LAN Switching (LAB Practice no. 4, 2022)

- LAB sessions to weeks 17 and 18 in 2022

- Include a report in your LabBook about the most significant results obtained **in this LAB practice**
  - Use the board pictures and screen dumps uploaded to the agora
- Solve the exercises included in this script in pages 15 and 19

What frames are accepted by a NIC

- A host NIC accepts frames having the following characteristics:
  - With DEST MAC equal to the NIC's MAC
  - With DEST MAC equal to the the Broadcast address (0xffffffffffff)
  - Frames sent to *any* Multicast addresses
- If we wish to have the NIC accept all of the frames delivered to the switch, we have to set the PROMISCUOUS MODE
  - With the ifconfig command
  - Programming with libpcap
  - Programming with Netlink sockets

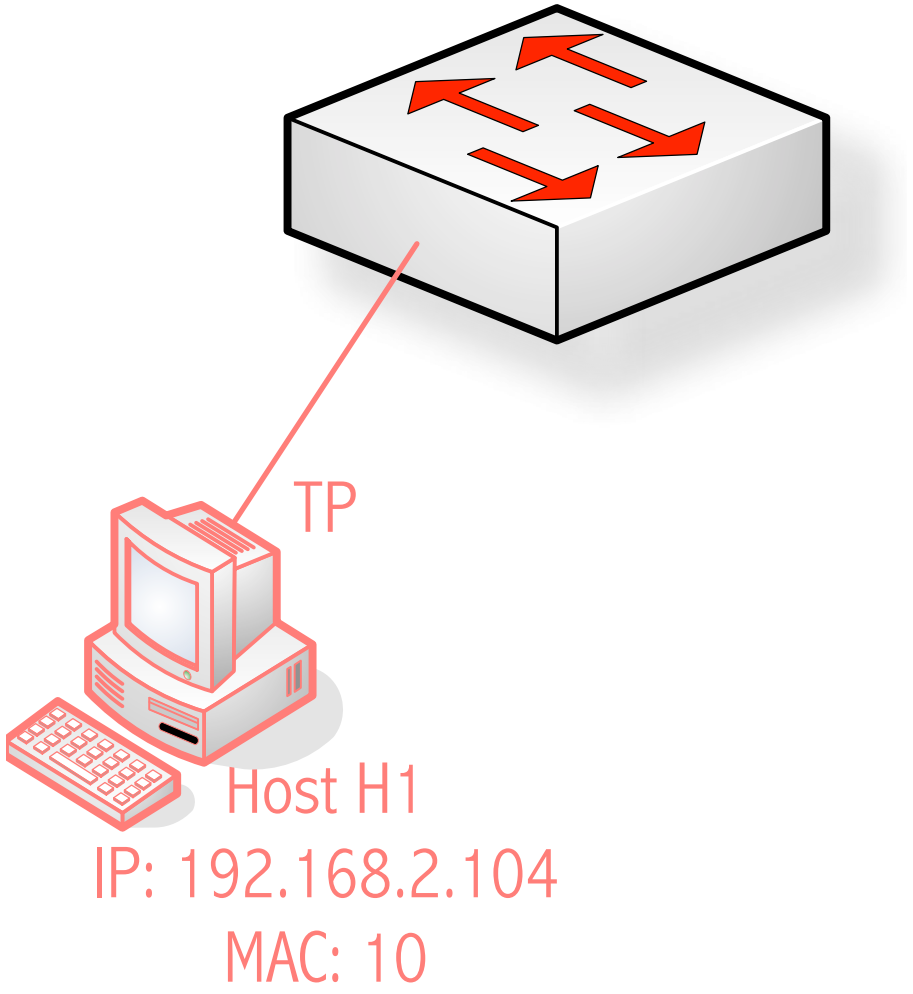

Communicating two max Ethernets

What network device can communicate E1 with E2 without violating the max distance limitation of Ethernet???

