**ICN – Examination date: 11 Dec 2015 – Round 1 – Text A – "Left spotlight"** 

**Available time: 80 minutes.**

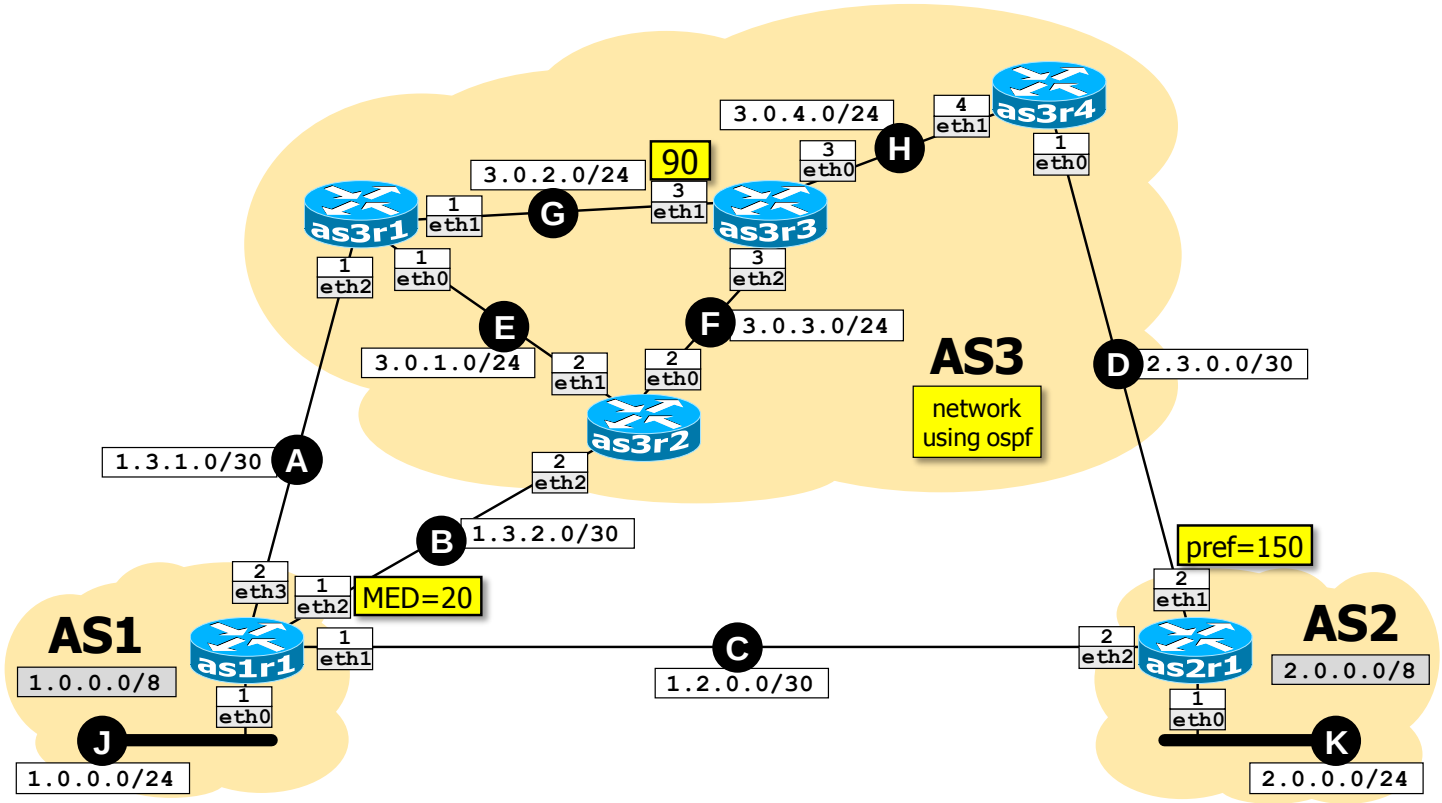


Using Netkit, implement the network depicted in the figure and described below.

