ICN – Examination date: 24-02-2015 – "HiFi"  **Available time: 180 minutes**

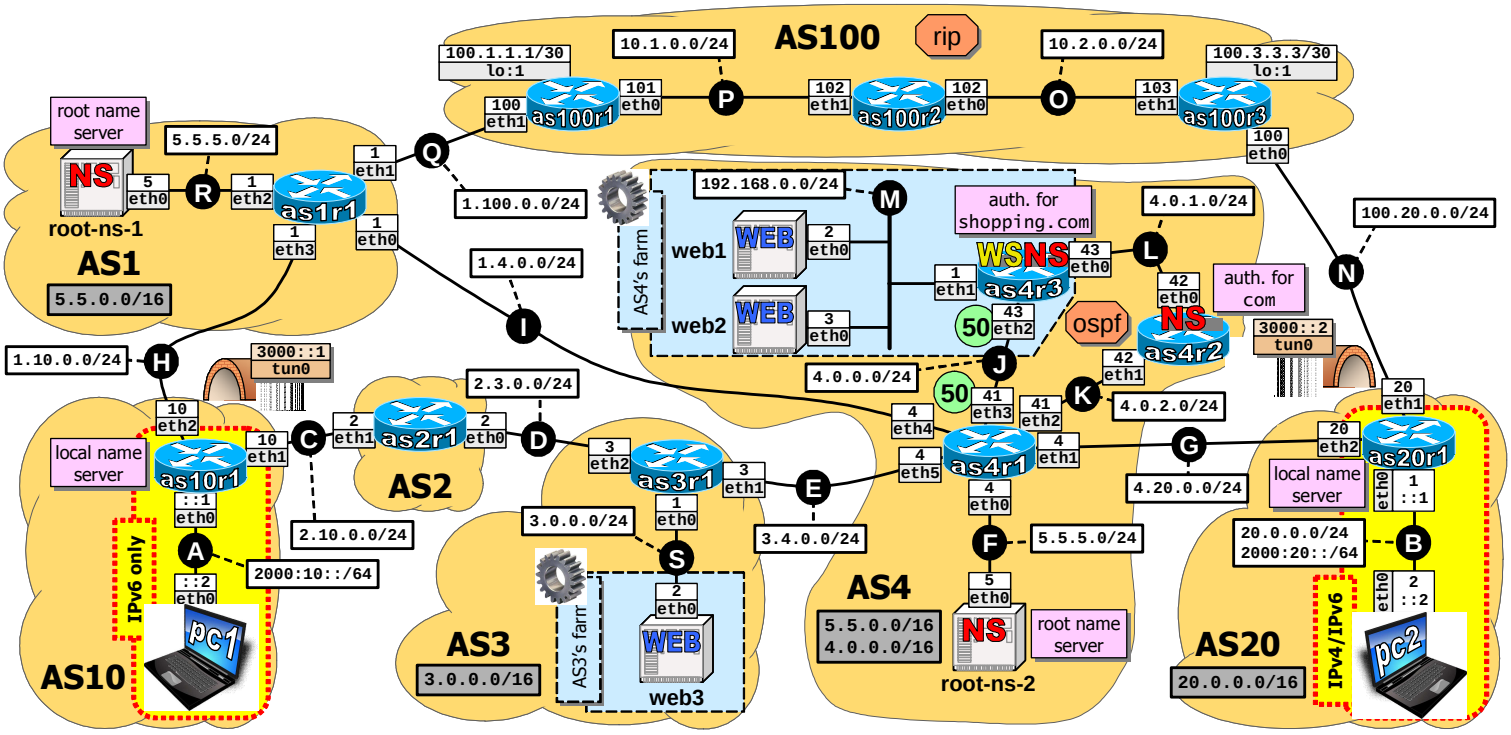

Using Netkit, implement the network depicted in the figure and described below (you can use the following as a checklist).