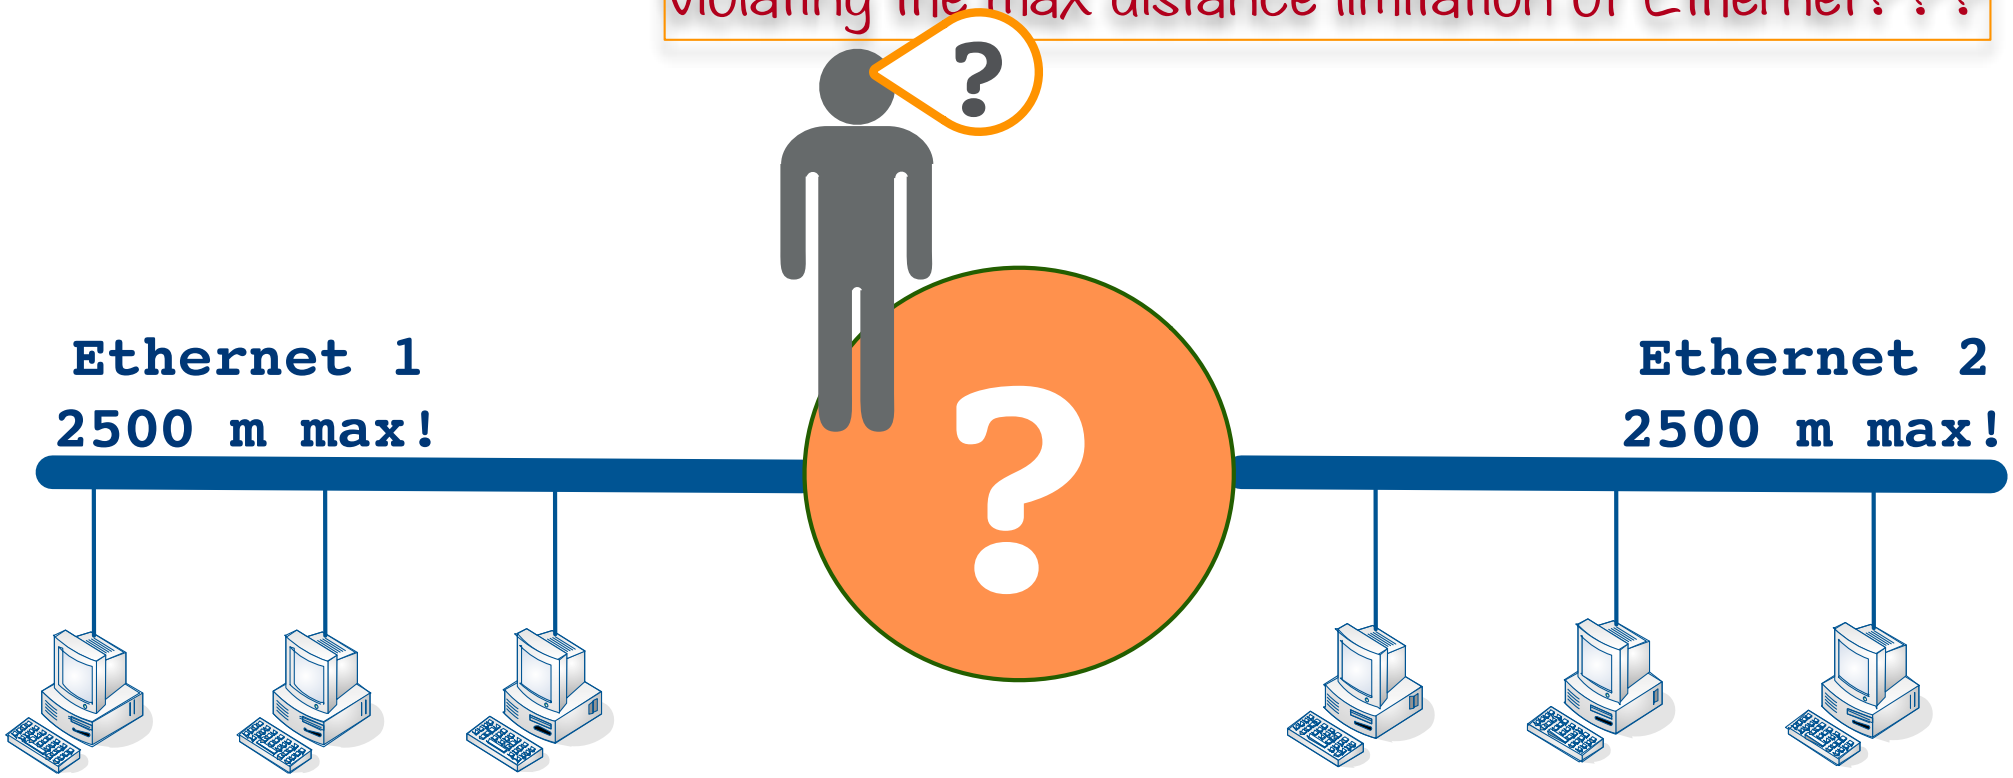

Connecting two Ethernets: Bridge

- A) Repeater in between them?
  - It might exceed the physical limitation of the Ethernet
  - 4 repeaters, < 2500 m
- B) Hubs regenerate **electrical** signals
  - Hubs are layer-1 devices (OSI)
- C) Bridge? New network equipment that forwards **frames** between two LANs
  - Bridges/switches are layer-2 devices (OSI)

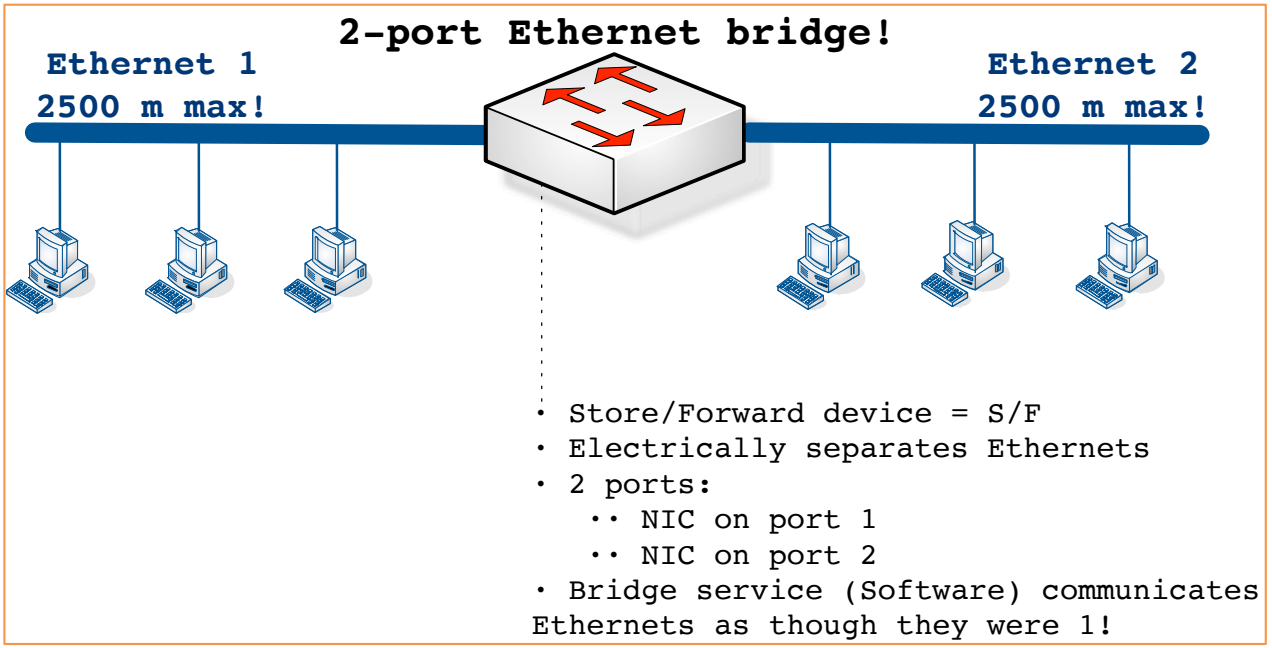

A Switched LAN

- **Switches learn** the location of each host connected to their network
  - Record the source MAC of each received frame
  - Into the MAC Table
- **Forward** each frame by looking up its **Destination MAC** into the **MAC Table**
  - Make their best effort to deliver the frame to its destination host
  - If Dest MAC is not contained in the MAC Table
    - **FLOOD** the frame

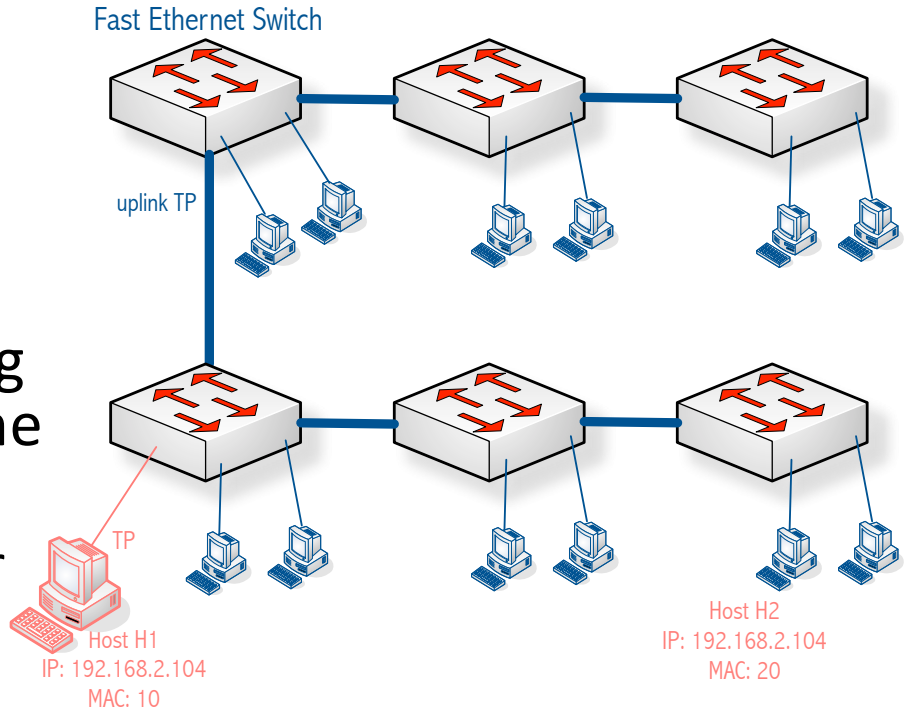

Open several tcpdump sessions

- Access paloalto with ssh at TCP port 50500
  - Then, hop at 3 or 4 hosts connected to the selected Cisco switch
  - administrator/19xxdpq16
- At each hop host, execute tcpdump and filter your personal ethertype

