

Open Infrastructure - Task #7180

Task # 7178 (In Progress): Replace routers: router1.place5, router2.place5, router1.place6, router2.place6

Phase in new router1.place6

09/28/2019 10:48 AM - Nico Schottelius

Status:	Closed	Start date:	09/28/2019
Priority:	Normal	Due date:	
Assignee:	Nico Schottelius	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			
PM Check date:			

Description

- router1-new is up and running

Steps

- Verify / update cdist configuration
- Checkout IP configuration
 - Add VRRP IPs to loopback interface (will be used instead of keepalived)
- Verify sysctl configuration
- Rerun cdist, exclude announcing of routes
- Adjust radvd configuration
 - lower interval
 - lower life time

cdist configuration

Was written for Devuan/keepalived. Need to check step-by-step.

- `__ungleich_bgp_router`: **TEMPFIX**
 - contains static route announcements
 - can only be used once the upstream networking is verified
- `__dcl_node_exporter`: still makes sense
 - Is debian/devuan specific, needs changes
- `__ungleich_ipv6_radvd --config router-$dc`
 - 80% done
 - restart needs to be implemented
- `__ungleich_nftables --config router-$dc`
- `--__dcl_router_resolvconf --search-path "$d"--`: done
 - no changes required
- `__dcl_tftp_http_pxe_bootserver`
 - ok, needs to be tested for alpine
 - needs to be updated for ipv6 only boot
 - could be moved to APUs
- ~~require="__dcl_router_network"__dcl_router_keepalived__master~~: done
 - Needs to be ported to loopback + bgp
 - ~~commented-out statement for place6 routers~~: done
- `__dcl_router_network --ipsuffix 5`
 - Need to add keepalived IPs
- `# __dcl_nat64 --mtu 9000`

Networking

- Update interfaces for alpine
 - ~~place6~~
 - place5

Implementing virtual IPs with BGP

- Announce the IPv6/128 IPv4/32 address
 - Or network alone is enough?
- ~~Add the addresses to the loopback interface:~~ done
 - all done locally on router1-new
 - Needs to be put back into cdist
- Add addresses for place5

Firewalling

Change router advertisements

We will have 2 active routers now with both advertising with radvd at the same time. This is part of replacing keepalived.

- ~~Change router advertisement lifetime to 10 seconds~~
- ~~Change advertisement interval to 3..5~~
- Modify cdist type to restart/reload radvd on changes

```
MinRtrAdvInterval 3;
MaxRtrAdvInterval 5;
AdvDefaultLifetime 10;
```

- Good documentation: <https://packetlife.net/blog/2011/apr/18/ipv6-neighbor-discovery-high-availability/>

IPv4 NAT session table

- With two active routers, the masquerading session information is only stored on one router
- If reply is received by other router, no session table match is found
- This only effects client devices
 - VMs are using direct public IPv4

Switch configuration

- Switches **might** need ipv4 bgp peering for providing virtual IPv4 address support
- Don't see traffic on router1-new from upstream
 - Need to verify trunk configuration on switches

```
switch5-place6(config)#vlan 100
switch5-place6(config-vlan-100)#name netstream
switch5-place6(config-vlan-100)#write
Copy completed successfully.
```

```
switch6-place6#configure
switch6-place6(config)#vlan 100
switch6-place6(config-vlan-100)#name netstream
switch6-place6(config-vlan-100)#write
Copy completed successfully.
```

Building jool

Requirements:

```
apk upgrade
apk add alpine-sdk
apk add libnl3-dev
apk add iptables-dev
apk add linux-vanilla-dev
apk add argp-standalone
reboot
```

Build:

```
wget https://github.com/NICMx/Jool/releases/download/v4.0.1/jool_4.0.1.tar.gz
tar xvfz jool_4.0.1.tar.gz
cd jool-4.0.1/
```

```
./configure && make && make install
```