- Routing within `AS3` is implemented by using OSPF, with the following setup:
  - All the routers belong to the backbone area.
  - BGP is redistributed into OSPF
    (note: enabling redistribution automatically injects into OSPF only routes learned from E-BGP).
  - Interface `eth1` of `as3r3` is assigned the indicated cost.
- The BGP configuration is as follows:
  - Border routers within `AS3` establish I-BGP peerings with each other
    (remember to establish the peerings between interfaces that are consistent with the internal routing of `AS3`).
  - All peering LANs are announced in BGP. Routers do not filter any updates. No routers announce the default route.
  - `AS1`'s and `AS2`'s border routers also announce their own subnet, as an aggregated prefix indicated in the gray box.
  - `as1r1` sets the indicated Multi Exit Discriminator (MED) value on outgoing updates.
  - `as2r1` sets the indicated local preference value on incoming updates.
- **Warning**: it may take several minutes (up to 5) for the peerings to be established and for the routing protocols to converge, even if the configuration is correct!

**Goals**: All BGP peerings as well as OSPF routing must operate correctly.
Packets from `as2r1` to `1.0.0.1` must traverse routers `as3r4`, `as3r3`, `as3r2`, and `as3r1`.
Internal routers of `AS3` do not need to be visible from outside the AS.

**ICN – Examination date: 11 Dec 2015 – Round 1 – Text B – "Eye"**

**Available time: 80 minutes.**

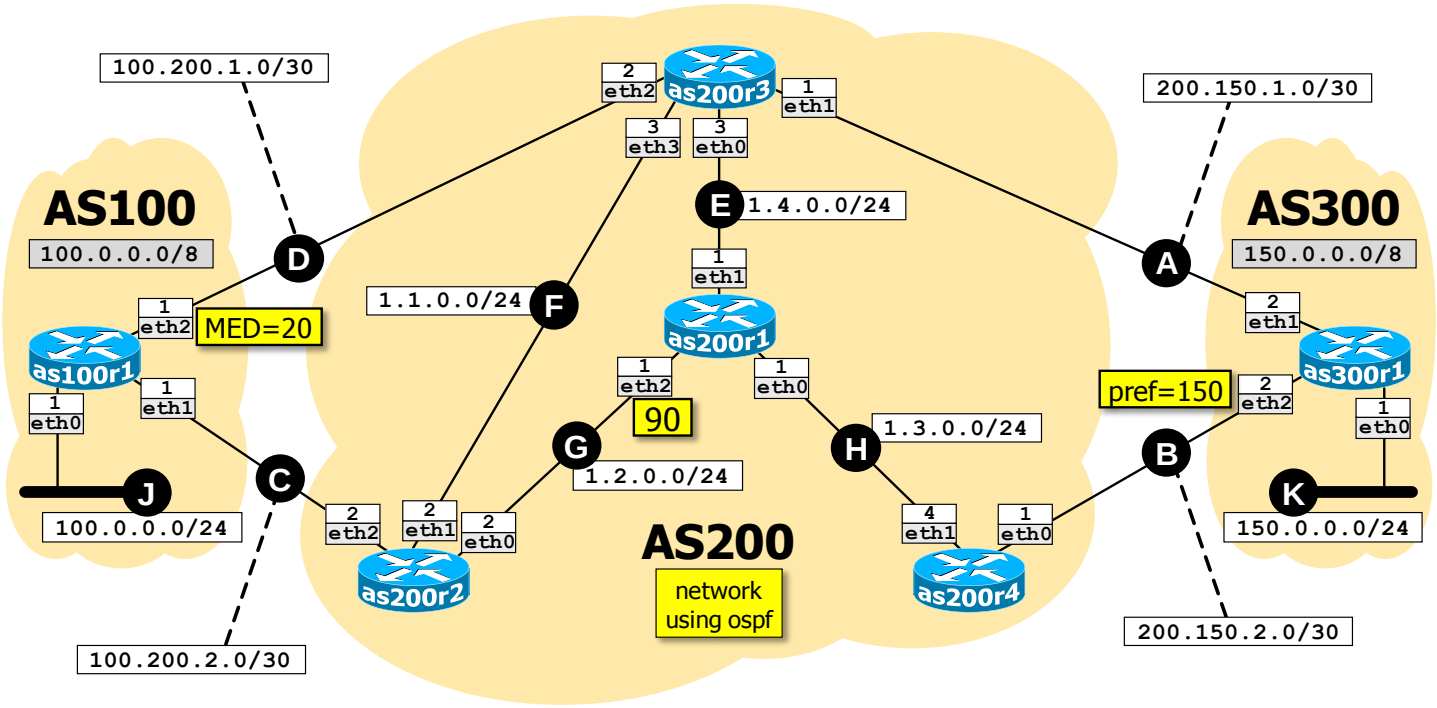


Using Netkit, implement the network depicted in the figure and described below.

- Routing within `AS200` is implemented by using OSPF, with the following setup:
  - All the routers belong to the backbone area.
  - BGP is redistributed into OSPF
    (note: enabling redistribution automatically injects into OSPF only routes learned from E-BGP).
  - Interface `eth2` of `as200r1` is assigned the indicated cost.
- The BGP configuration is as follows:
  - Border routers within `AS200` establish I-BGP peerings with each other
    (remember to establish the peerings between interfaces that are consistent with the internal routing of `AS200`).
  - All peering LANs are announced in BGP. Routers do not filter any updates. No routers announce the default route.
  - `as100r1` sets the indicated Multi Exit Discriminator (MED) value on outgoing updates.
  - `as300r1` sets the indicated local preference value on incoming updates.
- **Warning**: it may take several minutes (up to 5) for the peerings to be established and for the routing protocols to converge, even if the configuration is correct!

**Goals**: All BGP peerings as well as OSPF routing must operate correctly.
Packets from `as300r1` to `100.0.0.1` must traverse routers `as200r4`, `as200r1`, `as200r3`, and `as200r2`.
Internal routers of `AS200` do not need to be visible from outside the AS.