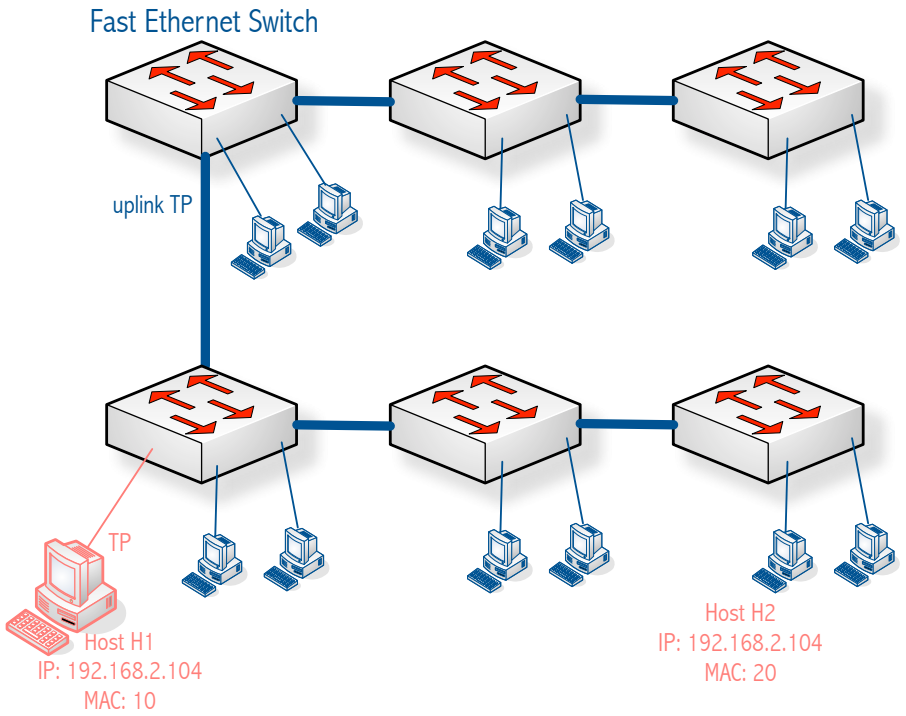

Exercise 1: Sending to the broadcast MAC address

- H_1 sends a frame to the broadcast MAC
  - (0xffffffffffff)
- Switches flood the frame
  - ***Flooding*** *as as consequence of the switch receiving a frame which dest address is the broadcast MAC address*

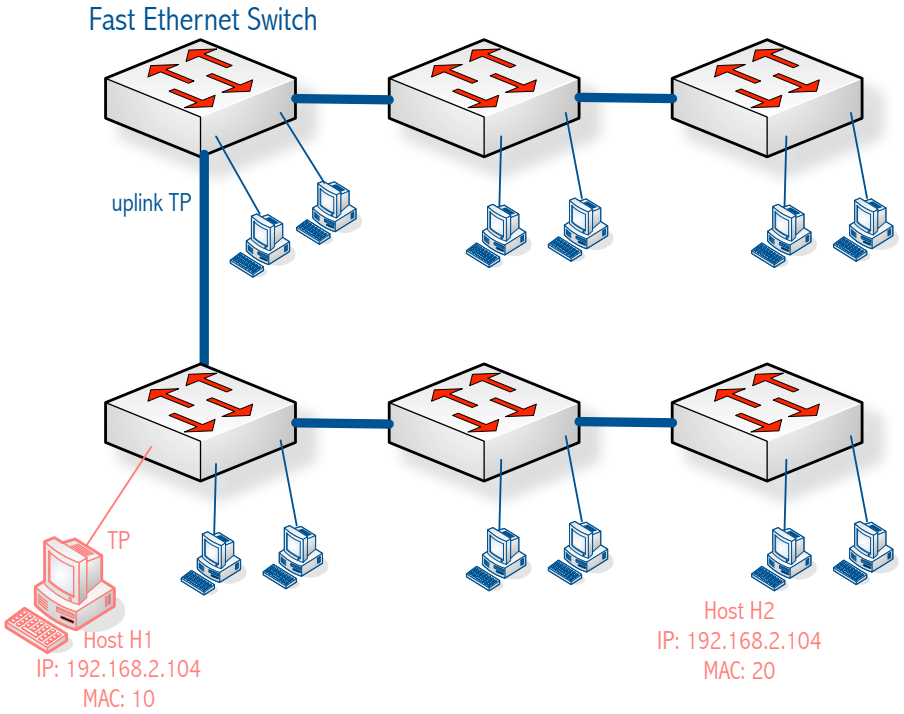

Exercise 2: Flooding

- ***Flooding*** *as as consequence of the switch not knowing some MAC address*
- Use the PF_PACKET send program that we made in practice 3.1
- H_1 sends a frame to a MAC address that hasn't been learned yet by any switch
  - Send your own assigned Ethertype
- No switch will have M recorded into the MAC table
- All switches flood the frame

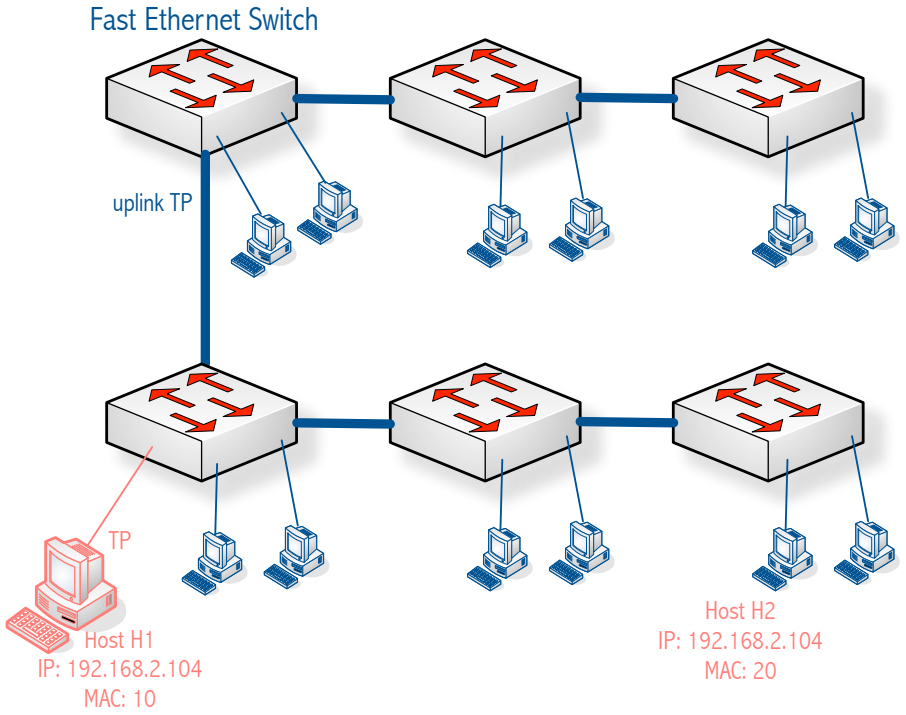

Exercise 3: Unicast traffic

- H_1 sends a frame to the MAC of H_2
  - Send your own assigned Ethertype
- Observe the frame with tcpdump as it is being transmitted in H_1 as well as in H_2

