Routing loops

Routing loops

```
[rbartos@RBs-MBP ~]$ traceroute google.com
traceroute: Warning: google.com has multiple addresses; using 74.125.39.104
traceroute to google.com (74.125.39.104), 64 hops max, 52 byte packets
 1  172.20.10.1 (172.20.10.1)  1.624 ms  1.192 ms  1.102 ms
 2  * * *
 3  10.126.180.51 (10.126.180.51)  777.291 ms  221.734 ms  225.316 ms
 4  fx-in-f104.1e100.net (74.125.39.104)  299.681 ms  211.961 ms  179.118 ms
 5  10.126.180.158 (10.126.180.158)  181.927 ms  179.536 ms
    10.126.180.156 (10.126.180.156)  180.695 ms
 6  62.141.11.129 (62.141.11.129)  179.466 ms
    62.141.11.133 (62.141.11.133)  392.432 ms  284.436 ms
 7  62.141.11.226 (62.141.11.226)  239.925 ms  420.173 ms
    10.126.196.6 (10.126.196.6)  193.896 ms
 8  nixcz.net.google.com (91.210.16.211)  179.074 ms  179.463 ms  180.535 ms
 9  216.239.46.55 (216.239.46.55)  207.427 ms  198.271 ms  221.075 ms
10  209.85.254.118 (209.85.254.118)  198.511 ms  225.061 ms
    209.85.254.112 (209.85.254.112)  211.958 ms
11  209.85.249.162 (209.85.249.162)  206.668 ms
    209.85.254.126 (209.85.254.126)  214.440 ms
    209.85.249.162 (209.85.249.162)  208.257 ms
12  fx-in-f104.1e100.net (74.125.39.104)  210.846 ms  220.091 ms  225.169 ms
```

Oops...