**ICN – Examination date: 11 Dec 2015 – Round 2 – Text A – "Right spotlight"** 

**Available time: 80 minutes.**

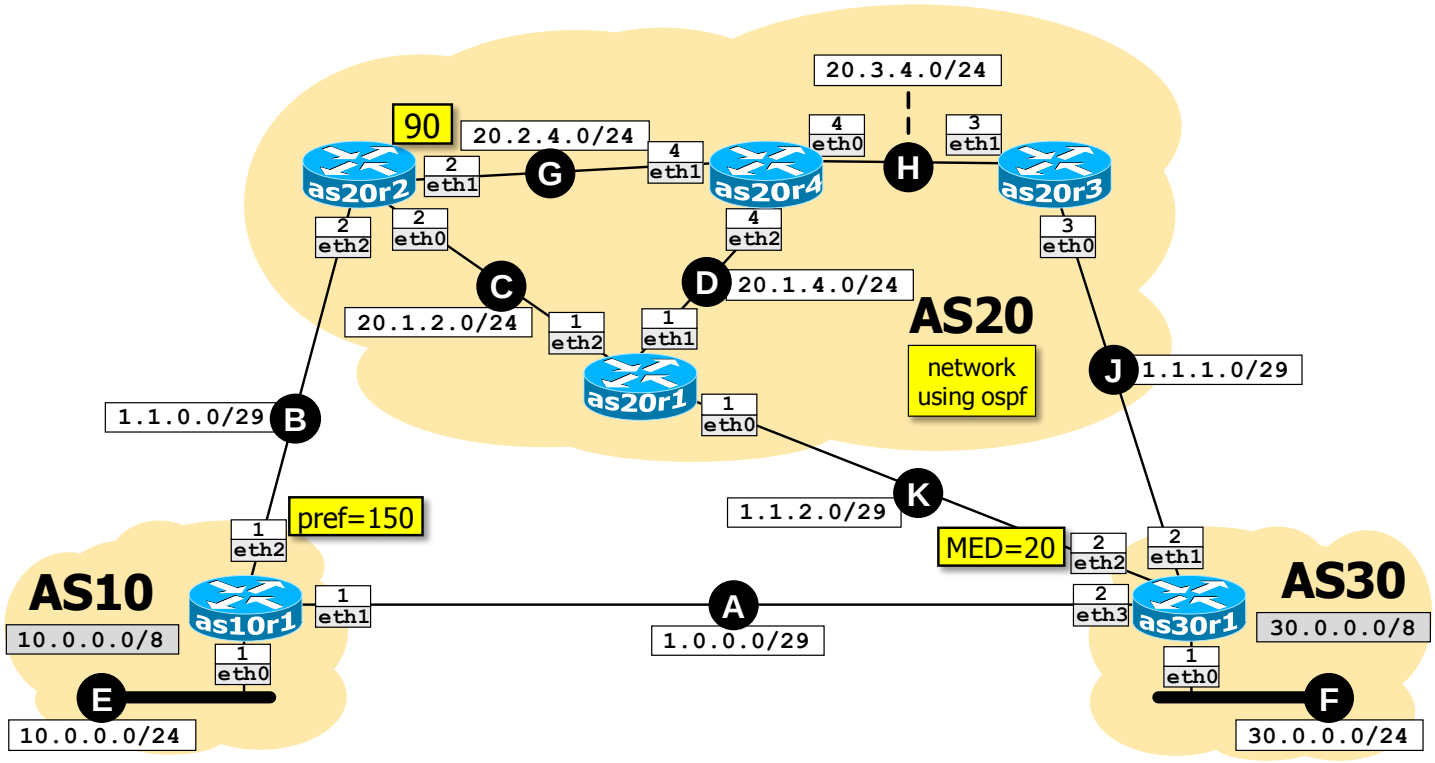


Using Netkit, implement the network depicted in the figure and described below.

- Routing within `AS20` is implemented by using OSPF, with the following setup:
  - All the routers belong to the backbone area.
  - BGP is redistributed into OSPF
    (note: enabling redistribution automatically injects into OSPF only routes learned from E-BGP).
  - Interface `eth1` of `as20r2` is assigned the indicated cost.
- The BGP configuration is as follows:
  - Border routers within `AS20` establish I-BGP peerings with each other (remember to establish the peerings between interfaces that are consistent with the internal routing of `AS20`; in particular, use the IP address of `as20r2`'s `eth0` interface for the peering between `as20r2` and `as20r3`).
  - All peering LANs are announced in BGP. Routers do not filter any updates. No routers announce the default route.
  - `AS10`'s and `AS30`'s border routers also announce their own subnet, as an aggregated prefix indicated in the gray box.
  - `as30r1` sets the indicated Multi Exit Discriminator (MED) value on outgoing updates.
  - `as10r1` sets the indicated local preference value on incoming updates.
- **Warning**: it may take several minutes (up to 5) for the peerings to be established and for the routing protocols to converge, even if the configuration is correct!

**Goals**: All BGP peerings as well as OSPF routing must operate correctly.
Packets from `as10r1` to `30.0.0.1` must traverse routers `as20r2`, `as20r1`, `as20r4`, and `as20r3`.
Internal routers of `AS20` do not need to be visible from outside the AS.

