

Communication Networks

Prof. Laurent Vanbever

Solution: Exercise 12 – Domain Name System (DNS)

12.1 Name it or Route it: pick one

In the course, we saw two ways to replicate and load-balance content: (i) using Anycast routing; or (ii) using DNS.

a) List and briefly justify three pros and cons of each.

Solution: DNS

(i) Advantages

- i. **Simplicity:** DNS-based load-balancing can be implemented by any CDN.
- ii. **(potentially) Fine-grained:** Decisions can be particularized on a per-source IP basis.
- iii. **Near real-time:** Assuming small TTL values (modulo the problem of load, see below), DNS load-balancing enables frequent load adaption.

(ii) Disadvantages

- i. **Infrastructure cost:** Good load-balancing requires small TTL values which induce a high load at the DNS server level, which in turn requires to dimension the DNS infrastructure accordingly.
- ii. **Location:** The source IP seen by the server and on which the load-balancing decision is made is the one from the resolver (e.g. Swisscom's one), not the direct client, meaning the resolver and the client can actually be far away from each other (think of Google's open DNS resolver).
- iii. **(potentially) Coarse-grained:** The DNS resolver source IP can actually serve a huge amount of clients which will therefore share the same load-balancing decisions.



Any clear winner?

Solution: (part 2) Anycast routing

(i) Advantages

- i. **Simplicity:** The required infrastructure is simple as the load-balancing is done by the network directly.
- ii. **Reliability:** Whenever a route fails, the network will simply converge to another replica (route).
- iii. **Expandability:** It is easy to add an additional replica, as simply an additional location needs to start advertising the respective prefix.

(ii) Disadvantages

- i. **Broken Connections:** As routing changes can happen at anytime, packets of the same TCP flow might arrive at different replicas effectively breaking the connection.
- ii. **Content and performance:** The load-balancing is not aware of the utilization of the replicas. Hence, it is not possible to offload work from heavily utilized replicas.
- iii. **Location:** Packets may not use the nearest replica due to routing policies.

- b)** We saw that CDNs often rely on DNS for distributing their load. Could they also use Anycast routing instead? Explain why or why not.

Solution: Yes, you can use Anycast routing instead to distribute the load. However, the load-balancing is neither content nor performance aware. You have to make sure that the same content is available on all replicas and none of the replicas is fully utilized.

12.2 Local DNS server

On Linux and Mac computers you can use the command line tool `dig` to perform DNS lookups. The corresponding tool for Windows is `nslookup`. First, perform a lookup for `nyu.edu` using your default DNS server by running the command `dig nyu.edu` or `nslookup nyu.edu`.

- What is the IP address of the server behind `nyu.edu`?

Solution: Note that the actual IP address can depend on the local DNS server you use. We got the following answer with `dig`:

```
;; ANSWER SECTION:  
nyu.edu. 60 IN A 216.165.47.10
```

Note that the format is slightly different for `nslookup`:

```
Non-authoritative answer:  
Name:   nyu.edu  
Address: 216.165.47.10
```

Now, perform the same lookup, but use one of the DNS root servers (e.g., `a.root-servers.net`) by running^a

```
dig @a.root-servers.net nyu.edu  
nslookup nyu.edu a.root-servers.net
```

^aIf the `nslookup` command does not yield helpful output for you, try adding `-type=soa` after `nslookup`. (For some, it may be the other way round—adding this option may hide the relevant part in the output.)

- Why does the answer differ compared to the one from your local DNS server?

Solution: The request is not sent to an open DNS resolver, but to a DNS server that only provides answers about its own zone. Therefore, the root DNS server only points you to the name servers responsible for the next zone in the hierarchy, the edu zone.

We got the following answer with dig:

```
;; AUTHORITY SECTION:
edu. 172800 IN NS a.edu-servers.net.
edu. 172800 IN NS c.edu-servers.net.
edu. 172800 IN NS d.edu-servers.net.
...
```

For nslookup, the output had this format (without -type=soa):

```
Served by:
- a.edu-servers.net
    192.5.6.30
    2001:503:a83e::2:30
    edu
- b.edu-servers.net
    192.33.14.30
    2001:503:231d::2:30
    edu
...
```

- How would you proceed with this answer to find the IP address behind nyu.edu?

Solution: Now that we know which servers are responsible for the edu zone, we can continue step-by-step just like your local DNS server would. Next, we would send a request to one of the edu name servers:

```
dig @a.edu-servers.net nyu.edu
```

The reply points us to the name servers in charge of the zone of NYU. By sending a request to them, we finally get the IP address behind the URL nyu.edu.

12.3 DNS Setup (Exam Question Summer 2021)

Assume you are the operator for the DNS sub-domain `ee.ethz.ch`. The responsible of the Networked Systems Group (NSG) reaches out to you as they want to create and independently maintain their own DNS sub-domain `nsg.ee.ethz.ch`. Among others, they want to create DNS names for:

- their group webpage `www.nsg.ee.ethz.ch`, pointing at `129.132.30.1`;
- a project webpage `fun.nsg.ee.ethz.ch`, also pointing at `129.132.30.1`;
- their chat room `chat.nsg.ee.ethz.ch`, pointing at `129.132.30.2`; and
- a mail server `mail.nsg.ee.ethz.ch`, pointing at `129.132.30.3` and which should receive any email sent to a `@nsg.ee.ethz.ch` email address.