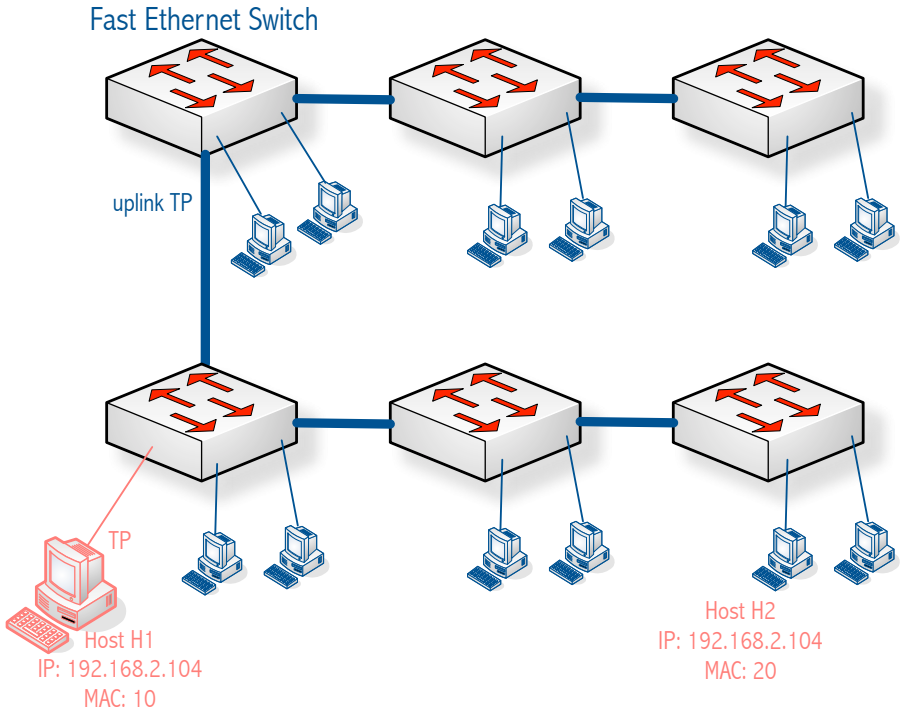

Exercise 4: Check mac address table of Switch from paloalto.unileon.es

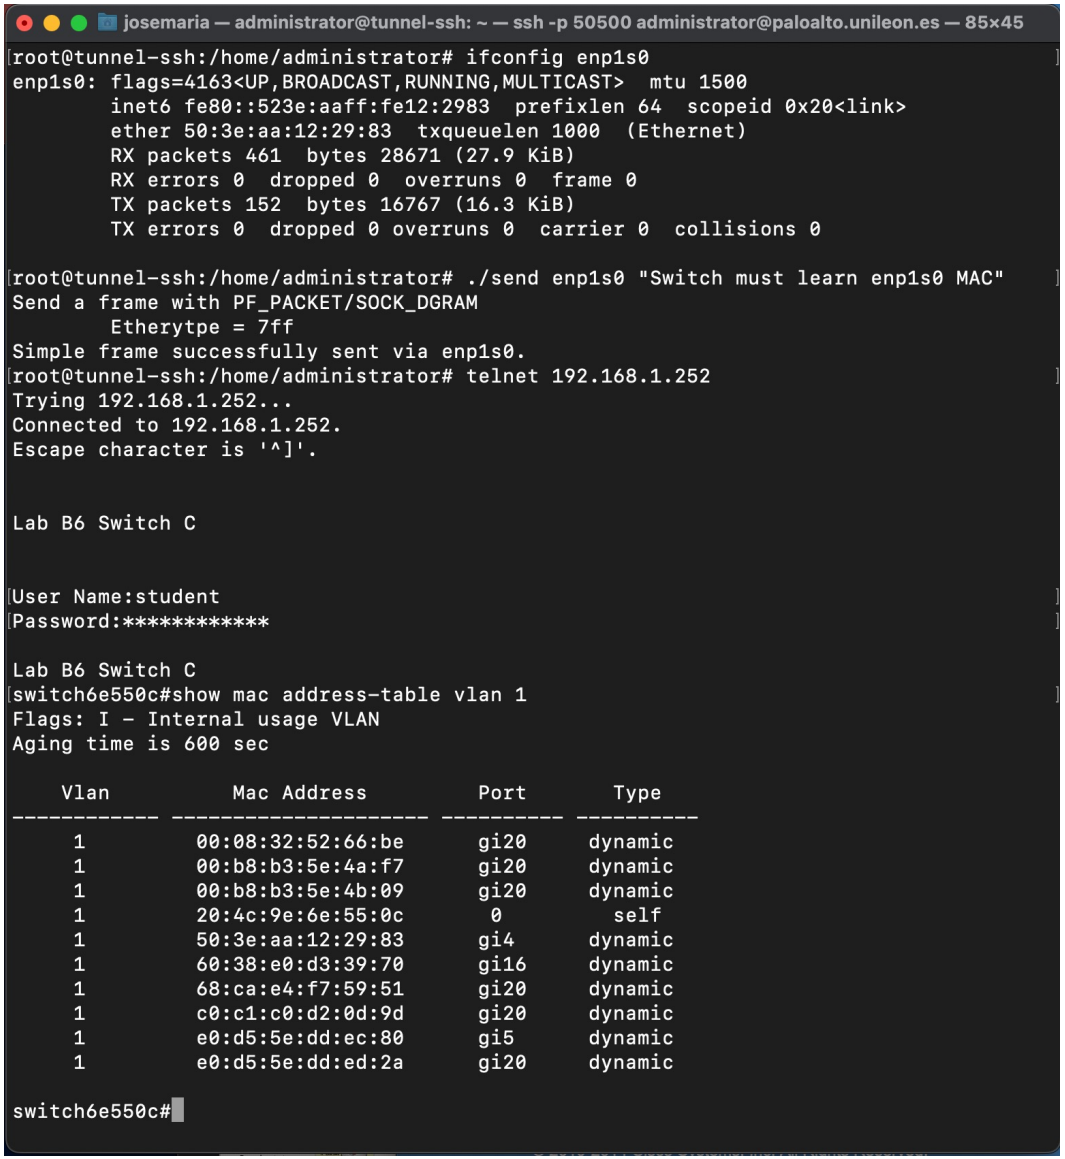

Broadcast storms

- If a loop is exists in a switched network frames may proliferate:
  - Frames which dest mac is the broadcast address
  - Frames whose dest mac is unknown
  - In both of the above cases, the switch will flood the frame
- **Flooding** will cause a *Broadcast Storm* if loops exist in the Extended LAN
- Broadcast storms can consume all the aggregated throughput in a switched network
- **SOLUTION: Spanning Tree Protocol (STP)**

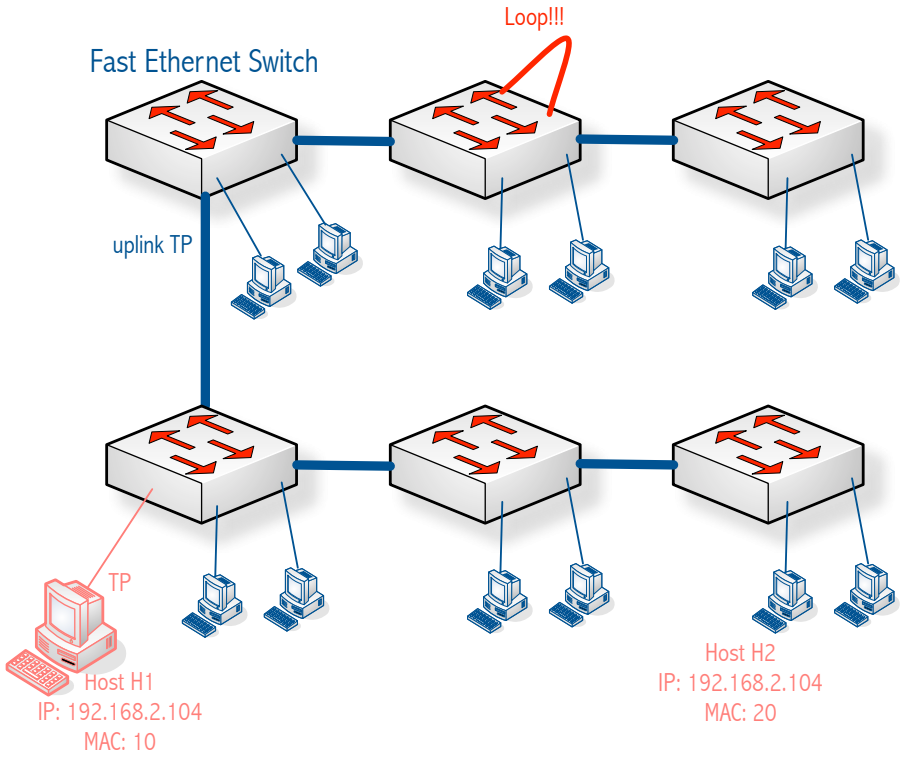

Simple switches and broadcast storms

- Simple, off-the-shelf TP-Link switches
- No Spanning Tree
- Loop => Broadcast storm

