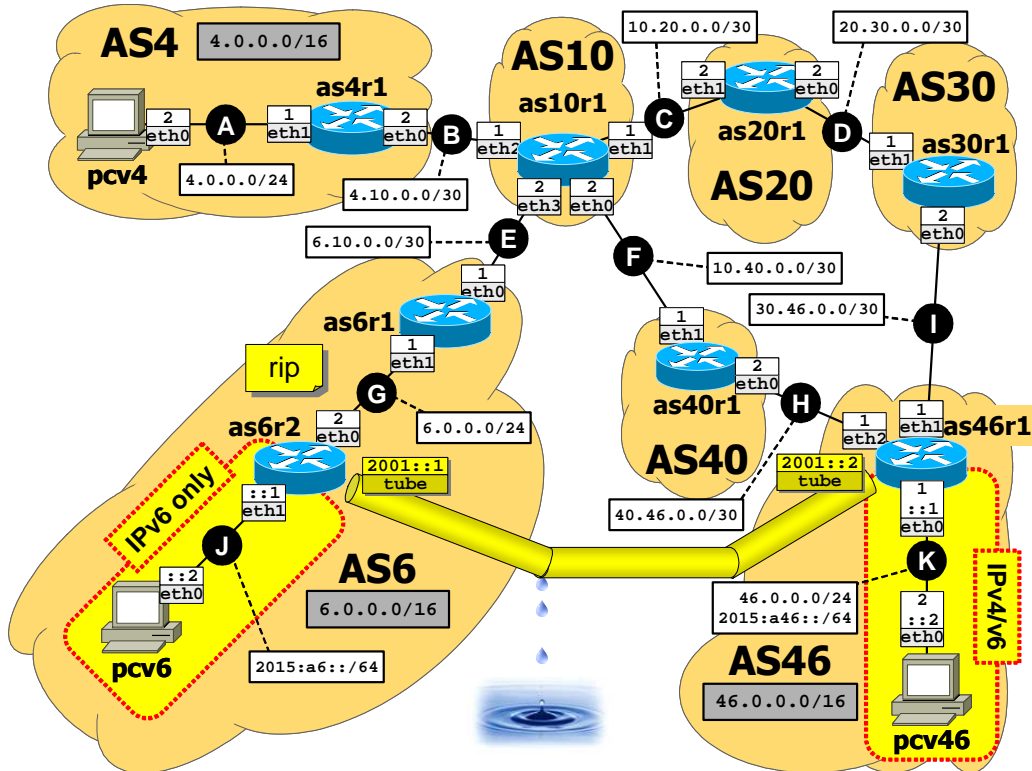




Available time: 90 minutes.



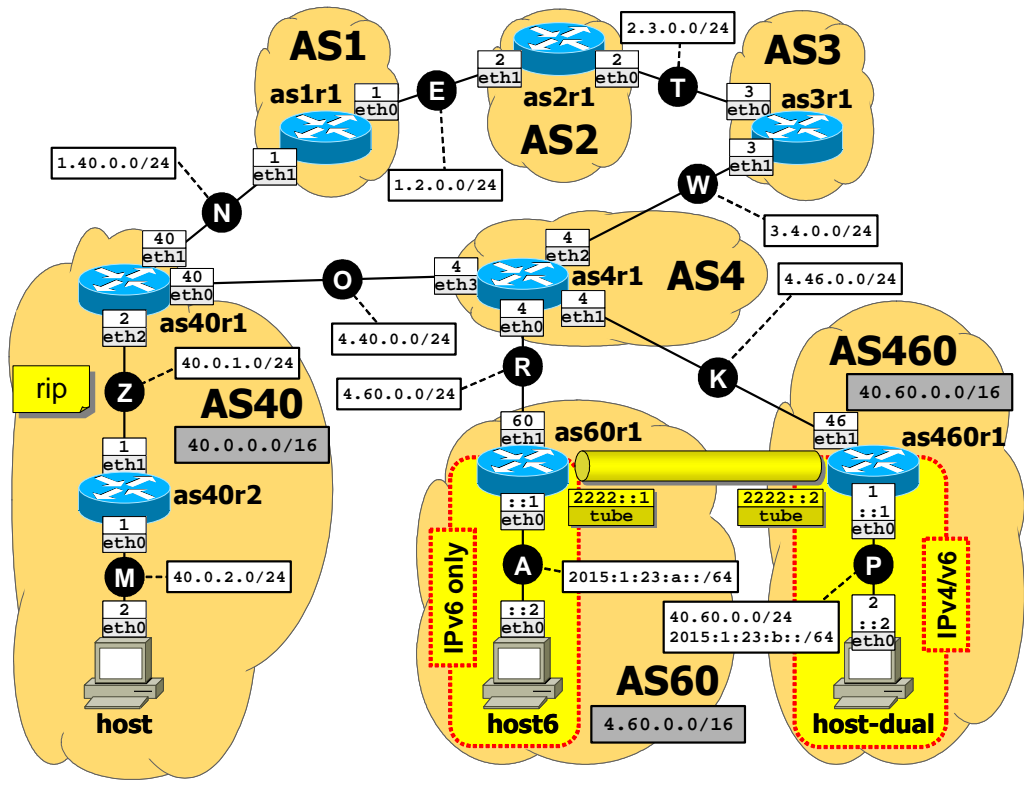
Using Netkit, implement the network depicted in the figure and described below (you can use the following items as a checklist).