Missing routes from iBGP

- new router1.place6 does not receive full routing table from peers in place5
- Error "Invalid NEXT_HOP attribute" appears in the config file
 - Good documentation:
 - <https://blog.cybozu.io/entry/bird> (use translate...)
 - It seems bird1.6 uses the peer's ip address for the route while
 - bird2 uses the "correct" next hop address inside the route
- Solution for bird2: "next hop self ebgp;"

Consul

- Missing: correct init script: not compatible w/ alpine
- Started manually -> works

```
[15:21] line:~% echo __dcl_consul_agent --datacenter place6 --server black1.place6.ungleich.ch --
server black2.place6.ungleich.ch --server black3.place6.ungleich.ch | cdist config -i - -vv -c ~
/vcs/ungleich-dot-cdist/ -j8 -p8 router1.place6.ungleich.ch
```

Node exporter / monit

- monit: need to modify /etc/monitrc to include include /etc/monit/conf.d/*
- need to create mkdir /etc/monit/conf.d -p
- Missing init script /etc/init.d/prometheus-node-exporter
 - need to add this into cdist and or alpine package
 - Alternative: modify the node-exporter monit configuration

```
[15:42] line:~% echo "__directory /etc/consul/conf.d; __dcl_node_exporter" | cdist config -i - -vv
-c ~/vcs/ungleich-dot-cdist/ -j8 router1.place6.ungleich.ch
```

History

#1 - 09/28/2019 10:55 AM - Nico Schottelius

Removed all static routes from new router temporarily:

```
protocol static place6_v6 { ipv6 {};  
}  
  
protocol static place6_v4 { ipv4 {};  
}
```

#2 - 09/28/2019 11:04 AM - Nico Schottelius

- Description updated

#3 - 09/28/2019 11:26 AM - Nico Schottelius

- Description updated

#4 - 09/28/2019 11:29 AM - Nico Schottelius

- Description updated

#5 - 09/28/2019 11:45 AM - Nico Schottelius

- Description updated

#6 - 09/28/2019 11:50 AM - Nico Schottelius

- Description updated

#7 - 09/28/2019 11:55 AM - Nico Schottelius

- Description updated

#8 - 09/28/2019 09:37 PM - Nico Schottelius

- Description updated

#9 - 09/29/2019 10:49 AM - Nico Schottelius

- Description updated

#10 - 09/29/2019 11:26 AM - Nico Schottelius

- Description updated

#11 - 09/29/2019 11:27 AM - Nico Schottelius

- Description updated

#12 - 09/29/2019 11:29 AM - Nico Schottelius

- Description updated

#13 - 09/29/2019 01:24 PM - Nico Schottelius

- Description updated

#14 - 09/29/2019 03:32 PM - Nico Schottelius

- Description updated

#15 - 09/29/2019 03:32 PM - Nico Schottelius

- Description updated

#16 - 09/29/2019 05:27 PM - Nico Schottelius

- Description updated

#17 - 09/30/2019 02:17 PM - Nico Schottelius

- Restarted bird with limitations:
 - no peering with iBGP routers
 - no announcements of routes to upstream

#18 - 09/30/2019 02:24 PM - Nico Schottelius

- added full routing table to router1.place6
 - Just "conf" was not enough to add it to the kernel, needed a full restart

#19 - 09/30/2019 03:12 PM - Il nu

- Description updated

#20 - 09/30/2019 03:26 PM - Nico Schottelius

- Description updated

#21 - 09/30/2019 03:26 PM - Nico Schottelius

- Description updated

#22 - 09/30/2019 03:36 PM - Nico Schottelius

- Description updated

#23 - 09/30/2019 03:49 PM - Nico Schottelius

- Description updated

#24 - 09/30/2019 03:53 PM - Nico Schottelius

- router1.place6 is hacked back into prometheus -> monitoring is active. We (=Balazs) need to make it clean

#25 - 08/09/2021 03:19 PM - Nico Schottelius

- Status changed from In Progress to Closed

Done long time ago.