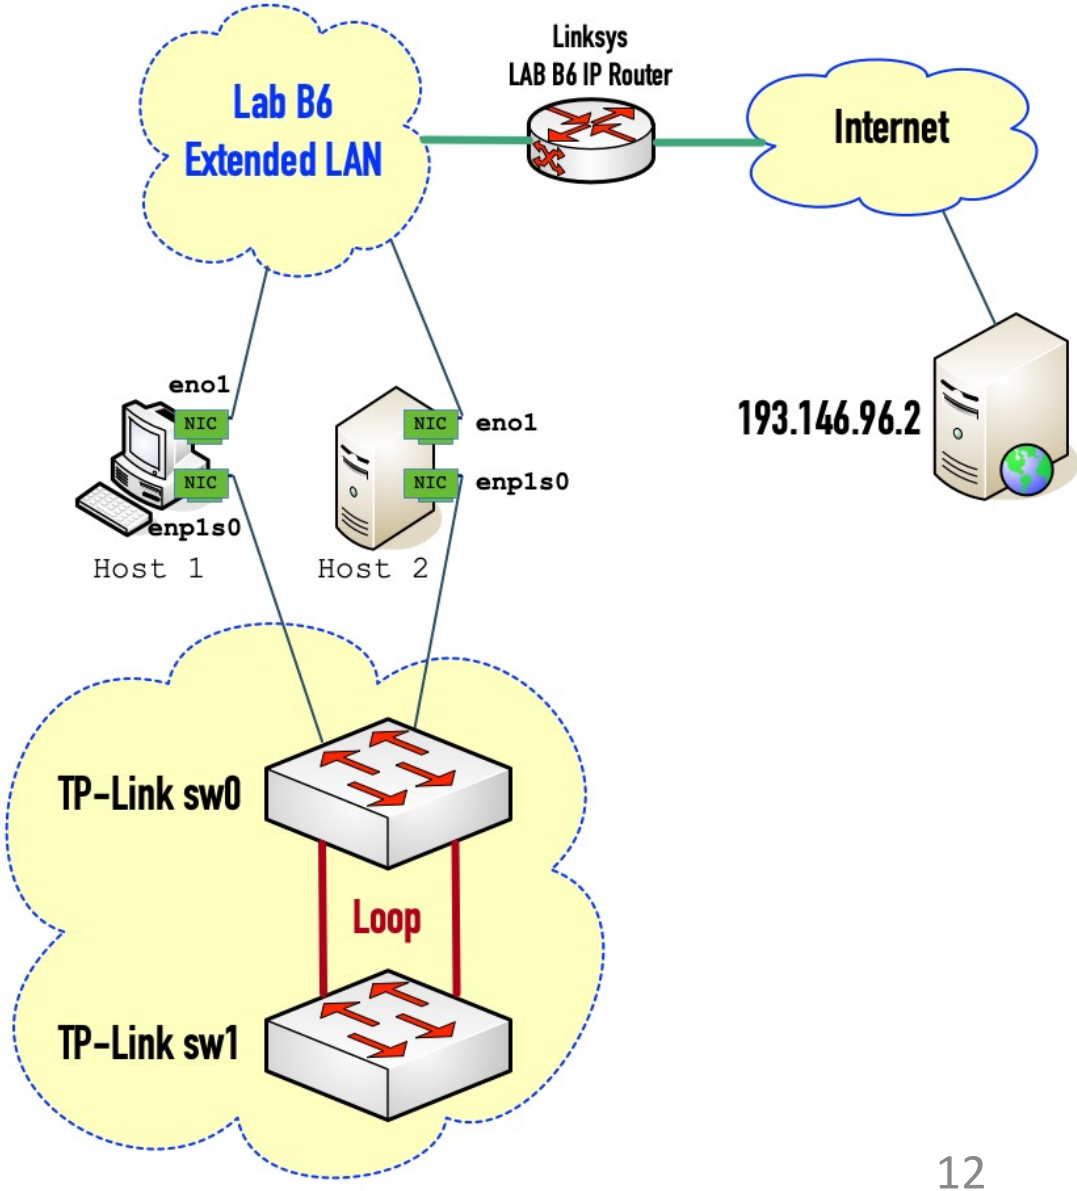

No loop: no broadcast storm

- Configure enp1s0 interfaces with ifconfig:
  - Host 1: 192.168.0.100
  - Host 2: 192.168.0.200
  - Netmask: 255.255.255.0
- Host1:
  - ping 192.168.0.200
  - Send single frame to broadcast
- Host 2:
  - Observe pings every second
  - Observe single broadcast frame

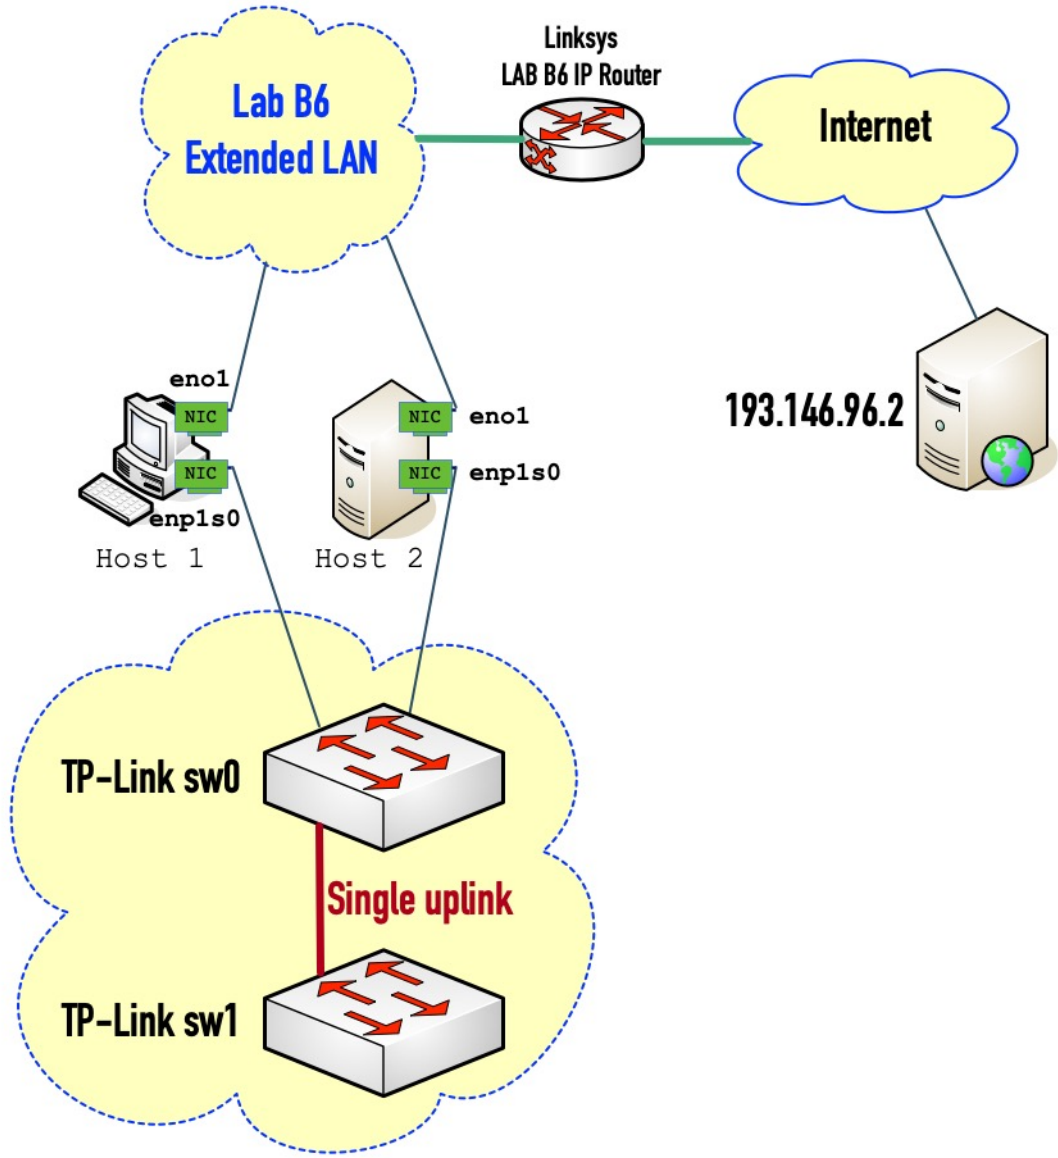

Loop: broadcast storm

- Close the loop
- Pings should be still normally received at host 2

- Send **single** frame to braoadcast
  - That frame should proliferate
  - Broadcast storm is created
  - A continual stream of proliferated copies of sinlge frame should be observed now at host 2
  - Where are the pings now?

