nominum™

What to do with SERVFAIL

The day wd2go.com fell off the Internet

What is wd2go.com

Personal cloud storage

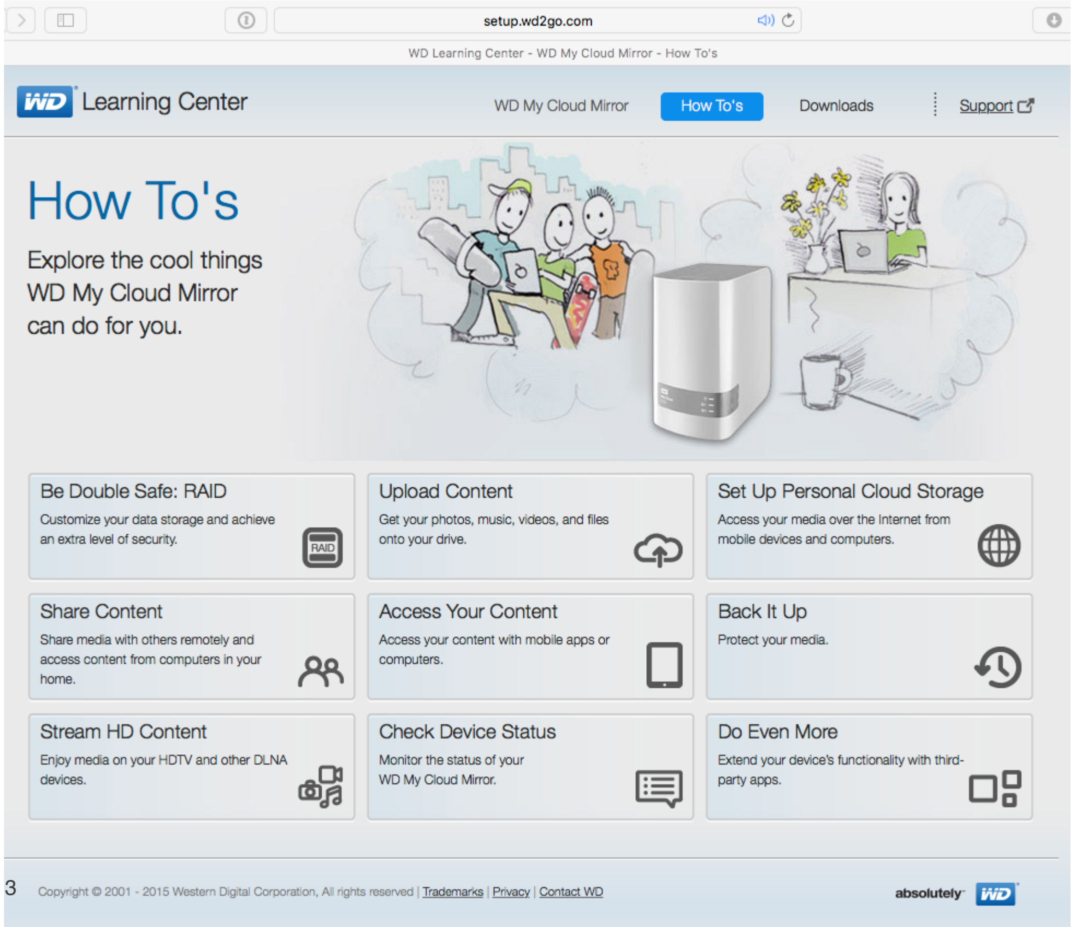

setup.wd2go.com

WD Learning Center - WD My Cloud Mirror - How To's

How To's

Explore the cool things WD My Cloud Mirror can do for you.

Be Double Safe: RAID
Customize your data storage and achieve an extra level of security.

Upload Content
Get your photos, music, videos, and files onto your drive.

Set Up Personal Cloud Storage
Access your media over the Internet from mobile devices and computers.

Share Content
Share media with others remotely and access content from computers in your home.

Access Your Content
Access your content with mobile apps or computers.

Back It Up
Protect your media.

Stream HD Content
Enjoy media on your HDTV and other DLNA devices.

Check Device Status
Monitor the status of your WD My Cloud Mirror.

Do Even More
Extend your device's functionality with third-party apps.

Copyright © 2001 - 2015 Western Digital Corporation, All rights reserved | Trademarks | Privacy | Contact WD

- Local Disk Storage for
  - Backups
  - Photos
  - Etc
- All content also accessible via Internet
- Acts as a cloud to your devices

DNS Parameters (on a normal day)

Names	
Name	%
www.wd2go.com.	59.47
web.wd2go.com.	11.77
api.wd2go.com.	1.44
discovery.wd2go.com.	0.94
relay-prod-apsoutheast2-3.wd2go.com.	0.29
relay-prod-apsoutheast2-4.wd2go.com.	0.28
device2998911-158c8433-local.wd2go.com.	0.13
device2998911-158c8433.wd2go.com.	0.13

Query types		
Query Type	Number	%
A	1	61.985895
AAAA	28	38.010420
DS	43	0.002190
DLV	32769	0.000940
NS	2	0.000523
CNAME	5	0.000019
TXT	16	0.000009

DNS Setup (before the event)

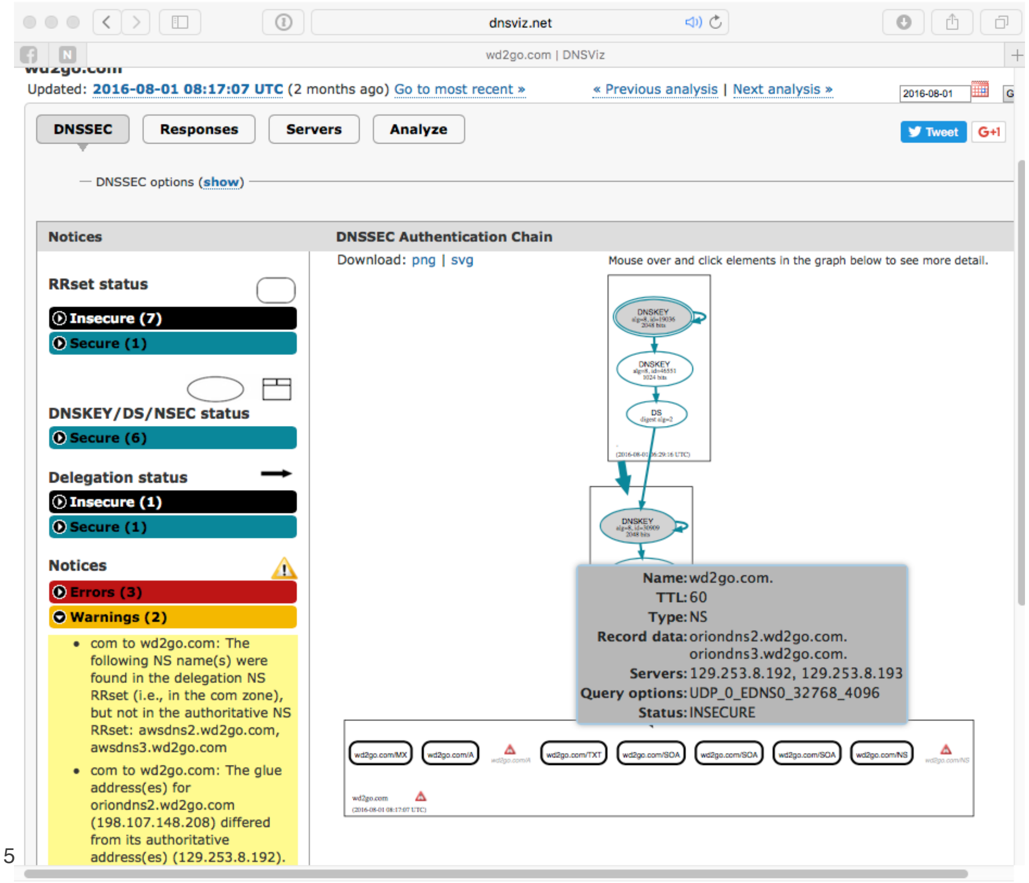

wd2go.com

Updated: 2016-08-01 08:17:07 UTC (2 months ago) Go to most recent » « Previous analysis | Next analysis » 2016-08-01

DNSSEC Responses Servers Analyze Tweet G+1

— DNSSEC options (show)

Notices

RRset status

- Insecure (7)
- Secure (1)

DNSKEY/DS/NSEC status

- Secure (6)

Delegation status

- Insecure (1)
- Secure (1)

Notices

- Errors (3)
- Warnings (2)
  - com to wd2go.com: The following NS name(s) were found in the delegation NS RRset (i.e., in the com zone), but not in the authoritative NS RRset: awsdns2.wd2go.com, awsdns3.wd2go.com
  - com to wd2go.com: The glue address(es) for oriondns2.wd2go.com (198.107.148.208) differed from its authoritative address(es) (129.253.8.192).

