DNS 101

Tuấn-Anh Nguyễn

December 8, 2021

DNS Basics

- Phone books (White/Yellow Pages) for machines
- UDP requests over port 53
- Specialized CLIs: `dig`, `host`, `nslookup`, `dscacheutil`
- OS and prog langs: getaddrinfo syscall

Record Types

- `A`: Name to IPs
- `CNAME`: Alias (like symlink)
- `MX`: Where to send mail to
- `TXT`: "Application-defined"
  - Domain ownership verification
  - Email stuff: `SPF`, `DKIM`, `DMARC`, `BIMI`
- `NS`: Authority name servers

Misc

- Hierarchical zones

```
dig NS parcelperform.com.
dig NS com.
dig NS vn.
```

- Caching
  - Clients, intermediate servers
  - TTL and "propagation"
- PTR records: IP to name

```
dig +short parcelperform.com
dig +short parcelperform.com | xargs -I {} dig +short -x {}
```

DNS in k8s

- Client: /etc/resolv.conf is injected in each pod

```
nameserver 172.20.0.10
search default.svc.cluster.local svc.cluster.local cluster.local eu-west-1.compute.internal
options ndots:5
```

- Server: CoreDNS
  - `cluster.local`
  - Forwards the rest to VPC DNS (X.X.X.2)
- Server: VPC DNS
  - `eu-west-1.compute.internal`
  - `amazonaws.com`
  - `pp-local.prod`
  - `parcelperform.com`

DNS in k8s - Troubles

- VPC DNS: throttles per EC2 instance
  - Spread CoreDNS pods around
- Client: IPv6 enabled in recent kernels, lots of wasted `AAAA`
  - Make CoreDNS drop them all
- Client: search domains, high ndots (Linux's default is only 1)
  - End with a dot to signify FQDN
  - Reduce `ndots`
  - Make CoreDNS drop `com.eu-west-1.compute.internal`

Misc - Troubleshooting

- Swiss-army container image for netshoot: nikolaka/netshoot

```
kubectl run my-pod --rm -i --tty --image nicolaka/netshoot -- /bin/bash
```

- tcpdump works on UDP as well

```
tcpdump -i eth0 -l -vvvv dst port 53 &
```

- Use `nc`, `curl`, `ping`, since DNS CLI tools don't use getaddrinfo

```
nc -vz pgbouncer-89c8148d0777dc19.elb.eu-west-1.amazonaws.com 5432
curl jobmanager-test.pp-flink-test:8081/jobs
ping pp-flink.pp-local.prod
```

- CoreDNS metrics