- IPv4 routing within **AS6** is implemented by using RIP.
- No routers announce **0.0.0.0/0** or IPv6 subnets.
- AS4**, **AS6**, and **AS46** announce the subnets in gray.
- AS10**, **AS20**, **AS30**, and **AS40** are transit ASes, and as such they do not announce any owned subnets.
- All peering LANs are announced.
- AS46** is a customer of **AS30** and **AS40**, and as such it must forbid transit traffic between **AS30** and **AS40**.
- AS46** prefers passing via **AS30** to reach **AS6**.
- Enable IPv6 forwarding on all IPv6 routers.
- IPv6 routing is implemented by static routes.
- An IPv6-in-IPv4 tunnel is established between interface **eth0** of **as6r2** and interface **eth1** of **as46r1** (the “tube” in the figure).

BGP	SETTING UP AN IPV6 ADDRESS
	<code>ifconfig INTERFACE up</code> <code>ifconfig INTERFACE add IPV6ADDR/NETMASK</code>
IPv6	ENABLING IPV6 FORWARDING
	<code>echo 1 >/proc/sys/net/ipv6/conf/all/forwarding</code>
	ADDING A STATIC IPV6 ROUTE
	<code>route -A inet6 add IPV6NET[/NETMASK] [gw IPV6ADDR] [dev INTERFACE]</code>
	CREATING AN IPV6-IN-IPV4 TUNNEL
	<code>ip tunnel add TUNNELNAME mode sit remote REMOTEIPV4 local LOCALIPV4 ttl 10</code>
	<code>ifconfig TUNNELNAME up</code> <code>ifconfig TUNNELNAME add LOCALIPV6[/NETMASK]</code> <code>route -A inet6 add default dev TUNNELNAME</code>
	IPV6 TRACEROUTE
	<code>traceroute6 -N 1 IPV6ADDR</code>

- Goals:**
- Every IPv4 address must be reachable from any IPv4-enabled network nodes. Same for IPv6.
 - Packet routing, especially among **AS46**, **AS4**, and **AS6**, must be consistent with BGP policies. **AS46** must forbid transit traffic.

Available time: 90 minutes.