Routing

- ☐ Remember to set up a default route where required.
- ☐ **AS100** is a transit AS whose internal routing is implementsd by using RIP; **as100r1** and **as100r3** redistribute eBGP and loopback interface addresses in RIP.
- ☐ Routing within **AS4** is implemented by using OSPF (all interfaces belong to area **0.0.0.0**); interface costs on **as4r1** and **as4r3** are assigned as indicated.
- ☐ No routers announce **0.0.0.0/0** or IPv6 subnets.

BGP

- ☐ All peering LANs are announced. **AS1**, **AS3**, **AS4**, and **AS20** also announce their own subnets (in gray).
- ☐ **AS100**'s border routers establish an iBGP peering using their loopback interfaces (use **update-source** *IPaddress*).
- ☐ **AS10** is a customer of **AS1** and **AS2** and, as such, forbids all transit traffic.
- ☐ **AS20** is a customer of **AS100** and a peer of **AS4**. As such, it forbids all transit traffic and it prefers using link **G**.
- ☐ **AS4** is a customer of **AS1**, a peer of **AS20**, and a provider of **AS3**. As such, it forbids any transit traffic between **AS1** and **AS20**, by suitably applying the following two access-lists to outgoing announcements:

```
ip as-path access-list announcementsForAS20 deny ^1_
ip as-path access-list announcementsForAS20 permit .*
neighbor NEIGHBORIP filter-list announcementsForAS20 out
```

```
ip as-path access-list announcementsForAS1 deny ^20_
ip as-path access-list announcementsForAS1 permit .*
neighbor NEIGHBORIP filter-list announcementsForAS1 out
```

Redistributing eBGP 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 up (adding) an IPv6 address

```
ifconfig INTERFACE up
ifconfig INTERFACE add IPV6ADDR/NETMASK
```

Enabling IPv6 forwarding

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

Configuring a static IPv6 route

```
route -A inet6 add IPV6NET[/NETMASK] [gw IPV6ADDR] [dev INTERFACE]
```

Creating an IPv6-in-IPv4 tunnel

```
ip tunnel add TUNNAME mode sit remote REMOTEIPV4 local LOCALIPV4 ttl 10
ifconfig TUNNAME up
ifconfig TUNNAME add LOCALIPV6[/NETMASK]
route -A inet6 add default dev TUNNAME
```

IPv6 traceroute

```
traceroute6 -N 1 IPV6ADDR
```

Enabling bind on IPv6 (in named.conf)

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

Web

- ☐ **web1**, **web2**, and **web3** are Web servers that run Apache and serve a default web page, with different contents for each server.
- ☐ **as4r3** is a layer 4 web switch with VIP **4.0.0.43**, with a round robin policy implemented by the following configuration:

```
iptables -t nat -A PREROUTING -d 4.0.0.43 -m statistic --mode nth --every 2 --jump DNAT --to-destination 192.168.0.2
iptables -t nat -A PREROUTING -d 4.0.0.43 --jump DNAT --to-destination 192.168.0.3
```

DNS

- ☐ **as10r1** is the local name server for **pc1** (reached over IPv6 – enable bind on IPv6); **as20r1** is the local name server for **pc2**.
- ☐ **root-ns-1** and **root-ns-2** are root name servers with anycast address **5.5.5.5**; **as4r2** is the authority for **com** (pick one of its IP addresses as name server address); **as4r3** is the authority for **shopping.com** (pick one of its IP addresses as name server address).
- ☐ **pc1.shopping.com** is mapped to **pc1**'s IPv6 address; **pc2.shopping.com** is mapped to **pc2**'s IPv4 and IPv6 addresses; on name **www.shopping.com** is implemented a round-robin load balancing between **AS4**'s farm (**4.0.0.43**) and **AS3**'s farm (**3.0.0.2**).

IPv6

- ☐ Enable IPv6 forwarding on network devices that act as IPv6 routers.
- ☐ IPv6 routing is implemented by static routes.
- ☐ An IPv6-in-IPv4 tunnel is established between **as10r1**'s interface **eth2** and **as20r1**'s interface **eth1**.

Goals