**ICN – Examination date: 11 Dec 2015 – Round 2 – Text B – "Trousers"** 

**Available time: 80 minutes.**

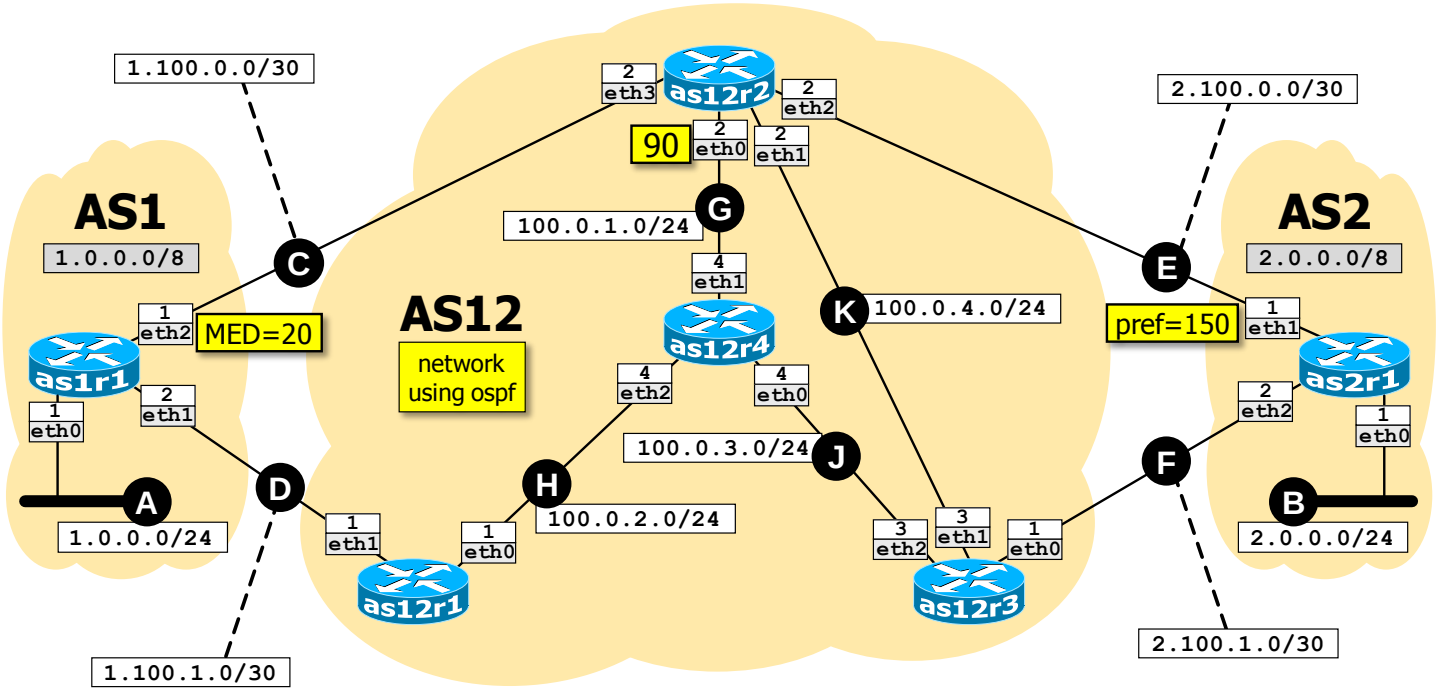


Using Netkit, implement the network depicted in the figure and described below.

- Routing within `AS12` is implemented by using OSPF, with the following setup:
  - All the routers belong to the backbone area.
  - BGP is redistributed into OSPF
    (note: enabling redistribution automatically injects into OSPF only routes learned from E-BGP).
  - Interface `eth0` of `as12r2` is assigned the indicated cost.
- The BGP configuration is as follows:
  - Border routers within `AS12` establish I-BGP peerings with each other (remember to establish the peerings between interfaces that are consistent with the internal routing of `AS12`; in particular, use the IP address of `as12r2`'s `eth1` interface for all the peerings with `as12r2`).
  - All peering LANs are announced in BGP. Routers do not filter any updates. No routers announce the default route.
  - `AS1`'s and `AS2`'s border routers also announce their own subnet, as an aggregated prefix indicated in the gray box.
  - `as1r1` sets the indicated Multi Exit Discriminator (MED) value on outgoing updates.
  - `as2r1` sets the indicated local preference value on incoming updates.
- **Warning**: it may take several minutes (up to 5) for the peerings to be established and for the routing protocols to converge, even if the configuration is correct!

**Goals**: All BGP peerings as well as OSPF routing must operate correctly.
Packets from `as2r1` to `1.0.0.1` must traverse routers `as12r2`, `as12r3`, `as12r4`, and `as12r1`.
Internal routers of `AS12` do not need to be visible from outside the AS.