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

- IPv4 routing within **AS40** is implemented by using RIP.
- No routers announce **0.0.0.0/0** or IPv6 subnets.
- AS40**, **AS60**, and **AS460** announce the subnets in gray.
- AS1**, **AS2**, **AS3**, and **AS4** are transit ASes, and as such they do not announce any owned subnets.
- All peering LANs are announced.
- AS40** is a customer of **AS1** and **AS4**, and as such it must forbid transit traffic between **AS1** and **AS4**.
- To reach **AS40**, **AS4** prefers passing via **AS3** and **AS3** prefers passing via **AS2**.
- Enable IPv6 forwarding on all IPv6 routers.
- IPv6 routing is implemented by static routes.
- An IPv6-in-IPv4 tunnel is established between interface **eth1** of **as60r1** and interface **eth1** of **as460r1** (the “tube” in the figure).

```

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

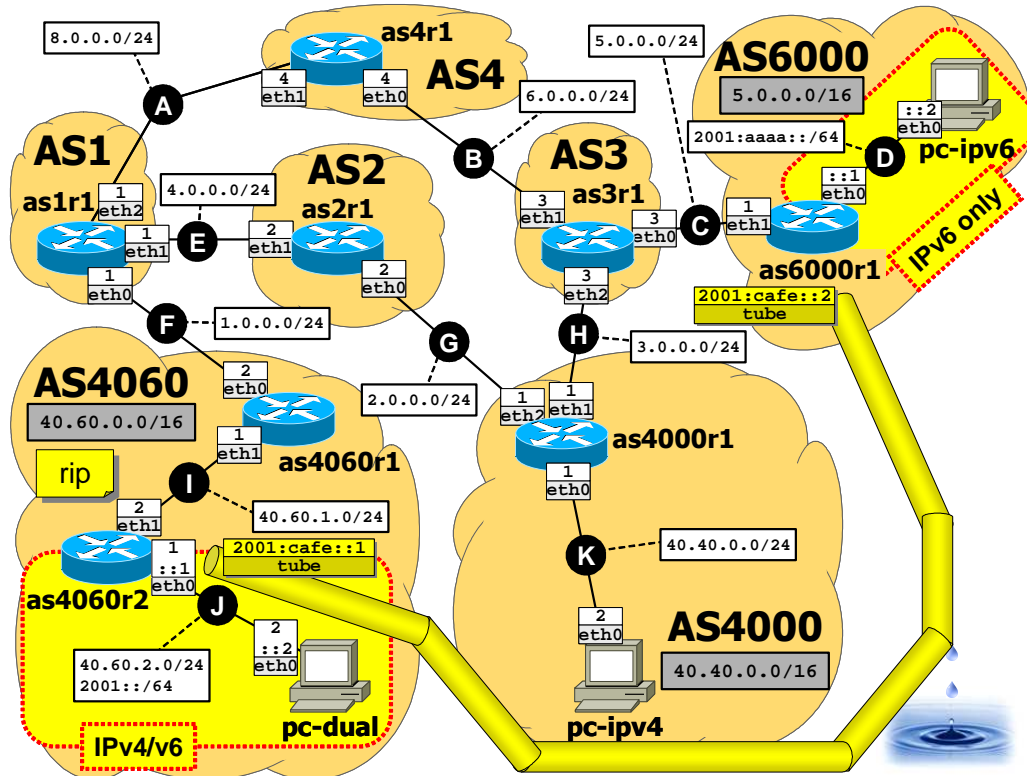
IPV6 TRACEROUTE
traceroute6 -N 1 IPV6ADDR
    
```

Goals:

- Every IPv4 address must be reachable from any IPv4-enabled network nodes. Same for IPv6.
- Packet routing, especially among **AS460**, **AS60**, and **AS40**, must be consistent with BGP policies. **AS40** must forbid transit traffic.



Available time: 90 minutes.



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

- IPv4 routing within **AS4060** is implemented by using RIP.
- No routers announce **0.0.0.0/0** or IPv6 subnets.
- AS4060**, **AS4000**, and **AS6000** announce the subnets in gray.
- AS1**, **AS2**, **AS3**, and **AS4** are transit ASes, and as such they do not announce any owned subnets.
- All peering LANs are announced.
- AS4000** is a customer of **AS2** and **AS3**, and as such it must forbid transit traffic between **AS2** and **AS3**.
- AS1** prefers passing via **AS4** to reach **AS4000**.
- Enable IPv6 forwarding on all IPv6 routers.
- IPv6 routing is implemented by static routes.
- An IPv6-in-IPv4 tunnel is established between interface **eth1** of **as4060r2** and interface **eth1** of **as6000r1** (the “tube” in the figure).