DNSSEC Authentication Chain

Download: png | svg

Mouse over and click elements in the graph below to see more detail.

Name: wd2go.com.
TTL: 60
Type: NS
Record data: oriondns2.wd2go.com.
oriondns3.wd2go.com.
Servers: 129.253.8.192, 129.253.8.193
Query options: UDP_0_EDNS0_32768_4096
Status: INSECURE

DNS Traffic (on a normal day)

Wd2go.com characterisitic	Result
Type of traffic	flat
Relative QPS	0.002% - 0.06%
QPS per client (avg)	0.003 – 0.01
Clients sending traffic	0.03% - 0.5%

nominum™

The August 3 Oooops

DNS Traffic on August 3

Type of Traffic	Spikey
Relative QPS	4% - 43%
QPS per client (avg)	3 – 20 qps
[illegible]	

Is this a problem

- Normal traffic 0.2 qps per sub
  - 200kqps per 1 million subs
- Now a few subs have more
  - 0.5% = 5000 clients
  - 5000 clients * 20 qps = 100kqps

Impact of problem

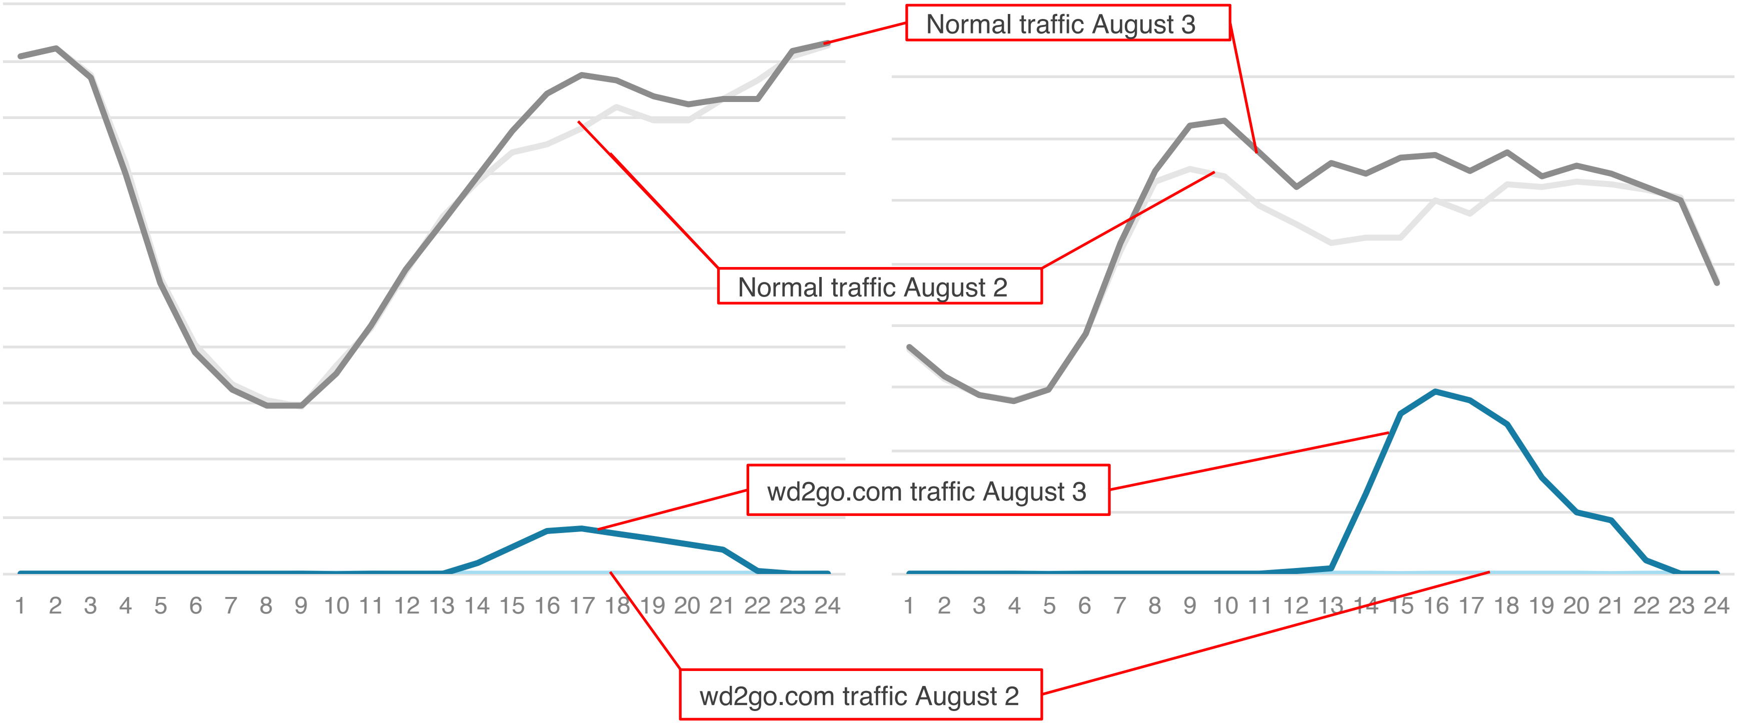

Normal traffic August 3

Normal traffic August 2

wd2go.com traffic August 3

wd2go.com traffic August 2

DNS Parameters (during incident)

Names	
Name	**%**
www.wd2go.com.	3.037
relay-prod-eucentral1-20.wd2go.com.	1.167
relay-prod-euwest1-12.wd2go.com..	1.162
web.wd2go.com.	0.665
relay-prod-useast1-13.wd2go.com.	0.265
relay-prod-apsoutheast2-4.wd2go.com.	0.283
device3263568-9cfc0eee.wd2go.com.	0.001
device1944395-96d9ffc1.wd2go.com.	0.001

Query types		
Query Type	**Number**	**%**
A	1	98.4618319
AAAA	28	1.5381624
A6	38	0.0000037
NS	2	0.0000007
TXT	16	0.0000005
DS	43	0.0000004
SOA	6	0.0000001

Summary of findings

- A medium size domain going off can cause a lot of strain on the infrastructure
  - Over provisioning still is good
  - No IPv6 if things go south
- Would be good if clients didn't retry as hard and back off
- How are servers responding to SERVFAIL storms caused by outages ?

How does different software handle this

How to replicate the event in a lab

- Create a domain where all the name servers are not reachable
- Create a dnsperf file with a name that of that domain
- Start a caching server
- Fire up dnsperf
  - Timout 5 seconds
  - -q large enough to still sending when server doesn't answer
  - -Q 10000 qps

```
dnsperf -Q 10000 -c 100 -T 10 -q 250000 -t 5 -l 60 -S 1 -d oneservfail.q –s I.P.I.P
```

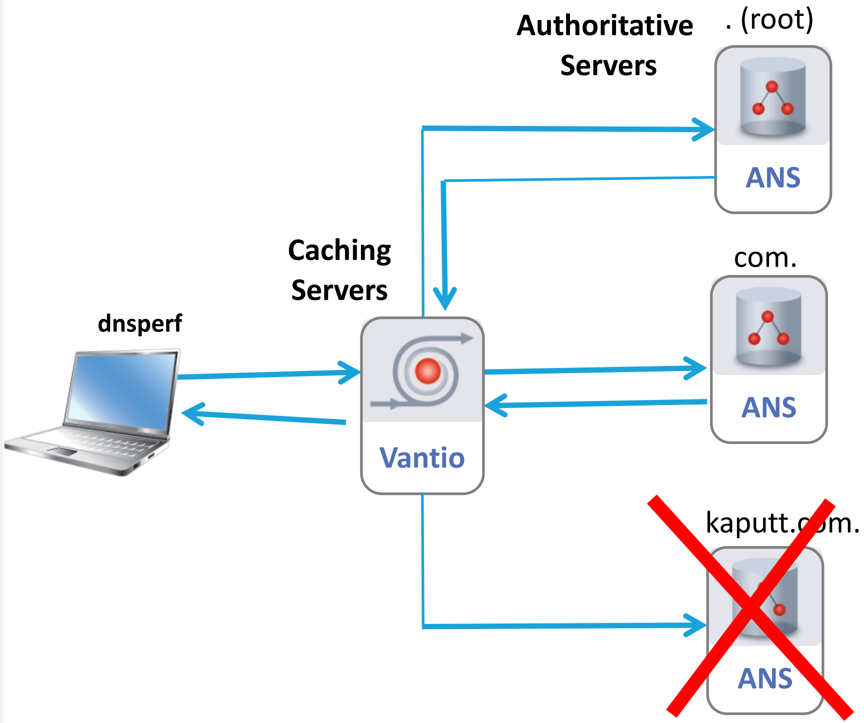

First results in

Software	Packets answered	Packets not answered
Bind-9.9	0 (0.00%)	600000 (100.00%)
Bind-9.10	0 (0.00%)	600000 (100.00%)
Bind-9.11	50000 (8.33%)	550000 (91.67%)
Cacheserve 7	598552 (99.76%)	1448 (0.24%)
Powerdns-4.0	599476 (99.91%)	524 (0.09%)
Unbund 1.5.10	66734 (11.12%)	533266 (88.88%)

