ICN – Examination date: 15-07-2015 – "Fireplace"  **Available time: 180 minutes**

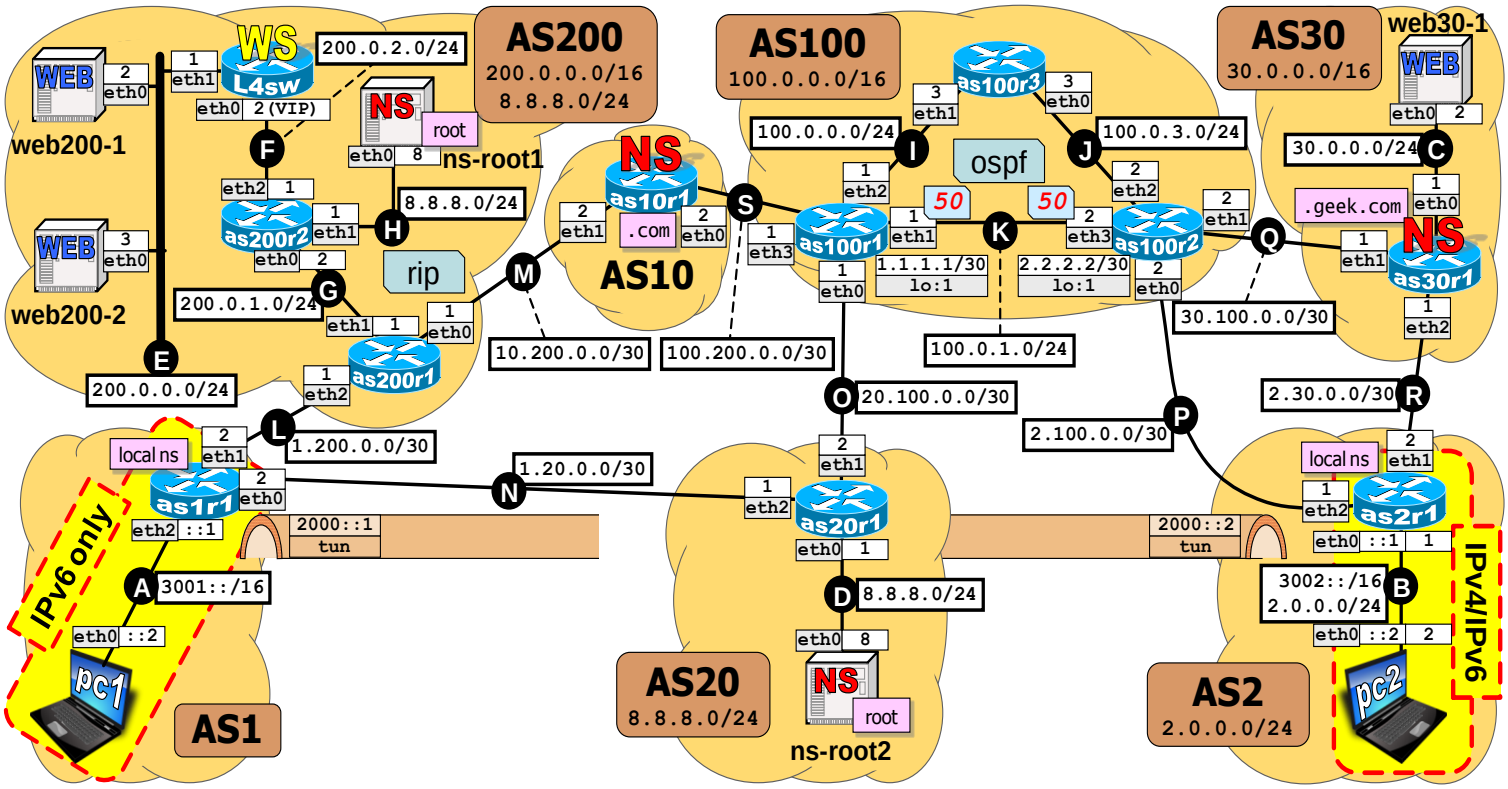

Using Netkit, implement the network depicted in the figure and described below (you can use the following items as a checklist).

Routing

- ☐ Remember to set up 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 using OSPF. All the interfaces belong to area **0.0.0.0**, and in particular the interfaces of **as100r1** and **as100r2** are assigned the indicated costs.
- ☐ **AS100** is a transit AS. Therefore, **as100r1** and **as100r2** redistribute in OSPF eBGP only, as well as the addresses of loopback interfaces.

BGP

- ☐ All peering LANs are announced in BGP. **AS2**, **AS20**, **AS30**, **AS100**, and **AS200** also announce their own subnets.
- ☐ **AS100**'s border routers establish an iBGP peering using loopback interfaces (use **update-source** *IPADDRESS*).
- ☐ **AS20** is a customer of **AS1** and **AS100**. As such, it forbids transit traffic.
- ☐ **AS2** is a customer of **AS100** and **AS30**. As such, it forbids transit traffic and prefers using link **R** for outgoing traffic.

Web

- ☐ **web30-1**, **web200-1**, and **web200-2** are Web servers running apache which serve a default page, different for each server.
- ☐ **L4sw** is a layer-4 Web switch with VIP **200.0.2.2**, which implements a round-robin load balancing policy using the following configuration:

```
iptables -t nat -A PREROUTING -d 200.0.2.2 -m statistic --mode nth --every 2 --jump DNAT --to-destination 200.0.0.2
iptables -t nat -A PREROUTING -d 200.0.2.2 --jump DNAT --to-destination 200.0.0.3
```

DNS

- ☐ **as1r1** is a local name server for **pc1** (reached via IPv6 – remember to configure bind to listen on IPv6); **as2r1** si a local name server for **pc2**.
- ☐ **ns-root1** and **ns-root2** are root name servers with anycast address **8.8.8.8**; **as10r1** is the authority for **com** (pick one of its IP addresses as the name server address); **as30r1** the authority for **geek.com** (pick one of its IP addresses as the name server address).
- ☐ **pc1.geek.com** is mapped to **pc1**'s IPv6 address; **pc2.geek.com** is mapped to both **pc2**'s IPv4 address and **pc2**'s IPv6 address; a round-robin load balancing policy is applied on name **www.geek.com**, involving **AS200**'s server farm (**200.0.2.2**) and **AS30**'s server farm (**30.0.0.2**).

IPv6

- ☐ Remember to enable IPv6 forwarding on the nodes that act as IPv6 routers.
- ☐ IPv6 routing is implemented using static routes.
- ☐ An IPv6-in-IPv4 tunnel is established between **as1r1**'s **eth0** interface and **as2r1**'s **eth2** interface.

Goals

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 AN IPV6 ADDRESS

```
ifconfig INTERFACE up
ifconfig INTERFACE add IPV6ADDR/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 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 TRACEOUTE

```
traceroute6 -N 1 IPV6ADDR
```

CONFIGURING BIND FOR LISTENING ON IPV6 (IN NAMED.CONF)

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