

Open Infrastructure - Task #7156

Add new sunrise link to infrastrucure, introduce policy routing

09/18/2019 07:11 AM - Nico Schottelius

Status:	Closed	Start date:	09/18/2019
Priority:	Normal	Due date:	
Assignee:	Nico Schottelius	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			
PM Check date:			
Description			
<ul style="list-style-type: none">• We get a new uplink from Sunrise<ul style="list-style-type: none">◦ configure interfaces / peering• We need to add the full routing table now<ul style="list-style-type: none">◦ We don't receive a default route anymore◦ Adding full routing table caused router1.place6 to go to 100% cpu usage on various softirqd threads<ul style="list-style-type: none">▪ Needs debugging◦ Phase in router2.place6 for taking some traffic part by part and see when/how/if cpu spikes there, too• Init7 link will be temporarily off for about 1-3 months<ul style="list-style-type: none">◦ Will be back as a 10 gbit/s link afterwards◦ Traffic will be handled by netstream and sunrise links meanwhile			
Related issues:			
Related to Open Infrastructure - Task #7160: Phase in 2 new routers to avoid ...		Closed	09/20/2019

History

#1 - 09/18/2019 07:13 AM - Nico Schottelius

- router2.place6 takes care of 100% of outgoing IPv6 traffic
- Phasing in IPv4 routing by preferring parts of the Internet step by step (see below)
 - We split the IPv4 Internet into 8 parts and change 1/8 of the traffic each time to go onto router2
 - Meanwhile running perf to find out the bottleneck
 - Might be related to https://www.linuxquestions.org/questions/linux-server-73/high-cpu-load-by-native_queued_spin_lock_slowpath-4175615354/

```
[14:08] line:~% sipcalc -s /3 0.0.0.0/0  
-[ipv4 : 0.0.0.0/0] - 0
```

```
[Split network]  
Network      - 0.0.0.0      - 31.255.255.255  
Network      - 32.0.0.0     - 63.255.255.255  
Network      - 64.0.0.0     - 95.255.255.255  
Network      - 96.0.0.0     - 127.255.255.255  
Network      - 128.0.0.0    - 159.255.255.255  
Network      - 160.0.0.0    - 191.255.255.255  
Network      - 192.0.0.0    - 223.255.255.255  
Network      - 224.0.0.0    - 255.255.255.255
```

#2 - 09/18/2019 07:14 AM - Nico Schottelius

- Description updated

#3 - 09/18/2019 07:14 AM - Nico Schottelius

- Description updated

#4 - 09/18/2019 08:36 AM - Nico Schottelius

- 2019-09-18, ca. 1530 KST: enabled 0.0.0.0/3 to go through router2.place6 with full routing table
- 2019-09-18, ca. 1600 KST: enabled 32.0.0.0/3 to go through router2.place6 with full routing table

#5 - 09/20/2019 12:45 AM - Nico Schottelius

- Related to Task #7160: Phase in 2 new routers to avoid softirqd problem added

#6 - 09/20/2019 07:00 PM - Nico Schottelius

Configuring network on router2.place5:

```
# Sunrise
auto bond0.102
iface bond0.102 inet6 static
    address 2001:1700:3500::12
    netmask 64
    post-up /sbin/ip link set $IFACE mtu 9000
    vlan-raw-device bond0
```

```
iface bond0.102 inet static
    address 195.141.230.103
    netmask 255.255.255.254
    vlan-raw-device bond0
```

```
[18:54:33] router2.place5:/etc/network/interfaces.d# ifup bond0.102
```

```
[18:54:42] router2.place5:/etc/network/interfaces.d#
```

Switch config

```
Et41                connected    1          a-full a-1G    1000BASE-SX
```

```
switch1-place5(config)#interface ethernet 41
switch1-place5(config-if-Et41)#switchport trunk allowed vlan 101,102
switch1-place5(config-if-Et41)#switchport mode trunk
switch1-place5(config-if-Et41)#write
Copy completed successfully.
switch1-place5(config-if-Et41)#
```

#7 - 09/20/2019 07:03 PM - Nico Schottelius

```
switch1-place5(config-if-Et41)#vlan 101
switch1-place5(config-vlan-101)#name sunrise_router1
switch1-place5(config-vlan-101)#vlan 102
switch1-place5(config-vlan-102)#name sunrise_router2
switch1-place5(config-vlan-102)#write
Copy completed successfully.
```

Receiving packets from peer:

```
64 bytes from 2001:1700:3500::11: icmp_seq=296 ttl=64 time=0.977 ms
64 bytes from 2001:1700:3500::11: icmp_seq=297 ttl=64 time=1.01 ms
64 bytes from 2001:1700:3500::11: icmp_seq=298 ttl=64 time=0.988 ms
64 bytes from 2001:1700:3500::11: icmp_seq=299 ttl=64 time=0.997 ms
64 bytes from 2001:1700:3500::11: icmp_seq=300 ttl=64 time=1.00 ms
```

```
64 bytes from 195.141.230.102: icmp_seq=474 ttl=255 time=0.937 ms
64 bytes from 195.141.230.102: icmp_seq=475 ttl=255 time=0.943 ms
```

#8 - 09/20/2019 07:15 PM - Nico Schottelius

Saving config, updating with new bgp type

```
[19:04:41] router2.place5:/etc/bird# rsync -av /etc/bird/ ~/bird20190921
sending incremental file list
created directory /root/bird20190921
./
bird.conf
bird.conf.ucf-dist
bird6.conf
bird6.conf.ucf-dist
envvars
general.conf
init7_v4.conf
init7_v6.conf
router1_place5_ungleich_ch_v4.conf
router1_place5_ungleich_ch_v6.conf
router1_place6_ungleich_ch_v4.conf
router1_place6_ungleich_ch_v6.conf
router2_place6_ungleich_ch_v4.conf
```

```
router2_place6_ungleich_ch_v6.conf
router_id.conf
switch1_place10_ungleich_ch_v6.conf
switch1_place5_ungleich_ch_v6.conf
switch1_place6_ungleich_ch_v6.conf
switch1_place8_ungleich_ch_v6.conf
switch2_place5_ungleich_ch_v6.conf
switch2_place6_ungleich_ch_v6.conf
switch2_place8_ungleich_ch_v6.conf
switch3_place6_ungleich_ch_v6.conf
switch4_place6_ungleich_ch_v6.conf
switch5_place6_ungleich_ch_v6.conf
switch6_place6_ungleich_ch_v6.conf
ungleich-routes-place5_v4.conf
ungleich-routes-place5_v6.conf
```

```
sent 21,681 bytes received 592 bytes 44,546.00 bytes/sec
total size is 19,302 speedup is 0.87
[19:04:48] router2.place5:/etc/bird#
```

#9 - 12/05/2019 06:09 PM - Nico Schottelius

- Status changed from In Progress to Closed

Added long time ago