- That looks weird…..
  - Am I doing something wrong?
  - Lets look at the details

Bind 9.11 details

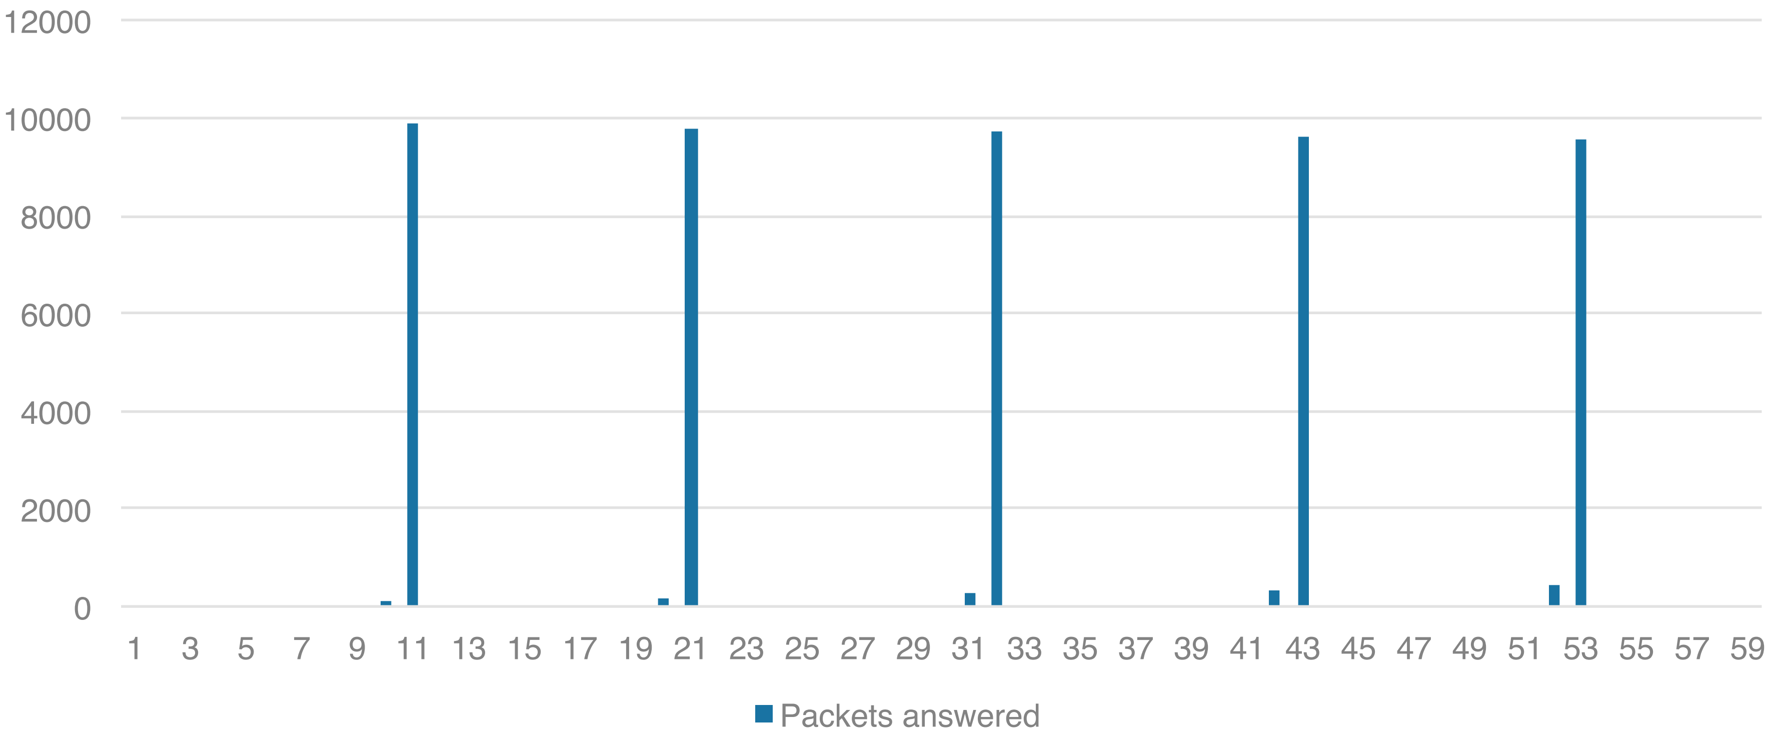

Cacheserve 7 details

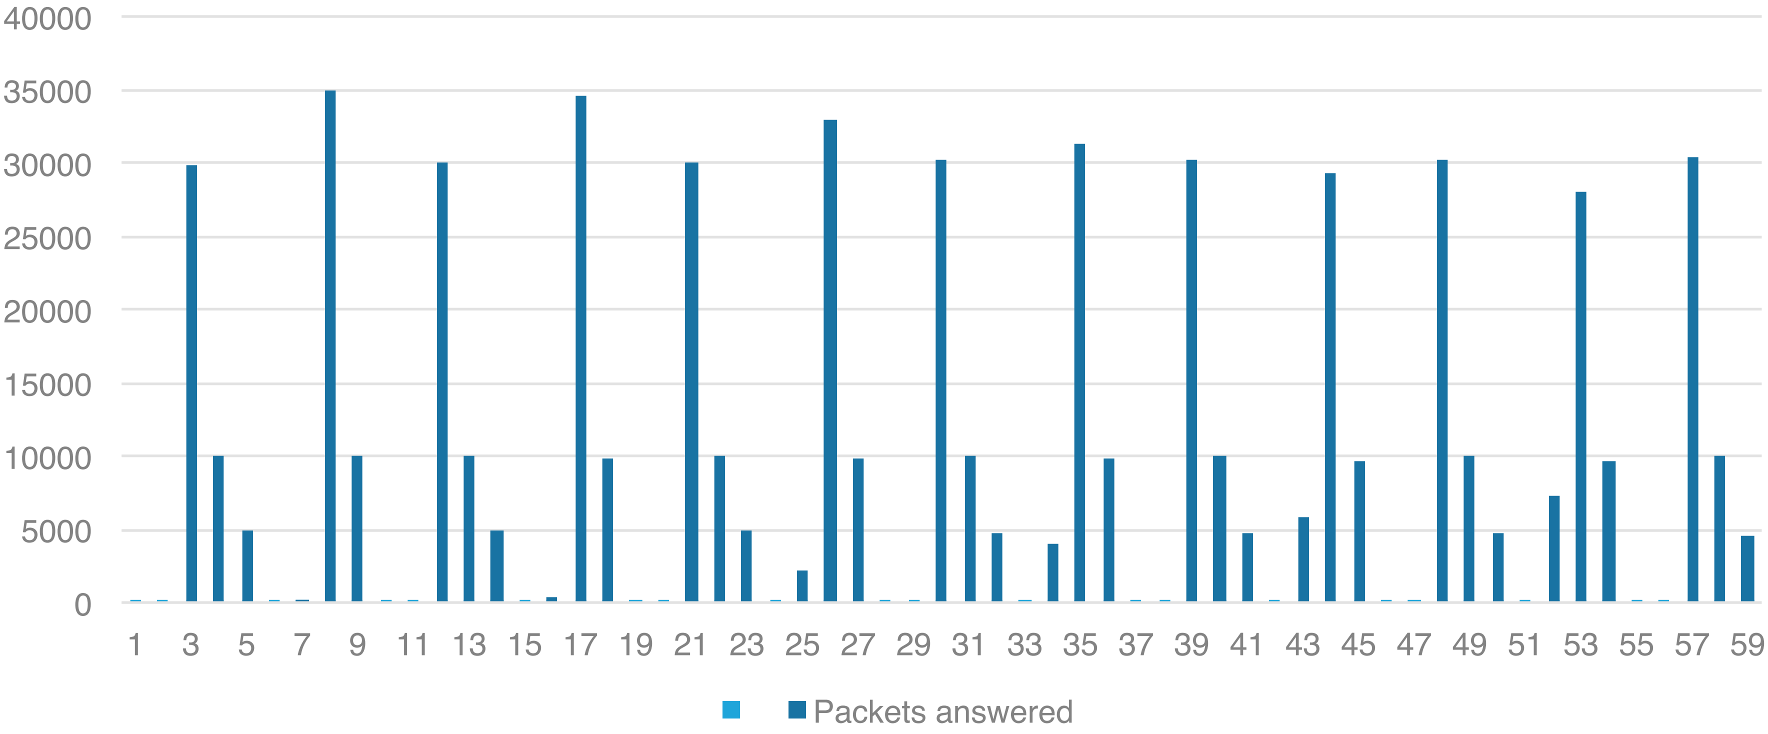

Powerdns 4.0 details

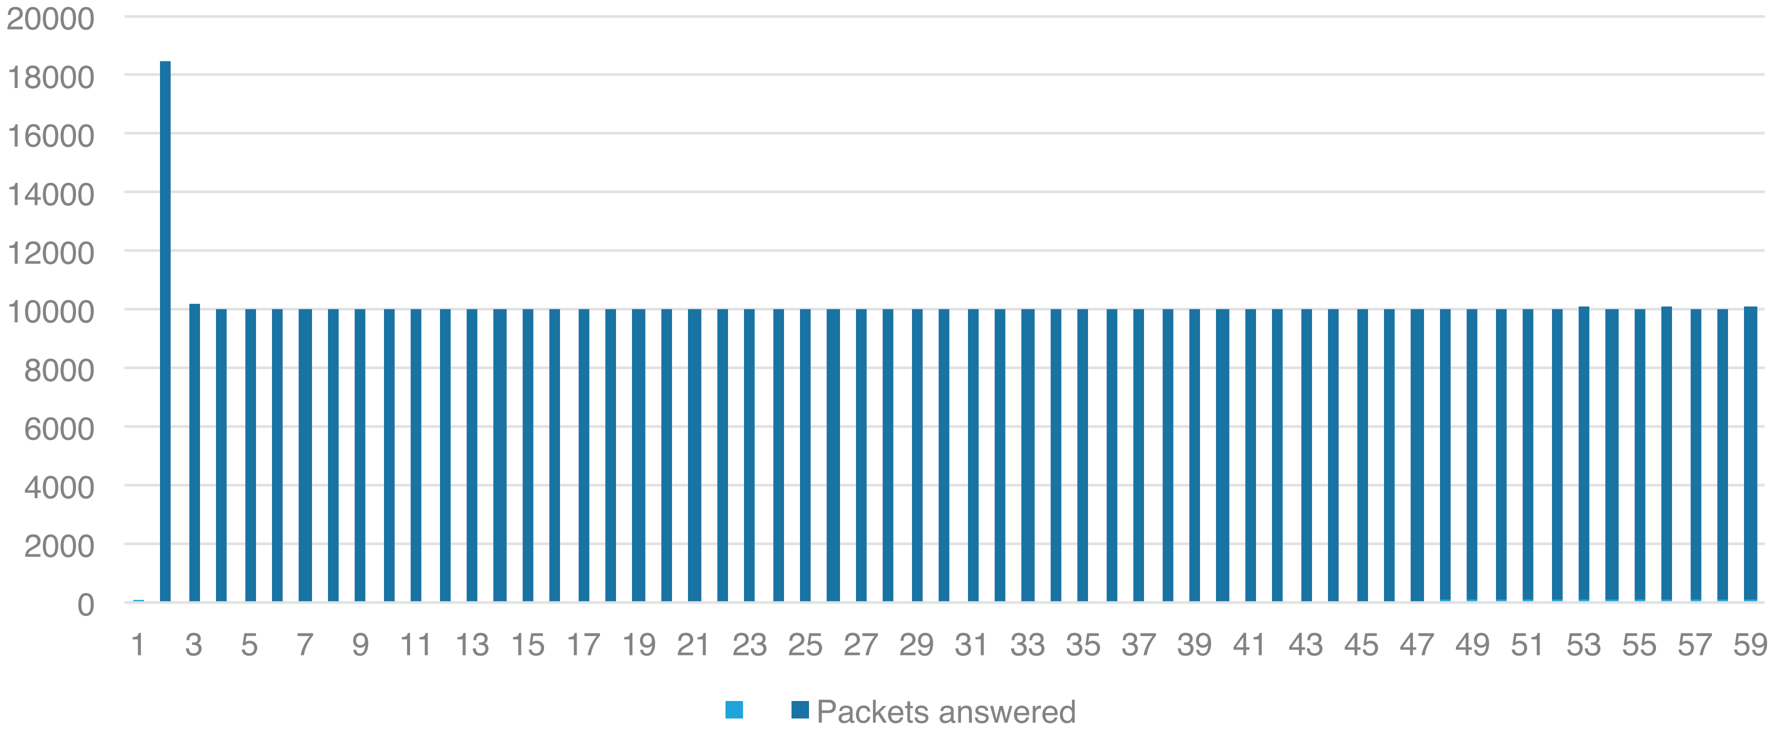

Unbound 1.5.10 details

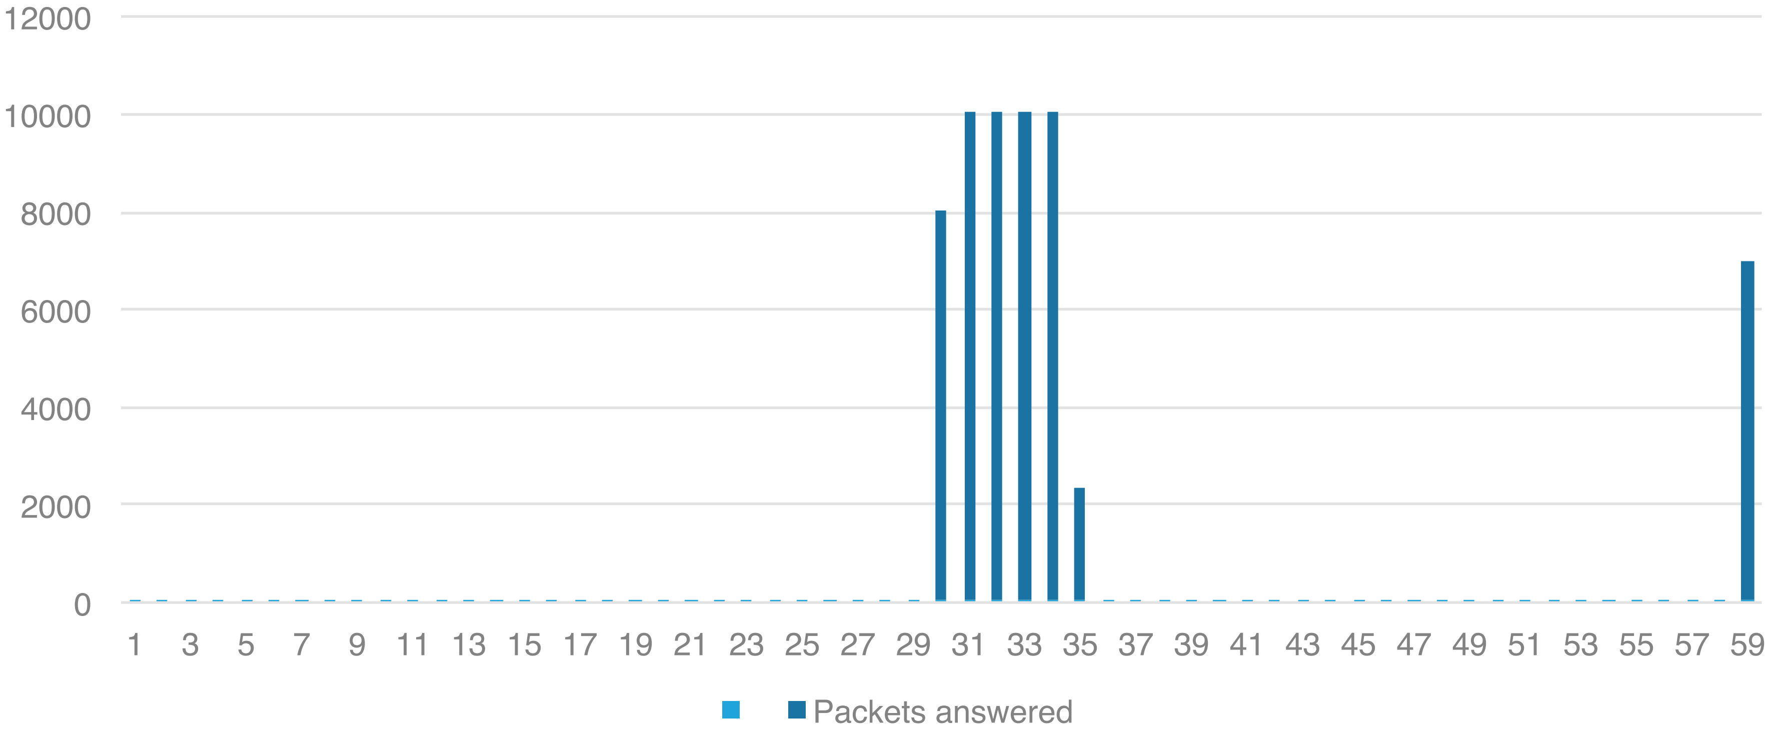

Thinking and more tests

- Looks like different server software has different personality
  - Different cache servfail timers
    - Longer timers means less work, but slower recovery
  - Different levels of persistence in trying to get an answer
- Test really didn't reflect the outage
  - More subdomains were asked
  - Refine test to ask 5000 different subdomains
- Fire