Unfortunately, the responsible of NSG does not seem to understand much about DNS. He needs your help to set this up.

You start by setting up two DNS name servers `ns1.nsg.ee.ethz.ch` and `ns2.nsg.ee.ethz.ch` which you host on `129.132.20.1` and `129.132.20.2`, respectively.

- a) Indicate the resource records that these two DNS servers should store. For each record, indicate its corresponding name, type, and value. *Note you might not need all 6 records.*

Solution:

- `www.nsg.ee.ethz.ch A 129.132.30.1`
- `fun.nsg.ee.ethz.ch A 129.132.30.1`
- `chat.nsg.ee.ethz.ch A 129.132.30.2`
- `nsg.ee.ethz.ch MX mail.nsg.ethz.ch`
- `mail.nsg.ethz.ch A 129.132.30.3`

The group webpage, project webpage, and chat room each need an IPv4 address record (A) s.t. the name server can map their respective domain names to the corresponding IPv4 address. For the mail address resolution, the name servers needs an MX record that points to the mail server's name, and again a corresponding IPv4 address record that resolves the mail server's name to an address.

- b) Indicate the resource records that need to be added (if any) to the name servers responsible for `ee.ethz.ch`. Briefly explain the usage of these extra records or why no such extra record is needed. *Note you might not need all 4 records.*

Solution:

- `nsg.ee.ethz.ch NS ns1.nsg.ee.ethz.ch`
- `nsg.ee.ethz.ch NS ns2.nsg.ee.ethz.ch`
- `ns1.nsg.ee.ethz.ch A 129.132.20.1`
- `ns2.nsg.ee.ethz.ch A 129.132.20.2`

- c) Indicate the resource records that need to be added (if any) to the name servers responsible for `ethz.ch`. Briefly explain the usage of these extra records or why no such extra record is needed. *Note you might not need all 4 records.*

Solution: No new entries are required: This name server only needs to know the IP address the name server for `ee.ethz.ch`. Since this subdomain was already in use beforehand, no new entries are needed. Concretely, the `ethz.ch` name server will point to the `ee.ethz.ch` name server during address resolution, and the `ee.ethz.ch` name server will then take care of pointing to the NSG name servers. The NSG name servers can then resolve the custom NSG subdomains to the correct IP address. (If any of the involved name servers already have the record cached, they can of course directly reply with it and the resolution stops.)

- d) While setting up the records, you observe that NSG wants to host two webservers (for `www` and `fun`) on the same IP address (`129.132.30.1`). Is that possible? If so, briefly explain how this works. If not, briefly explain why.

Solution: Yes, it is possible with virtual hosting: Here, the requesting side puts the target host name into the HTTP Host header. With this, the server hosting the two webservers can distinguish between the requests and deliver the correct response.

- e) It looks like you underestimated NSG's popularity! It appears the two name servers `ns1.nsg.ee.ethz.ch` and `ns2.nsg.ee.ethz.ch` are completely overloaded with (legitimate) DNS requests.

Explain two *distinct* techniques you could use to scale the system further.

Solution: Technique 1: For each name server, one can add more devices with different IPs. Then, one installs new A records in the `ee.ethz.ch` name server, s.t. the NSG name servers' domains map to several IPs. The `ee.ethz.ch` name server then round-robins between the different IPs. This has a load-balancing effect.

Technique 2: One can add another nameserver (`ns3`), and again add the corresponding A records to the `ee.ethz.ch` name server.

(Many more techniques are possible, e.g. using a CDN)

12.4 Local vs. authoritative DNS server

Perform a DNS query for `uzh.ch` using first the authoritative DNS server (`ns1.uzh.ch`) and then your local server.

Note: When using `nslookup` on Windows, you need to specify the `-debug` flag to get the relevant information for this task. For example:

```
nslookup -debug <Domain Name> <DNS-Server>
```

- Compare the ANSWER SECTION of the responses. Can you see differences between the answers from your local DNS server and the authoritative server? Run the query to your local server multiple times to make the differences more obvious.

Solution: The answers differ in the time to live (TTL). While the TTL is constant in the replies from the authoritative DNS server, it varies in the replies from the local server.

- What is the reason for this difference?

Solution: The local DNS server caches replies to requests. To ensure that it does not keep outdated information in its cache, each authoritative name server attaches a TTL to its replies. The TTL tells the local DNS server how long it can store the reply in the cache and use it to reply to requests.

- As you have seen in the lecture, DNS can be used to balance the incoming load. What are the considerations one has to make when using DNS load balancing with respect to the TTL?

Solution: With low TTLs we can ensure that we can shift the load quickly. However, low TTLs also mean that our authoritative DNS server will get many more requests.