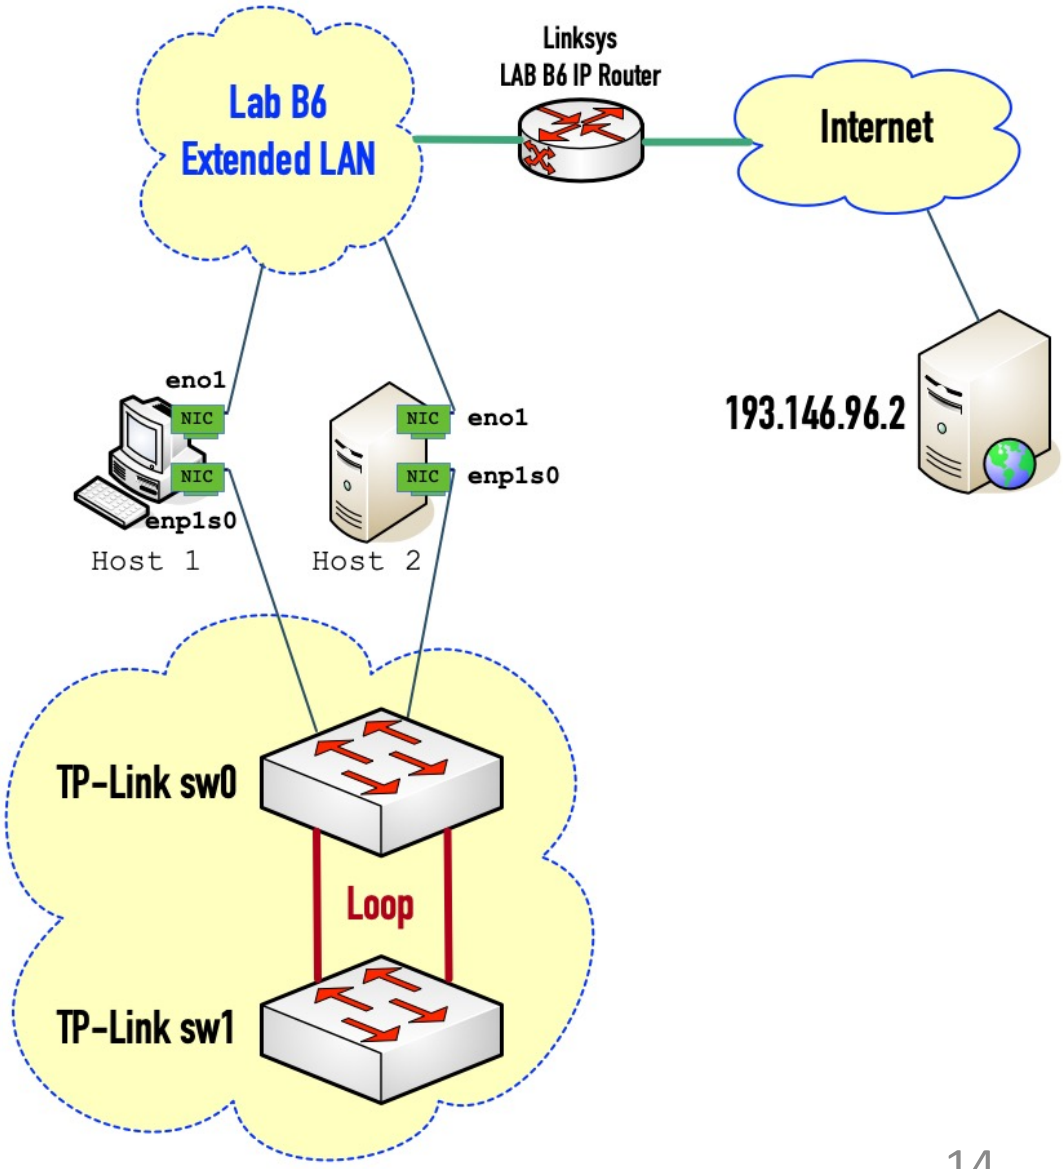

No loop, storm ends

- Open the loop
- Pings should be normally received at host 2, again

- Send **single** frame to braoadcast
  - That frame should not proliferate anymore
  - Broadcast storm must fade away now
  - Only one copy of the single sent frame should be received by host 2

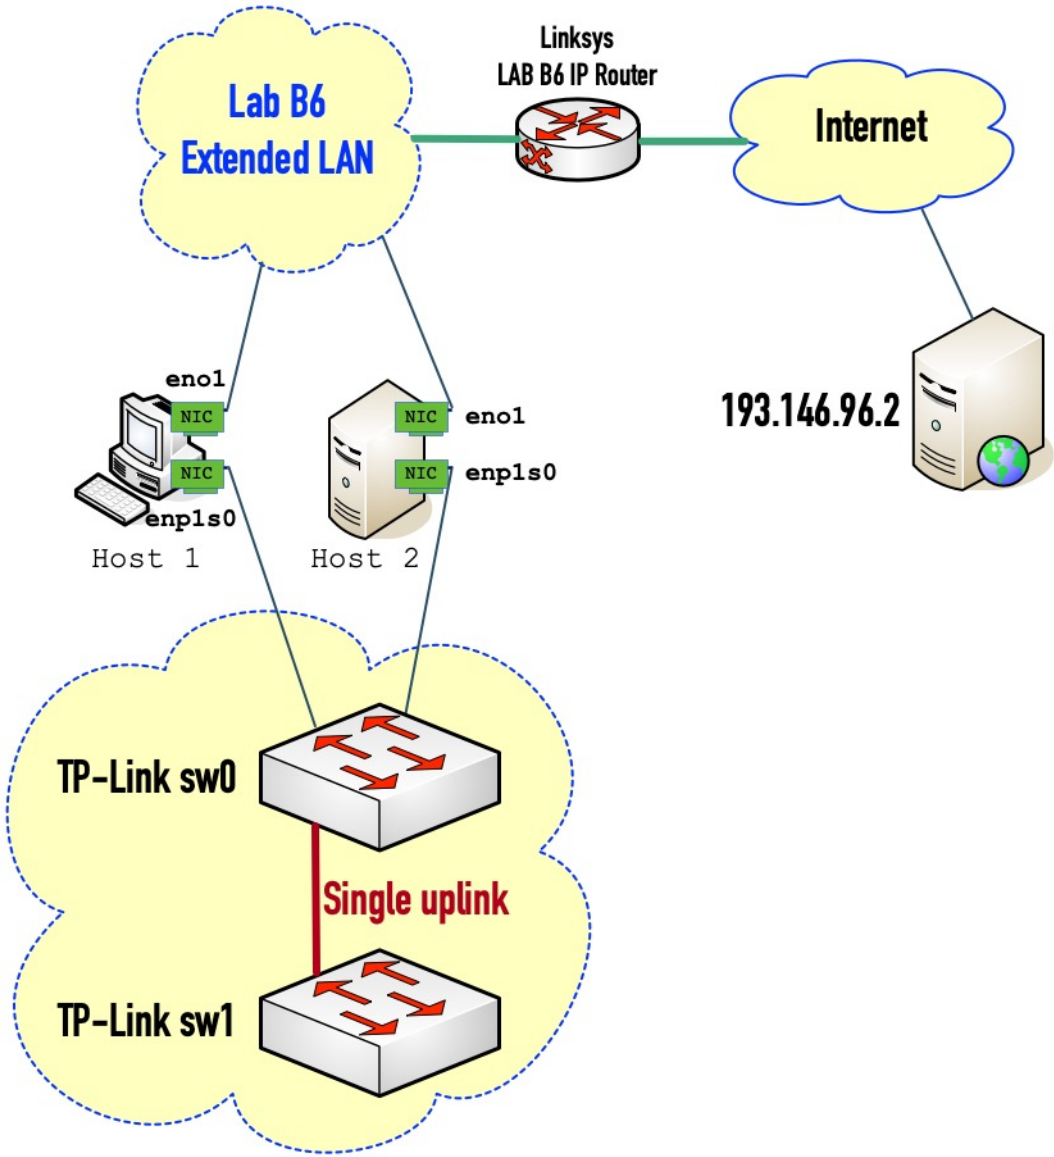

Exercises for Lab Book

- 5. Describe the experiment that we've have done on slides 10-14.
  - Use a terminology appropriate for Switched LANs and STP