More results in (5000 different subdomains)

Software	Packets answered	Packets not answered
Bind-9.9	260597 (43.43%)	339403 (56.57%)
Bind-9.10	288966 (48.16%)	311034 (51.84%)
Bind-9.11	599100 (99.85%)	900 (0.15%)
Cacheserve 7	591961 (98.66%)	8039 (1.34%)
Powerdns-4.0	587090 (97.85%)	12910 (2.15%)
Unbund 1.5.10	348456 (58.08%)	251544 (41.92%)

- That looks more consistent…..
  - With the exception of older bind and unbound
  - Lets look at the details

Bind 9.9 details

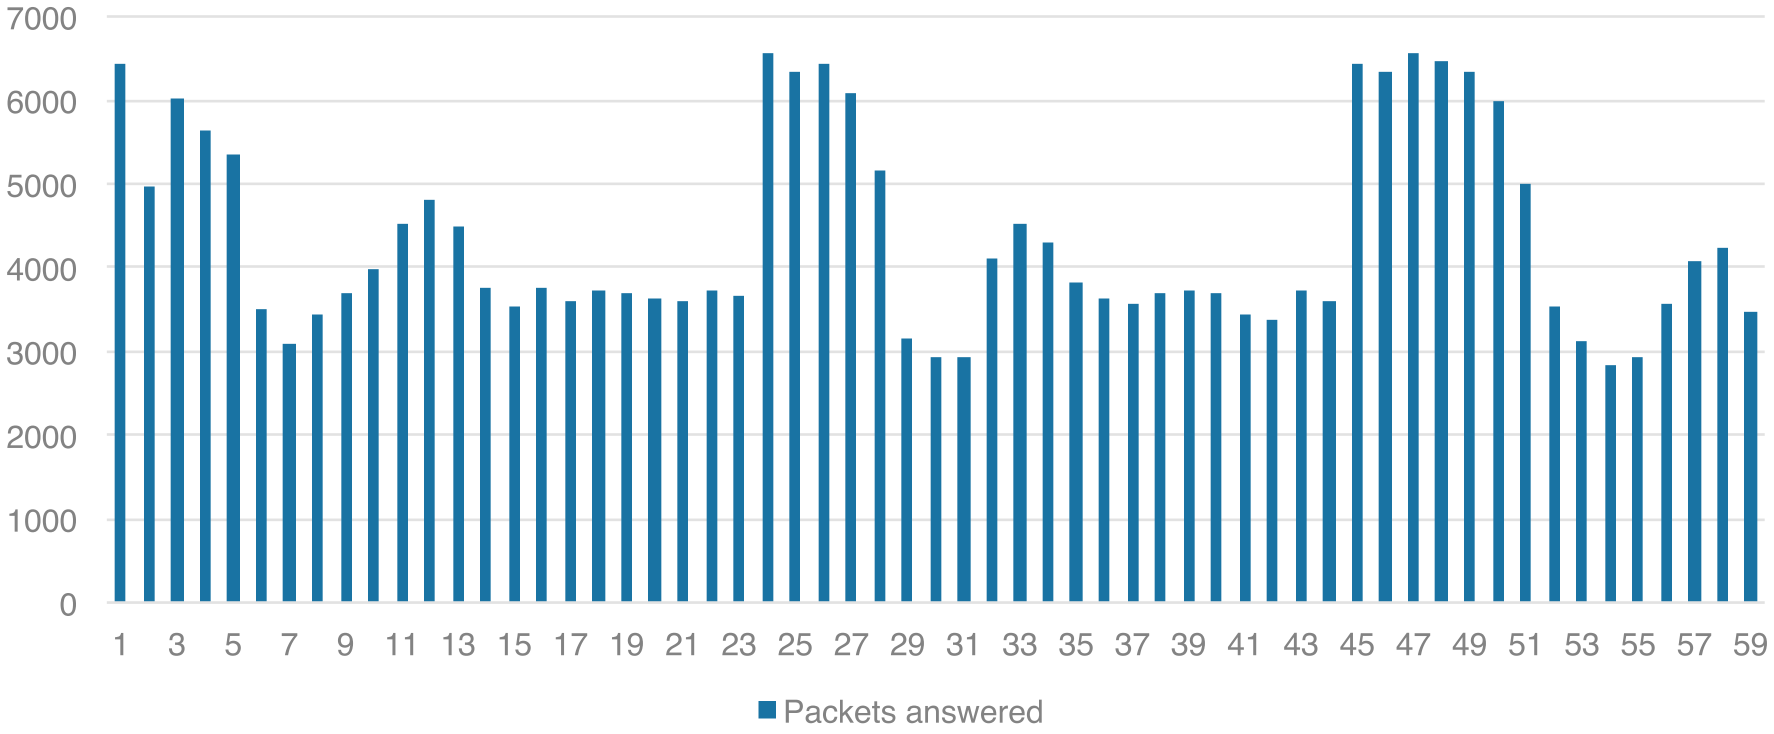

Bind 9.10 details

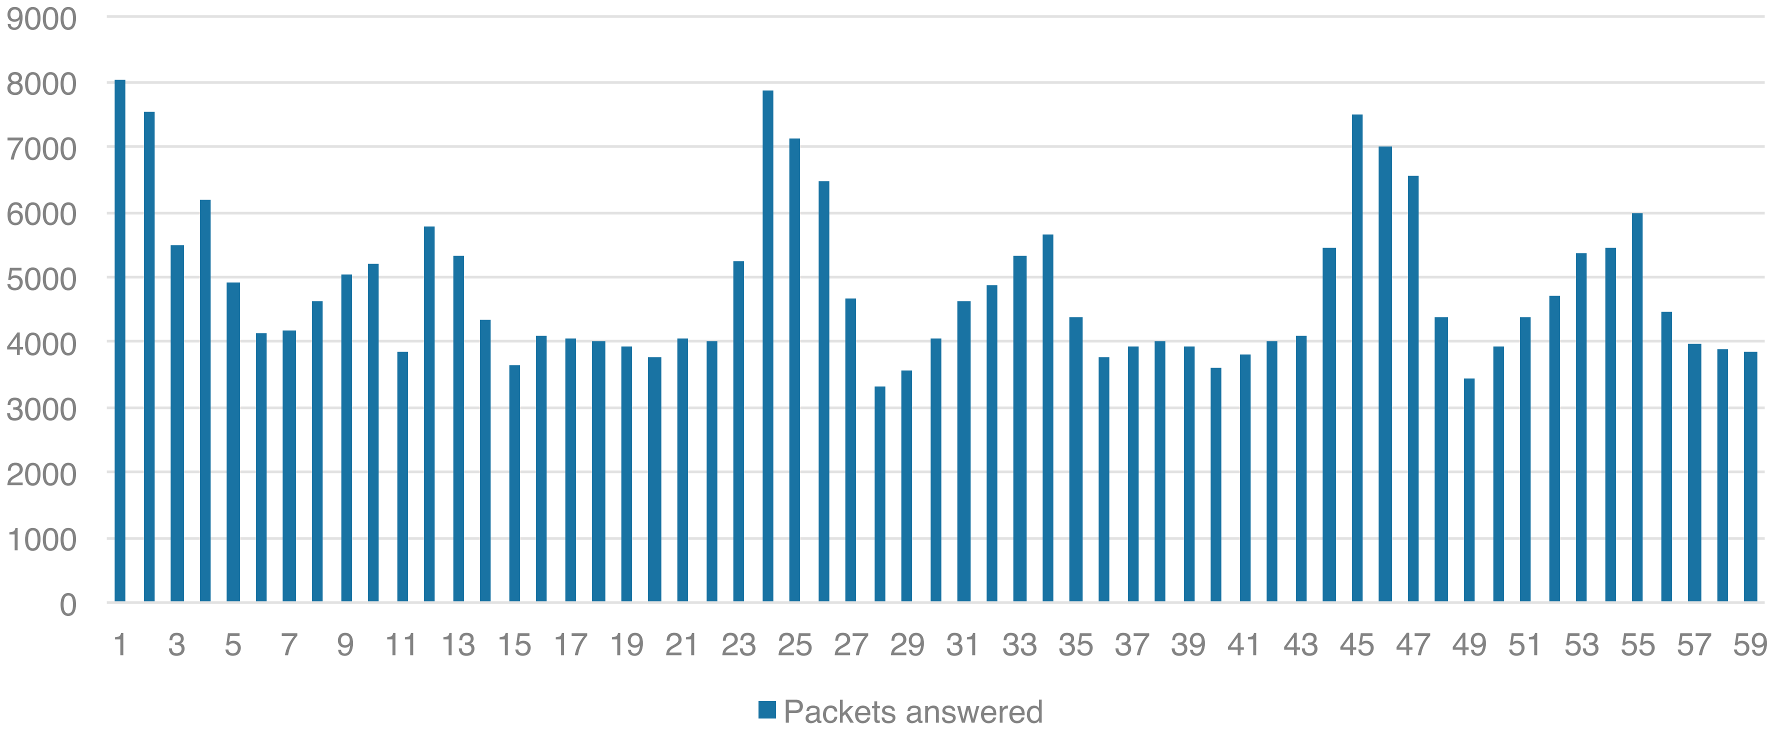

