OSS diversity is good

(aka why we run various DNS servers flavors in CentOS.org infra)

Fabian Arrotin

arrfab@centos.org, @arrfab

/whois arrfab

['ops','infra','floor sweeper'] @ centos.org

Agenda

CentOS.org Infra DNS usage:

- Bind
- PowerDNS
- Unbound
- Dnsmasq

Ansible Roles to deploy/manage those

Bind

Legacy

```
whois centos.org|egrep "Creation Date"
Creation Date: 2003-12-04T12:28:30Z
```

Bind

Started small , only one zone : centos.org

- Few donated machines
- Small needs
- Small changes
- Managed by Puppet/SVN

Bind (today)

```
[arrfab@ns1 ~]$ rpm -q bind-chroot
bind-chroot-9.11.4-9.P2.el7.x86_64
```

- Several zones
- Ansible controlled
- Delegation for some Wildcards
  - (example) apps.ci.centos.org (Openshift cluster for CI)

Bind

Specific delegation for ACME protocol

LetsEncrypt use-case

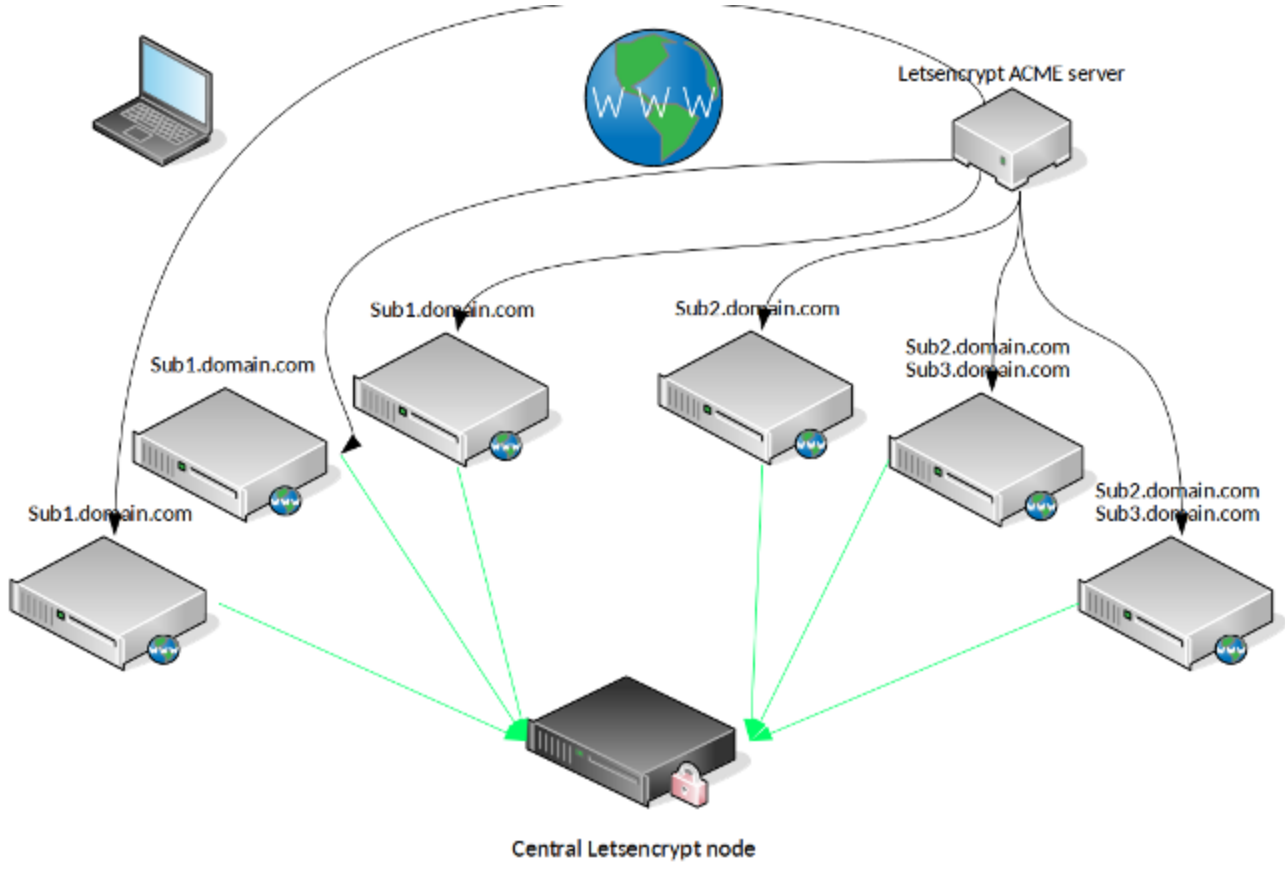

https://arrfab.net/posts/2016/May/03/generating-multiple-certificates-with-letsencrypt-from-a-single-instance/

Bind

Enters ACME v2 (including Wildcard)

https://github.com/Neilpang/acme.sh

- supports http challenges, but also DNS
- supports CNAME too !

