ICN – Examination date: 30-09-2015 – "Balloons" **Available time: 180 minutes**

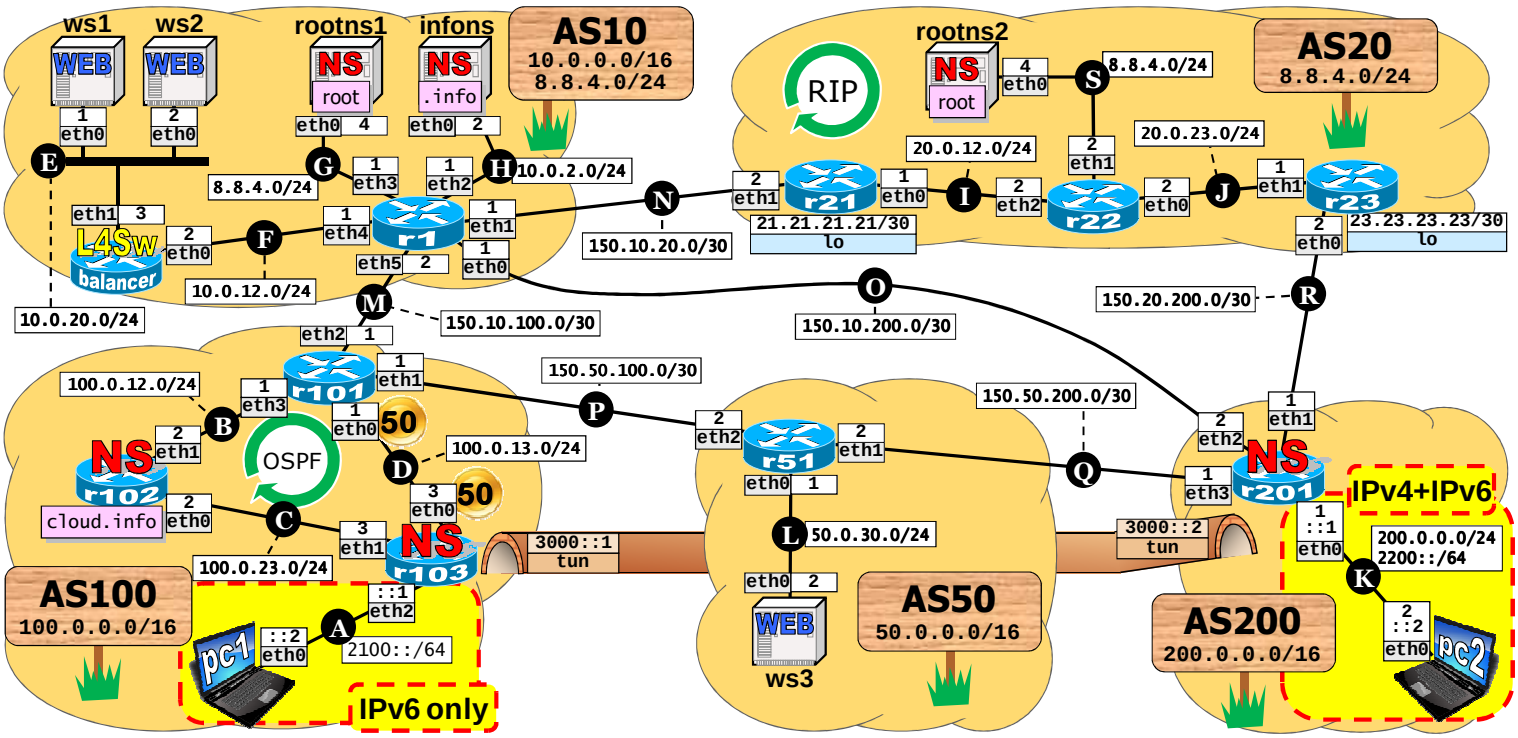

Using Netkit, implement the network scenario depicted in the above figure and described below (you can use the following items as a checklist).

Routing & IGPs

- ☐ Remember to configure a default route where required.
- ☐ No routers announce the default route `0.0.0.0/0` or any IPv6 subnets.
- ☐ Routing within **AS100** is implemented by using OSPF. All the interfaces belong to area **0.0.0.0**, and the indicated costs are assigned to the interfaces of **r101** and **r103**.
- ☐ **AS20** is a transit AS that uses RIP. Therefore, **r21** and **r23** redistribute eBGP as well as loopback addresses within RIP.

BGP

- ☐ Every border router announces the peering subnets besides the subnets of the AS it belongs to (indicated in the AS number sign).
- ☐ **AS20**'s border routers establish an iBGP peering using loopback interfaces (use **update-source** *IPaddress*).
- ☐ **AS50** is a customer of **AS100** and **AS200**. As such, it forbids transit traffic.
- ☐ **AS200** prefers using link **R** for outgoing traffic.

Web

- ☐ **ws1**, **ws2**, and **ws3** are Web servers running Apache; they serve a default Web page, different for each server.
- ☐ **balancer** is a layer-4 switch having VIP **10.0.12.2**, which realizes a round-robin policy implemented by the following configuration:

```
iptables -t nat -A PREROUTING -d 10.0.12.2 -m statistic --mode nth --every 2 --jump DNAT --to-destination 10.0.20.1
iptables -t nat -A PREROUTING -d 10.0.12.2 --jump DNAT --to-destination 10.0.20.2
```

DNS

- ☐ **r103** is **pc1**'s local name server (reached over IPv6 – remember to tell bind to listen on IPv6); **r201** is **pc2**'s local name server.
- ☐ **rootns1** and **rootns2** are root name servers with anycast address **8.8.4.4**; **infons** is the authority for **info**; **r102** is the authority for **cloud.info** (pick one of its IP addresses as the name server's address).
- ☐ **pc1.cloud.info** is associated with **pc1**'s IPv6 address; **pc2.cloud.info** is associated with **pc2**'s IPv4 and IPv6 addresses; a DNS-based round-robin load balancing is implemented on **www.cloud.info**, between **AS10**'s server farm (**10.0.12.2**) and **AS50**'s server farm (**50.0.30.2**).

IPv6

- ☐ Enable IPv6 forwarding on network nodes that act as IPv6 routers.
- ☐ IPv6 routing is implemented using static routes.
- ☐ An IPv6-in-IPv4 tunnel is established between **r103**'s **eth0** interface and **r201**'s **eth3** interface.

Goal

REDISTRIBUTING EBGP ROUTING INFORMATION IN AN IGP

```
redistribute bgp route-map RouteMapName
route-map RouteMapName permit 10
   match ip next-hop prefix-list PrefixListName
ip prefix-list PrefixListName permit nextHop/32
```

SETTING UN AN IPV6 ADDRESS

```
ifconfig interface up
ifconfig interface add IPv6address/netmask
```

ENABLING IPV6 FORWARDING

```
echo 1 >/proc/sys/net/ipv6/conf/all/forwarding
```

ADDING A STATIC IPV6 ROUTE

```
route -A inet6 add IPv6net[/netmask] [gw IPv6address] [dev interface]
```

ESTABLISHING AN IPV6-IN-IPV4 TUNNEL

```
ip tunnel add tunName mode sit remote RemoteIPv4 local LocalIPv4 ttl 10
ifconfig tunName up
ifconfig tunName add LocalIPv6address[/netmask]
route -A inet6 add default dev tunName
```

IPV6 TRACEROUTE

```
traceroute6 -N 1 IPv6address
```

TELLING BIND TO LISTEN ON IPV6 (IN NAMED.CONF)

```
options { listen-on-v6 { ::/0; }; };
```