Bind 9.11 details

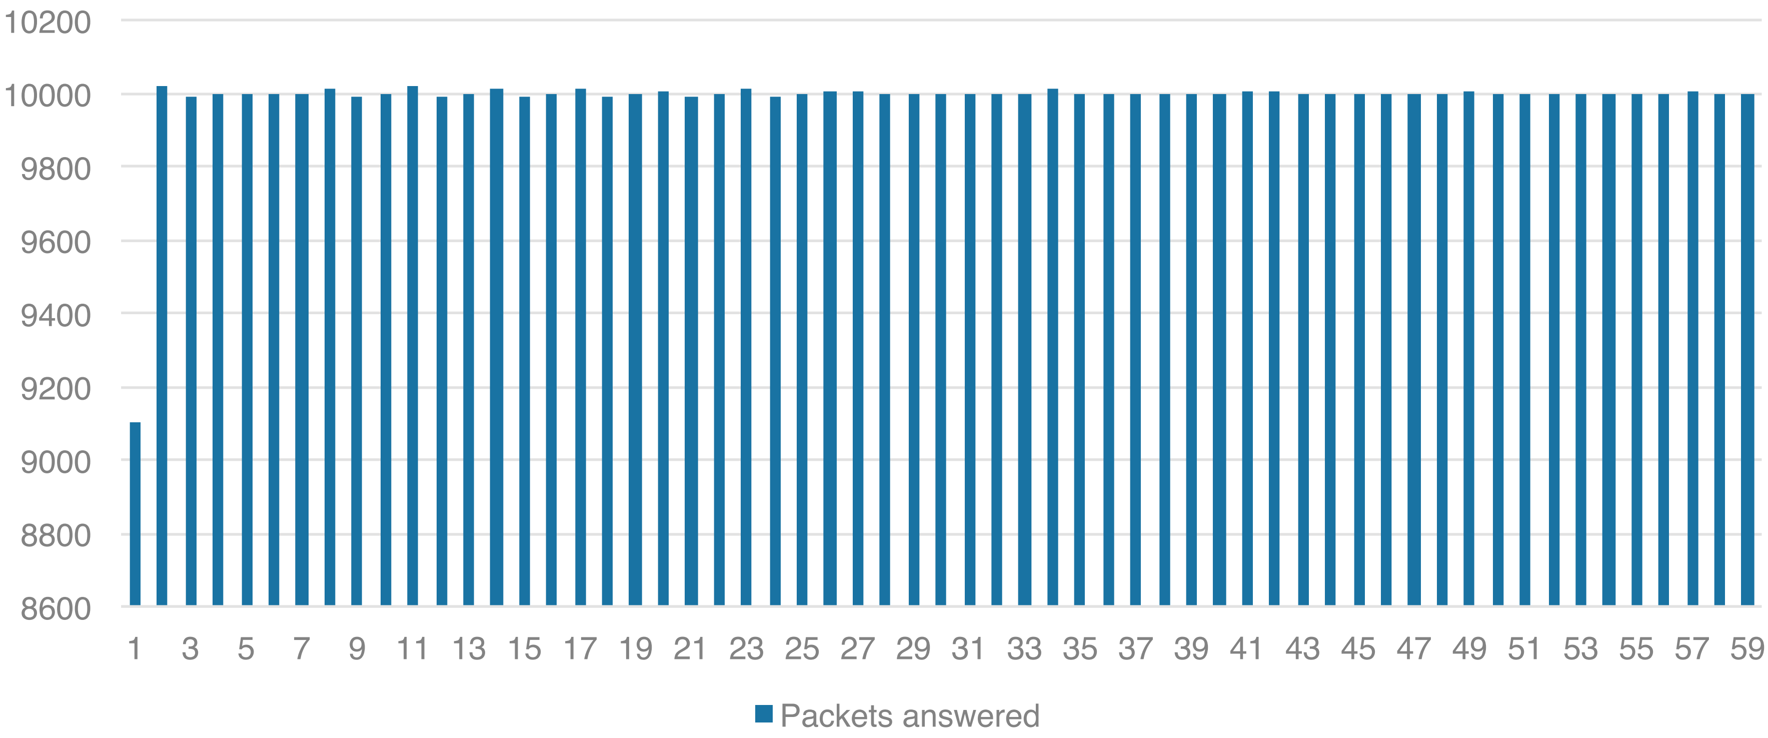

Cacheserve 7 details

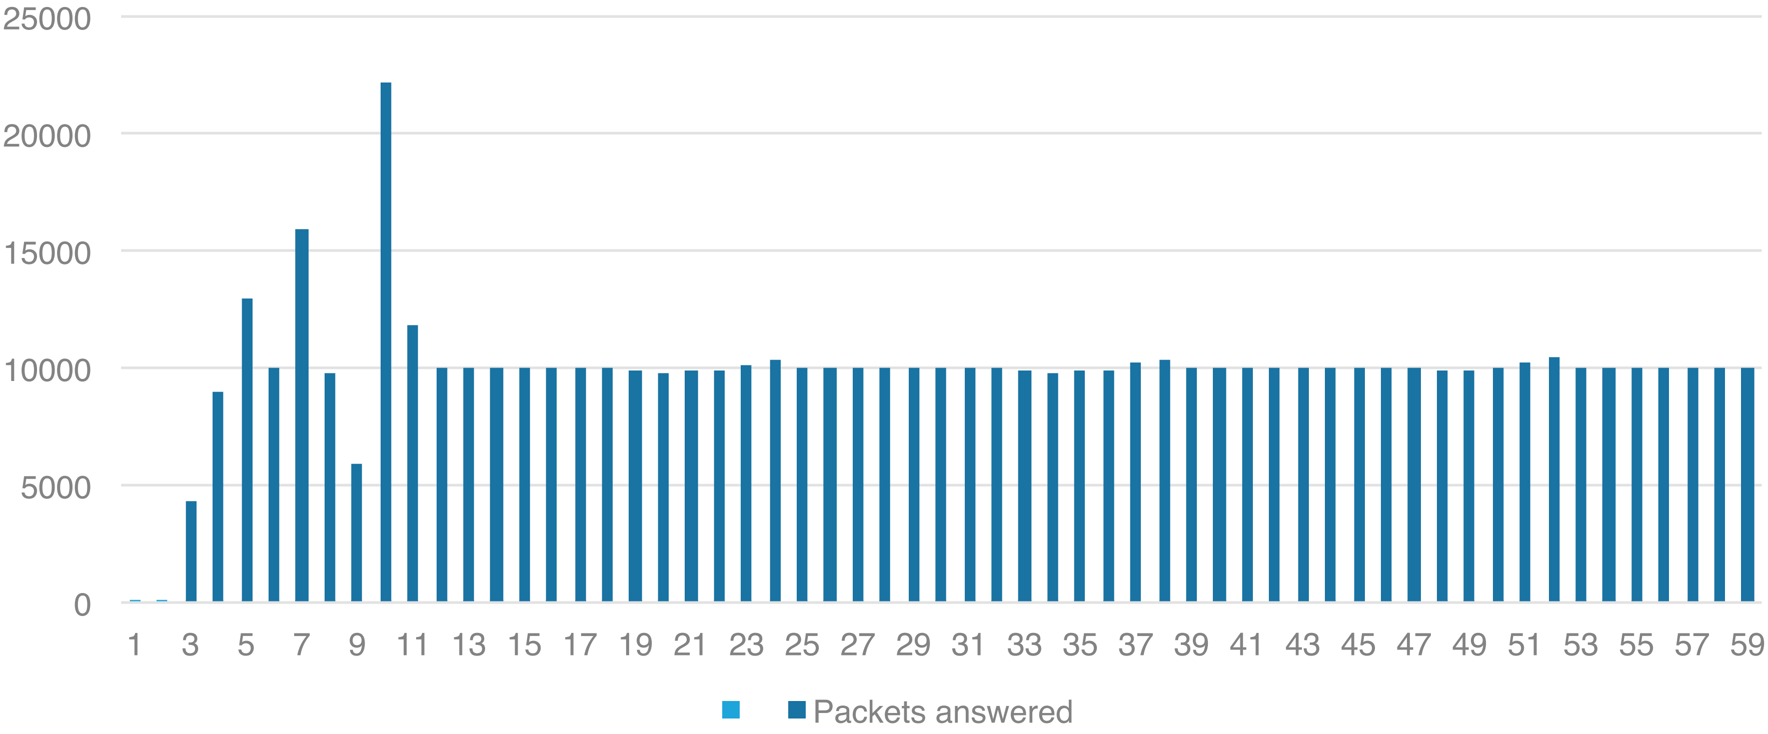

Powerdns 4.0 details

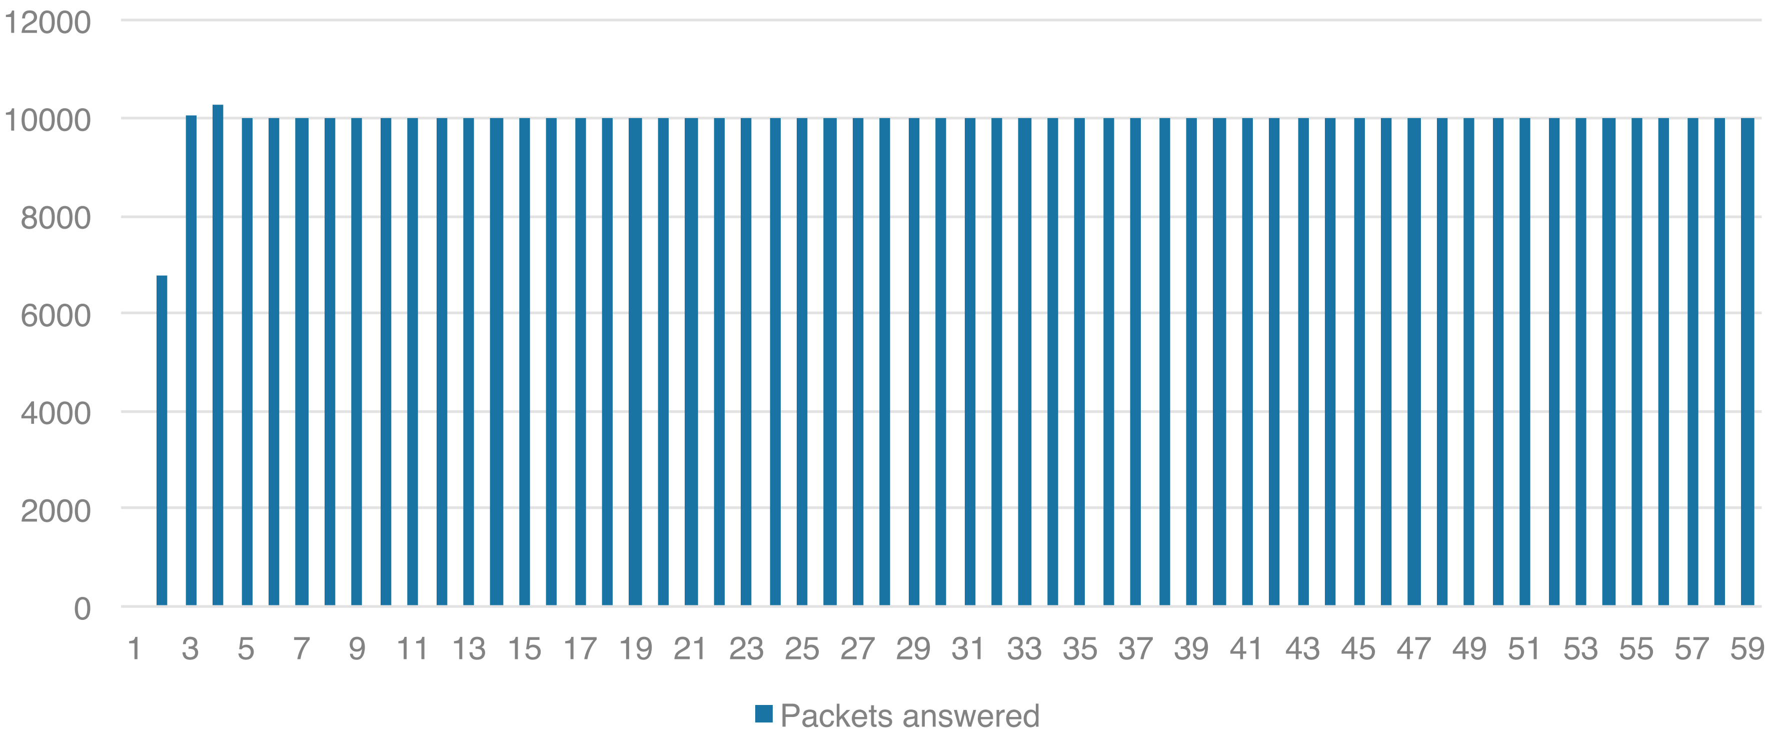

Unbound 1.5.10 details

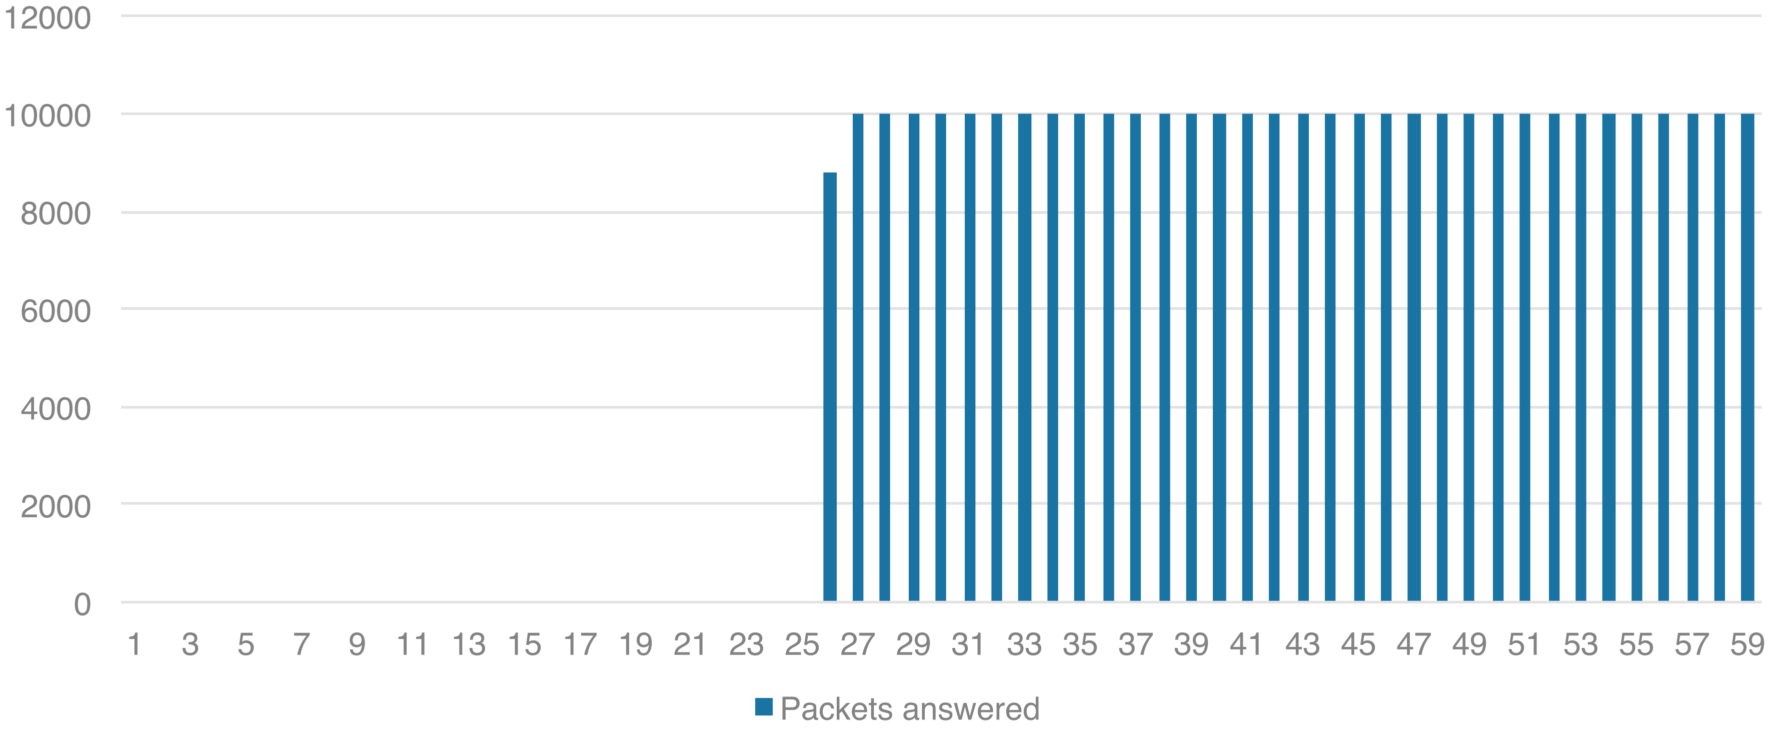

Summary