- 6. Can the *broadcast storm propagate on to Lab B6 extended LAN?*
  - What conditions should be met such that the *broadcast storm initiated in the TP-Link network effectively propagates on to Lab B6 extended LAN?*

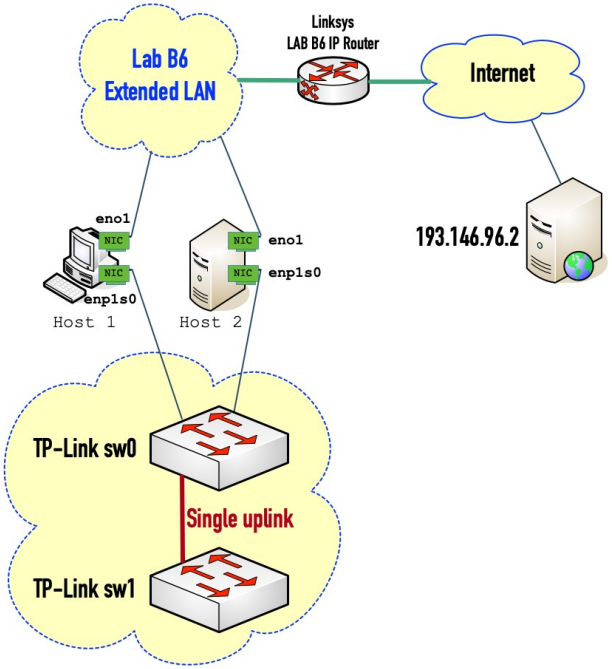

STP disables ports for breaking loops

- STP at Switch B0 has disabled port 2

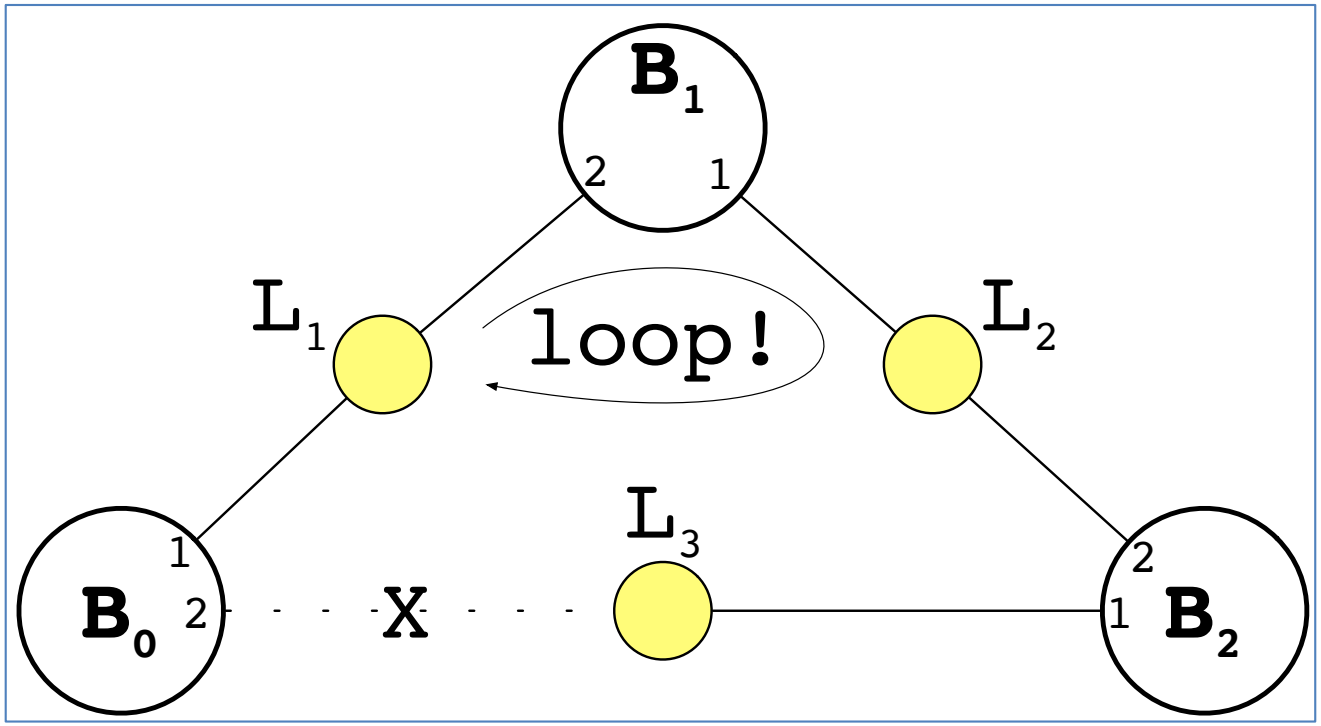

Loops in Cisco switches running STP

- **Observe STP traffic between the Cisco switches in the lab**
- **Tcpdump**

  `# tcpdump -i eno1 -etn -XX -vvv stp`
- **IEEE 802.1D BPDUs encap into IEEE 802.3 frames**
- **The Multicast MAC** address reserved for *all the bridges connected to this extended LAN:*

  0x0180C2000000

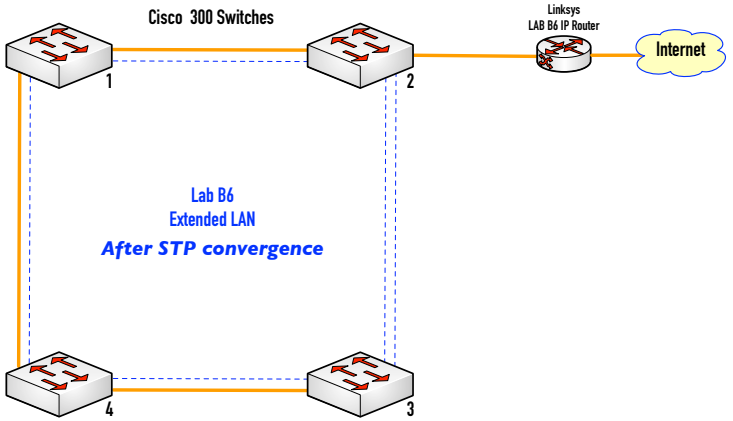

Example

```
20:4c:9e:6e:55:11 > 01:80:c2:00:00:00, 802.3, length 39: LLC, dsap STP (0x42)
Individual, ssap STP (0x42) Command, ctrl 0x03: STP 802.1w, Rapid STP, Flags
[Learn, Forward, Agreement], bridge-id 8000.20:4c:9e:6e:55:0c.8035, length 43
      message-age 2.00s, max-age 20.00s, hello-time 2.00s, forwarding-delay 15.00s
      root-id 8000.00:08:32:52:66:be, root-pathcost 40000, port-role Designated
      0x0000:  0180 c200 0000 204c 9e6e 5511 0027 4242  .......L.nU..'BB
      0x0010:  0300 0002 027c 8000 0008 3252 66be 0000  .....|....2Rf...
      0x0020:  9c40 8000 204c 9e6e 550c 8035 0200 1400  .@...L.nU..5....
      0x0030:  0200 0f00 0000 0000 0000 0000            ............
```

