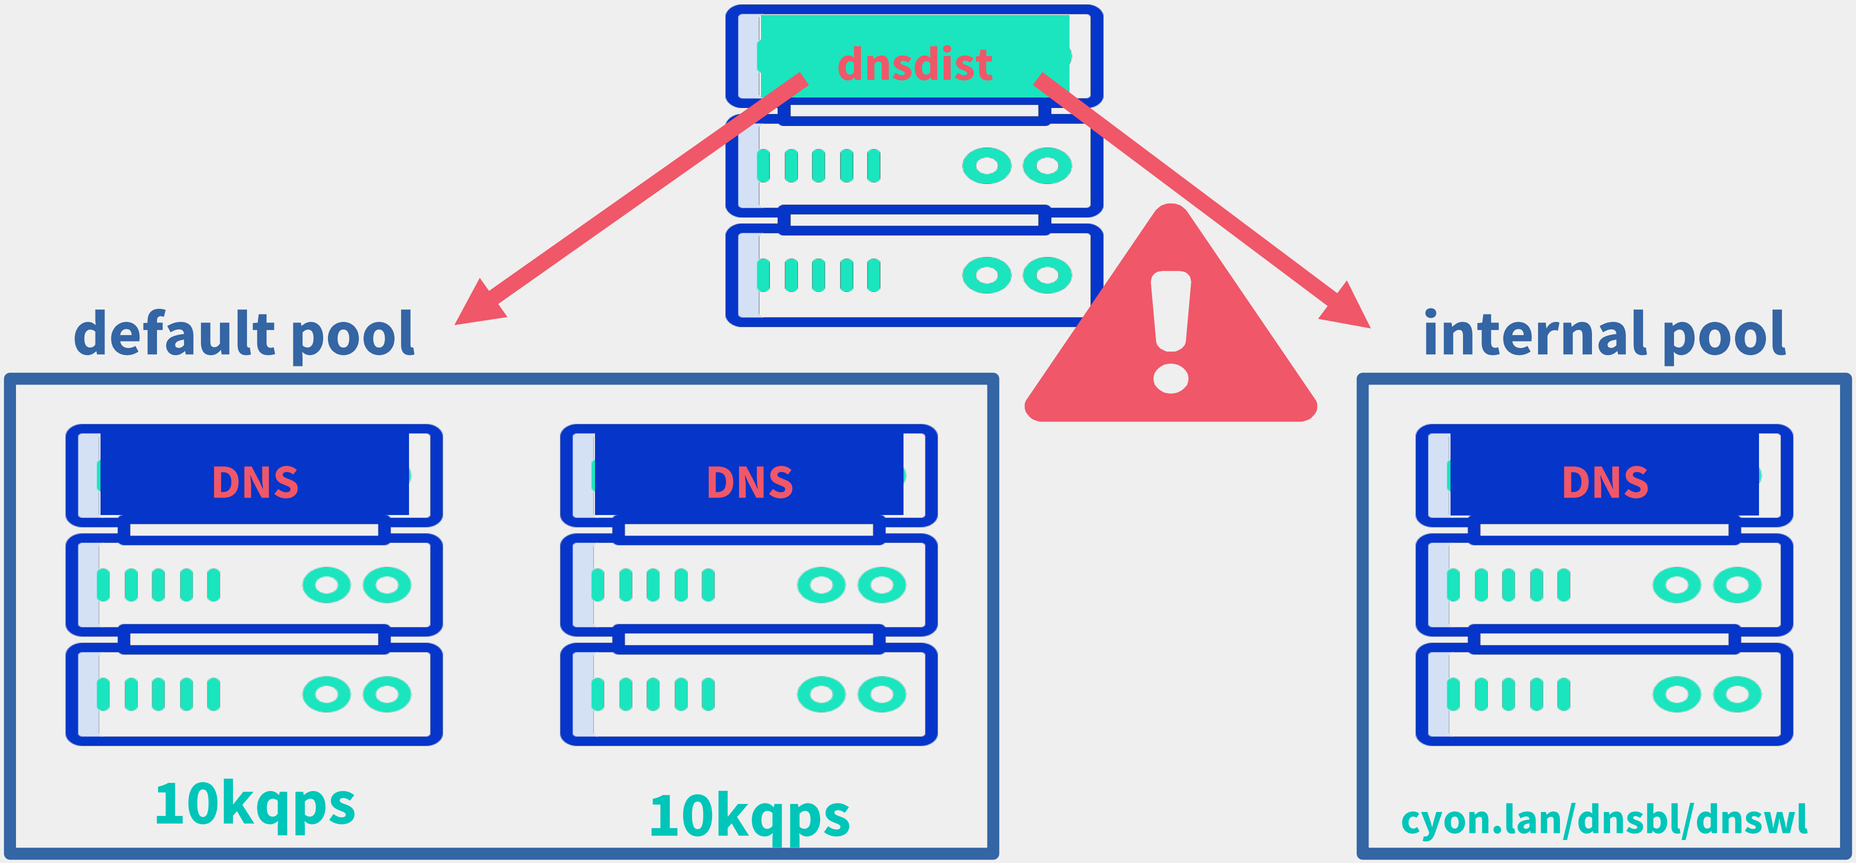
dnsdist
default pool
internal pool
DNS
DNS
DNS
10kqps
10kqps
cyon.lan/dnsbl/dnswl

~1000 q/s

82% <1ms | 90% < 10ms

It's fast 👍

~70% cache hits

Absorbs traffic

dnsdist 1.4

New TCP stack

Event-based TCP handling

Privacy

Pseudonymize IP addresses

DOH

DNS-over-HTTPS

Security

Systemd sandboxing, drop capabilities

Stale data

Keep stale data on cache purge

More stats

New "frontend" caching stats

Current state

- Dependent on vendor
- Old soft~~ware~~
- ~~Software has bugs~~
  - ~~NS delegation~~
  - ~~Max size of TXT record~~
- ~~No caching~~
- ~~Minimal fault tolerance~~
- ~~No stats or insights~~

LESSIONS LEARNED

- Put dnsdist in front of you DNS server
- Safely iterate
- Gather stats and health information

POWERDNS

Thanks!

Any questions?

You can find me and my coworker at:

dol@cyon.ch / @drdol

aw@cyon.ch / Andreas Willi