- Given enough traffic to bad domains most software will answer a lot
  - Not answering is bad for most use cases
- Machines will keep on retrying
- Traffic increases
  - Still good to over provision
- Client coders really should do backoff strategies for SERVFAIL

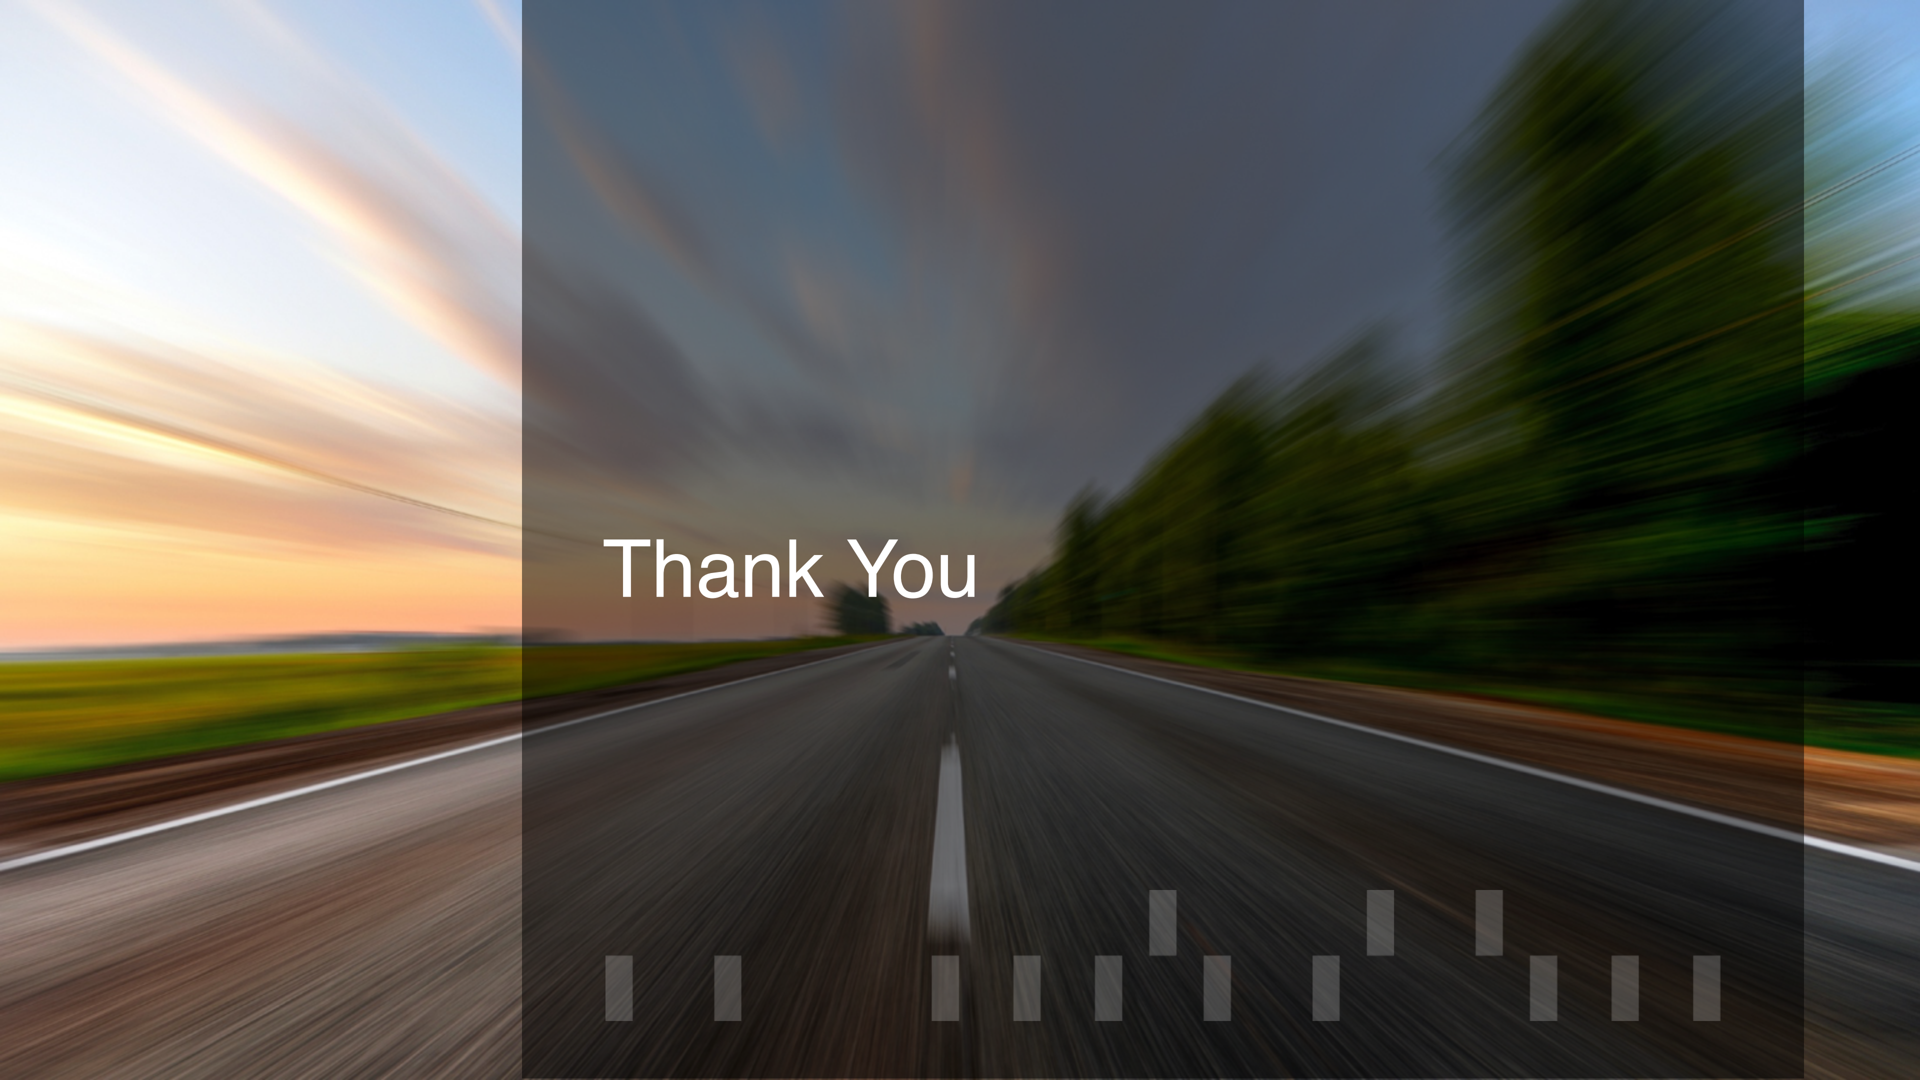

Thank You