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

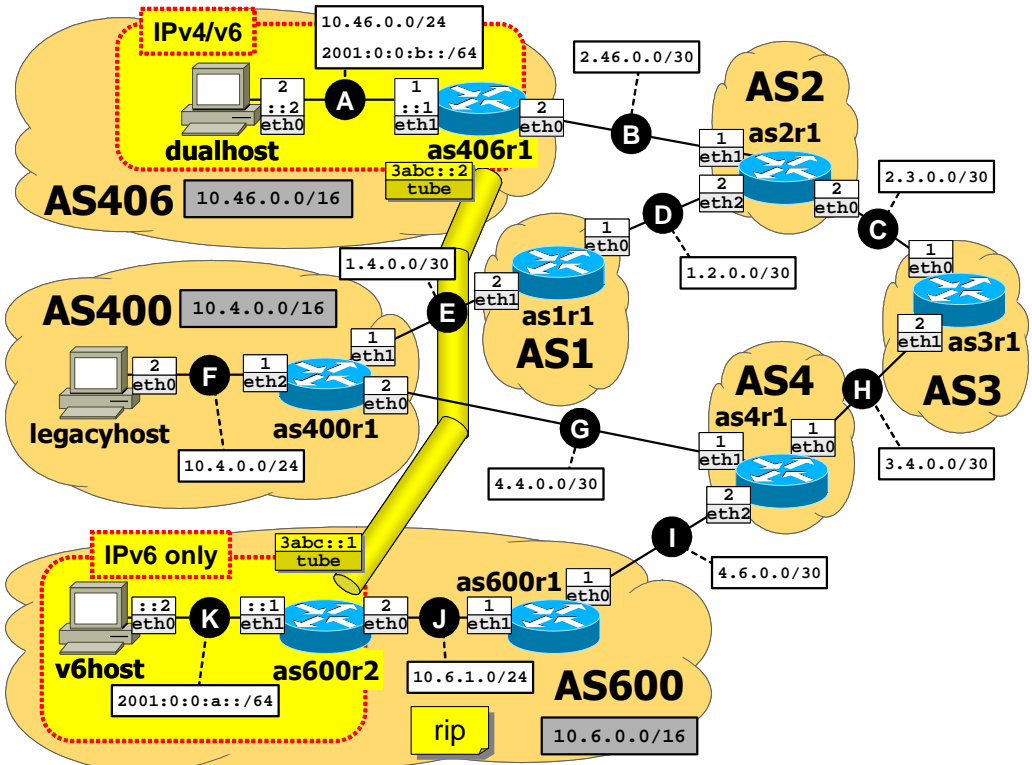
IPV6 TRACEROUTE

```
traceroute6 -N 1 IPV6ADDR
```

Goals:

- Every IPv4 address must be reachable from any IPv4-enabled network nodes. Same for IPv6.
- Packet routing, especially among **AS4060**, **AS4000** ed **AS6000**, must be consistent with BGP policies. **AS4000** must forbid transit traffic.

Available time: 90 minutes.



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

- IPv4 routing within **AS600** is implemented by using RIP.
- No routers announce **0.0.0.0/0** or IPv6 subnets.
- AS406**, **AS400**, and **AS600** announce the subnets in gray.
- AS1**, **AS2**, **AS3**, and **AS4** are transit ASes, and as such they do not announce any owned subnets.
- All peering LANs are announced.
- AS400** is a customer of **AS1** and **AS4**, and as such it must forbid transit traffic between **AS1** and **AS4**.
- AS2** prefers passing via **AS3** to reach **AS400**.
- Enable IPv6 forwarding on all IPv6 routers.
- IPv6 routing is implemented by static routes.
- An IPv6-in-IPv4 tunnel is established between interface **eth0** of **as600r2** and interface **eth0** of **as406r1** (the “tube” in the figure).

```

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

IPV6 TRACEROUTE
traceroute6 -N 1 IPV6ADDR
    
```

Goals:

- Every IPv4 address must be reachable from any IPv4-enabled network nodes. Same for IPv6.
- Packet routing, especially among **AS406**, **AS400**, and **AS600**, must be consistent with BGP policies. **AS400** must forbid transit traffic.