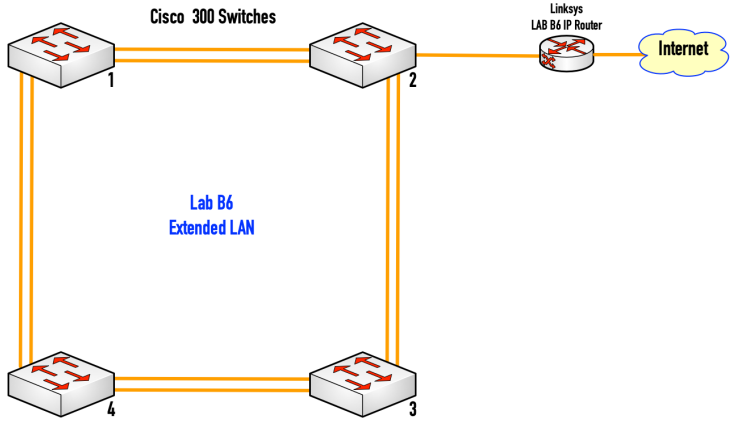

BPDUs are encapsulated into 802.3

802.1D frame encapsulated into 802.3

```
20:4c:9e:6e:55:11 > 01:80:c2:00:00:00, 802.3, length 39: LLC, dsap STP (0x42)
Individual, ssap STP (0x42) Command, ctrl 0x03: STP 802.1w, Rapid STP, Flags
[Learn, Forward, Agreement], bridge-id 8000.20:4c:9e:6e:55:0c.8035, length 43
      message-age 2.00s, max-age 20.00s, hello-time 2.00s, forwarding-delay 15.00s
      root-id 8000.00:08:32:52:66:be, root-pathcost 40000, port-role Designated
      0x0000:  0180 c200 0000 204c 9e6e 5511 0027 4242  .......L.nU..'BB
      0x0010:  0300 0002 027c 8000 0008 3252 66be 0000  .....|....2Rf...
      0x0020:  9c40 8000 204c 9e6e 550c 8035 0200 1400  .@...L.nU..5....
      0x0030:  0200 0f00 0000 0000 0000 0000            ............
```

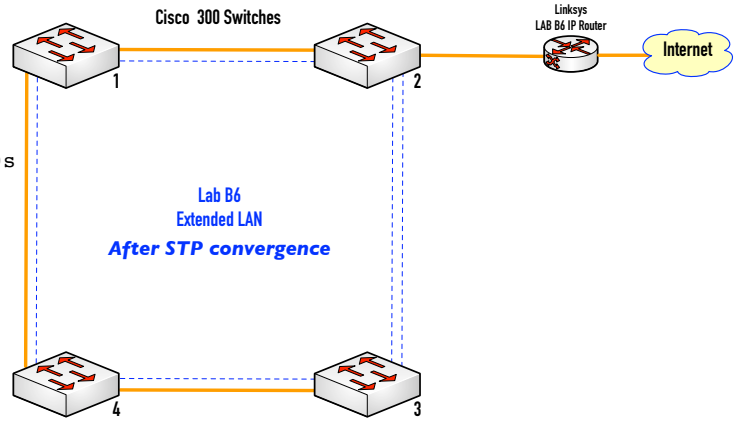

Ethernet *vs.* IEEE 802.3 Frames:

Dest MAC(48) + Src MAC(48) + **length(16)** + SCM(Variable length) + CRC(32)

```
if length <= 1500
      frame is IEEE 802.3
else
      frame is Ethernet
```

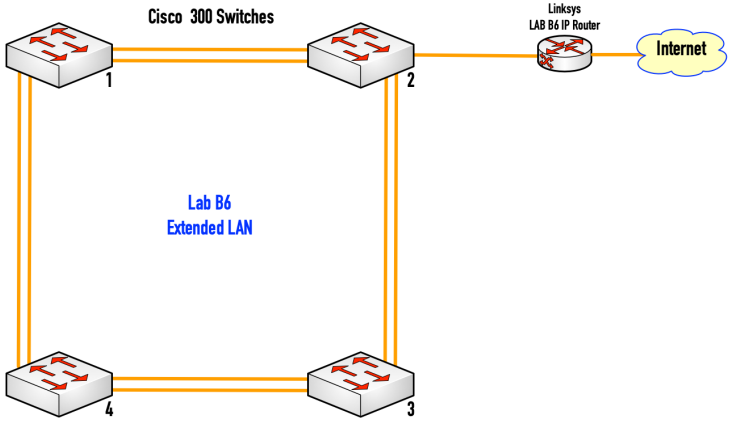

802.3 multiplexing key:

- DSAP(8) + SSAP(8) each one of them must contain the binary value 0x42 which represents the STP.

Exercises for Lab Book

7. Disconnect a few uplinks from the switch loop ascertaining that all switches are still connected

8.Observe the convergence of the STP by using tcpdump from your student account in paloalto.unileon.es

- Explain the most significant fields of the received BPDUs

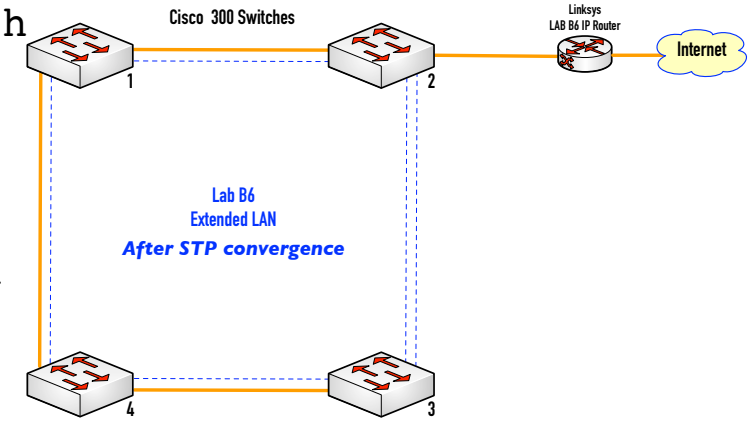

9. What's the DEST MAC in each BPDU? Is it a Multicast address?

- Check that the bits from the binary representation of the received Destination MAC do represent a Multicast MAC

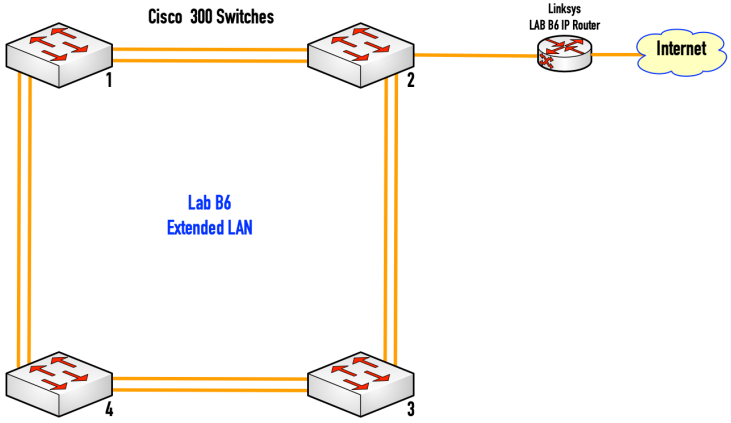

10. Why is the case that every PC in the lab is *accepting* STP traffic

11. Check the received frame.
    a. Is it IEEE 802.3 or is it Ethernet?
    b. Explain its multiplexing key

THE END