Bind

Idea :

- Still have main centos.org 'static' zone served from git/ansible
- Create CNAME for _acme-challenge.centos.org => _acme-challenge.acme.centos.org
- Delegates acme.centos.org
- Have that acme.centos.org zone 'dynamic' and controlled by acme.sh

Bind

Normal cert with multiples SANs

```
acme.sh --issue  \
  -d  centos.org --challenge-alias acme.centos.org \
  --dns dns_nsupdate \
  -d  node1.centos.org \
  -d  node2.centos.org
```

Bind

Wildcard example :

```
acme.sh --issue  \
  -d  *.stg.centos.org --challenge-alias acme.centos.org \
  --dns dns_nsupdate \
  -d  *.dev.centos.org
```

Bind

Wildcard cert obtained/deployed through dns challenge

```
host="id.dev.centos.org"
openssl s_client -host $host -port 443 -showcerts </dev/null 2>/dev/null \
  | sed -n '/BEGIN CERTIFICATE/,/END CERT/p' \
  | openssl x509 -text -noout|grep DNS

  DNS:*.dev.centos.org, DNS:dev.centos.org
```

Some Acme.sh useful links for DNS acme challenges

- https://github.com/Neilpang/acme.sh/wiki/dnsapi
- https://github.com/Neilpang/acme.sh/wiki/DNS-alias-mode

PowerDNS

first usage (time machine !)

mirror.centos.org case

- More and more nodes donated to Project
- All around the world
- How to avoid just RR and use GeoIP ?

PowerDNS

Delegation of some host/role to pdns auth servers

```
dig @ns1.centos.org mirror.centos.org

; <<>> DiG 9.11.4-P2-RedHat-9.11.4-17.P2.el8_0.1 <<>> @ns1.centos.org mirror.cent
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 51460
;; flags: qr rd; QUERY: 1, ANSWER: 0, AUTHORITY: 3, ADDITIONAL: 4
;; WARNING: recursion requested but not available

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
; COOKIE: c6426bf327b75ff964a7f1565ddf63dffddacdb4d00679fa (good)
;; QUESTION SECTION:
;mirror.centos.org.        IN    A

;; AUTHORITY SECTION:
mirror.centos.org.    600    IN    NS    pdns1.centos.org.
mirror.centos.org.    600    IN    NS    pdns2.centos.org.
mirror.centos.org.    600    IN    NS    pdns3.centos.org.
```

PowerDNS

Entered PowerDNS with Pipe backend !

(https://doc.powerdns.com/md/authoritative/backend-pipe/)

At that time perl based backend (based on the backend.pl example)

PowerDNS today (switched in 2019)

Still Pipe, but switched to python :

- Central DB (also controlled by Zabbix/monitoring)
- Add/remove hosts from DB (one "hostname" = one column in DB with true/false)
- Reproduces backend.json / GPG encrypt it
- PowerDNS pipe detects new backend.json, reload

PowerDNS backend.json

```
{
    "mirror": {
        "AF": {
            "ipv4": [],
            "ipv6": []
        },
        "NA": {
            "ipv4": [
                "192.168.1.1",
                "192.168.2.2"
            ],
            "ipv6": [
                "::2",
                "::3"
            ]
        },
      }
}
```

https://github.com/CentOS/pdns-custom-geoip-backend

Unbound

What about resolvers ?

- Very fast
- good caching as resolver
- supports DoT
- Specific record overrides (no need for full zone)

Unbound

record overrides

```
local-data: "mirror.centos.org. IN A 172.31.234.10"
local-data-ptr: "172.31.234.10 mirror.centos.org"

local-data: "mirrorlist.centos.org. IN A 172.31.234.11"
local-data-ptr: "172.31.234.11 mirrorlist.centos.org"
```

Disclaimer : Bind can now also do that (since > 9.8.1)

RPZ (Response Policy Zone)

```
zone “rpz” {
 (…)
}

options {
 (…)
 response-policy { zone “rpz”; } ;
}
```

Unbound

- iptables/firewalled so that only centos.org infra nodes can use those resolvers
- roadmap:
  - enable DoT (RFC7858 / Dns Over TLS)
  - use stubby (getdns-stubby pkg on CentOS/epel) for now
  - use systemd-resolved (eventually, but needs > v239)

Dnsmasq

For very small environments Initially we used it for small DC setup example : ci.centos.org when it started (small)

- dhcp
- tftp
- dns
  - adding host was just 'echo $host $ip >> /etc/hosts && service dnsmasq reload'

Ansible roles for CentOS.org infra:

- https://github.com/CentOS/ansible-role-bind
- https://github.com/CentOS/ansible-role-pdns-pipe
- https://github.com/CentOS/ansible-role-unbound

Q&A