No route to host

Photo courtesy of Chip McNaughton, UNH-CS

HISTORICAL NOTE — CLASSES OF IP ADDRESSES

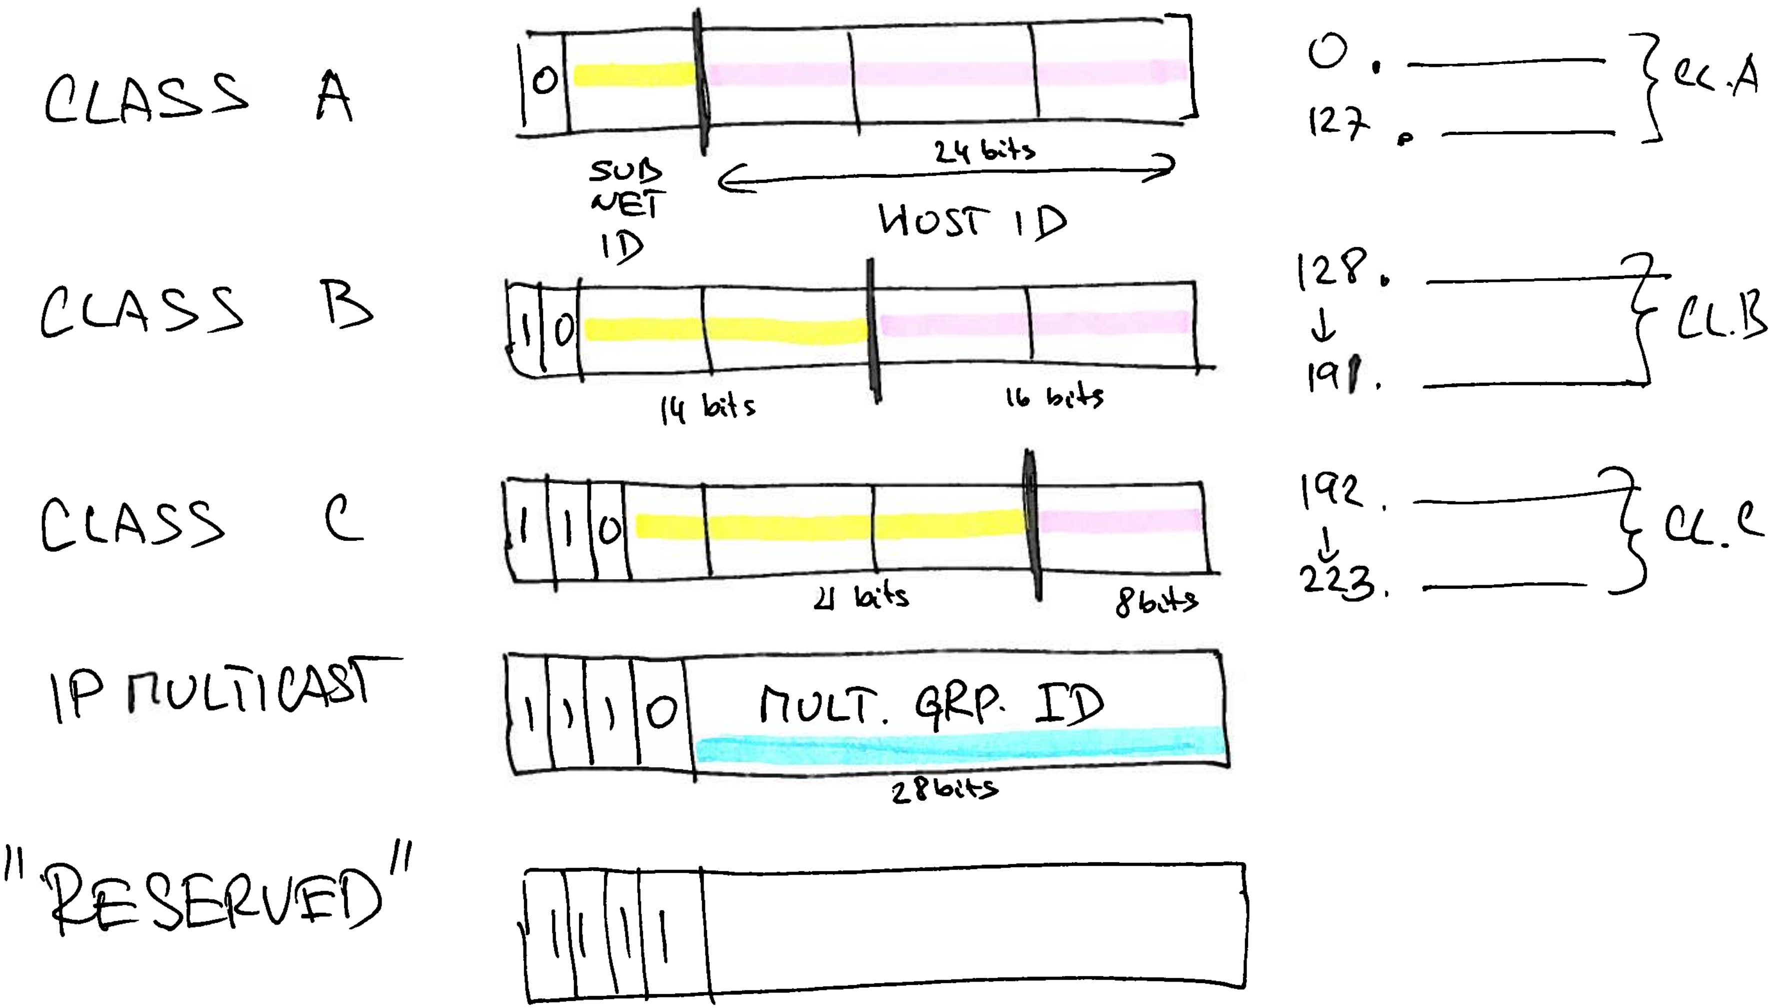

SPEC. IP ADDRESSES:

- 127.0.0.1 — LOCALHOST
- 192.168.0.0/16
- 172.16.0.0/12
- 10.0.0.0/8

} PRIVATE IP ADDR'S

- 255.255.255.255 — IP BROADCAST

CIDR Addresses

notation	addrs/block	# blocks	
n.n.n.n/32	1	4294967296	"host route"
n.n.n.x/31	2	2147483648	"p2p link"
n.n.n.x/30	4	1073741824	
n.n.n.x/29	8	536870912	
n.n.n.x/28	16	268435456	
n.n.n.x/27	32	134217728	
n.n.n.x/26	64	67108864	
n.n.n.x/25	128	33554432	
n.n.n.0/24	256	16777216	legacy "Class C"
n.n.x.0/23	512	8388608	
n.n.x.0/22	1024	4194304	
n.n.x.0/21	2048	2097152	
n.n.x.0/20	4096	1048576	
n.n.x.0/19	8192	524288	
n.n.x.0/18	16384	262144	
n.n.x.0/17	32768	131072	
n.n.0.0/16	65536	65536	legacy "Class B"
n.x.0.0/15	131072	32768	
n.x.0.0/14	262144	16384	
n.x.0.0/13	524288	8192	
n.x.0.0/12	1048576	4096	
n.x.0.0/11	2097152	2048	
n.x.0.0/10	4194304	1024	
n.x.0.0/9	8388608	512	
n.0.0.0/8	16777216	256	legacy "Class A"
x.0.0.0/7	33554432	128	
x.0.0.0/6	67108864	64	
x.0.0.0/5	134217728	32	
x.0.0.0/4	268435456	16	
x.0.0.0/3	536870912	8	
x.0.0.0/2	1073741824	4	
x.0.0.0/1	2147483648	2	
0.0.0.0/0	4294967296	1	"default route"

From **RFC 4632**:

n is an 8-bit decimal octet value.

x is a 1- to 7-bit value, based on the prefix length, shifted into the most significant bits of the octet and converted into decimal form; the least significant bits of the octet are zero.

An Alternative: Bridging

- **Motivation**
  - Routers are expensive and require configuration
- **Approach** - extend the reach of L2
- **Problem** - scalability (L2 is broadcast-based)
- **Solution** - limit the scope of broadcasting (bridging)

Historical evolution

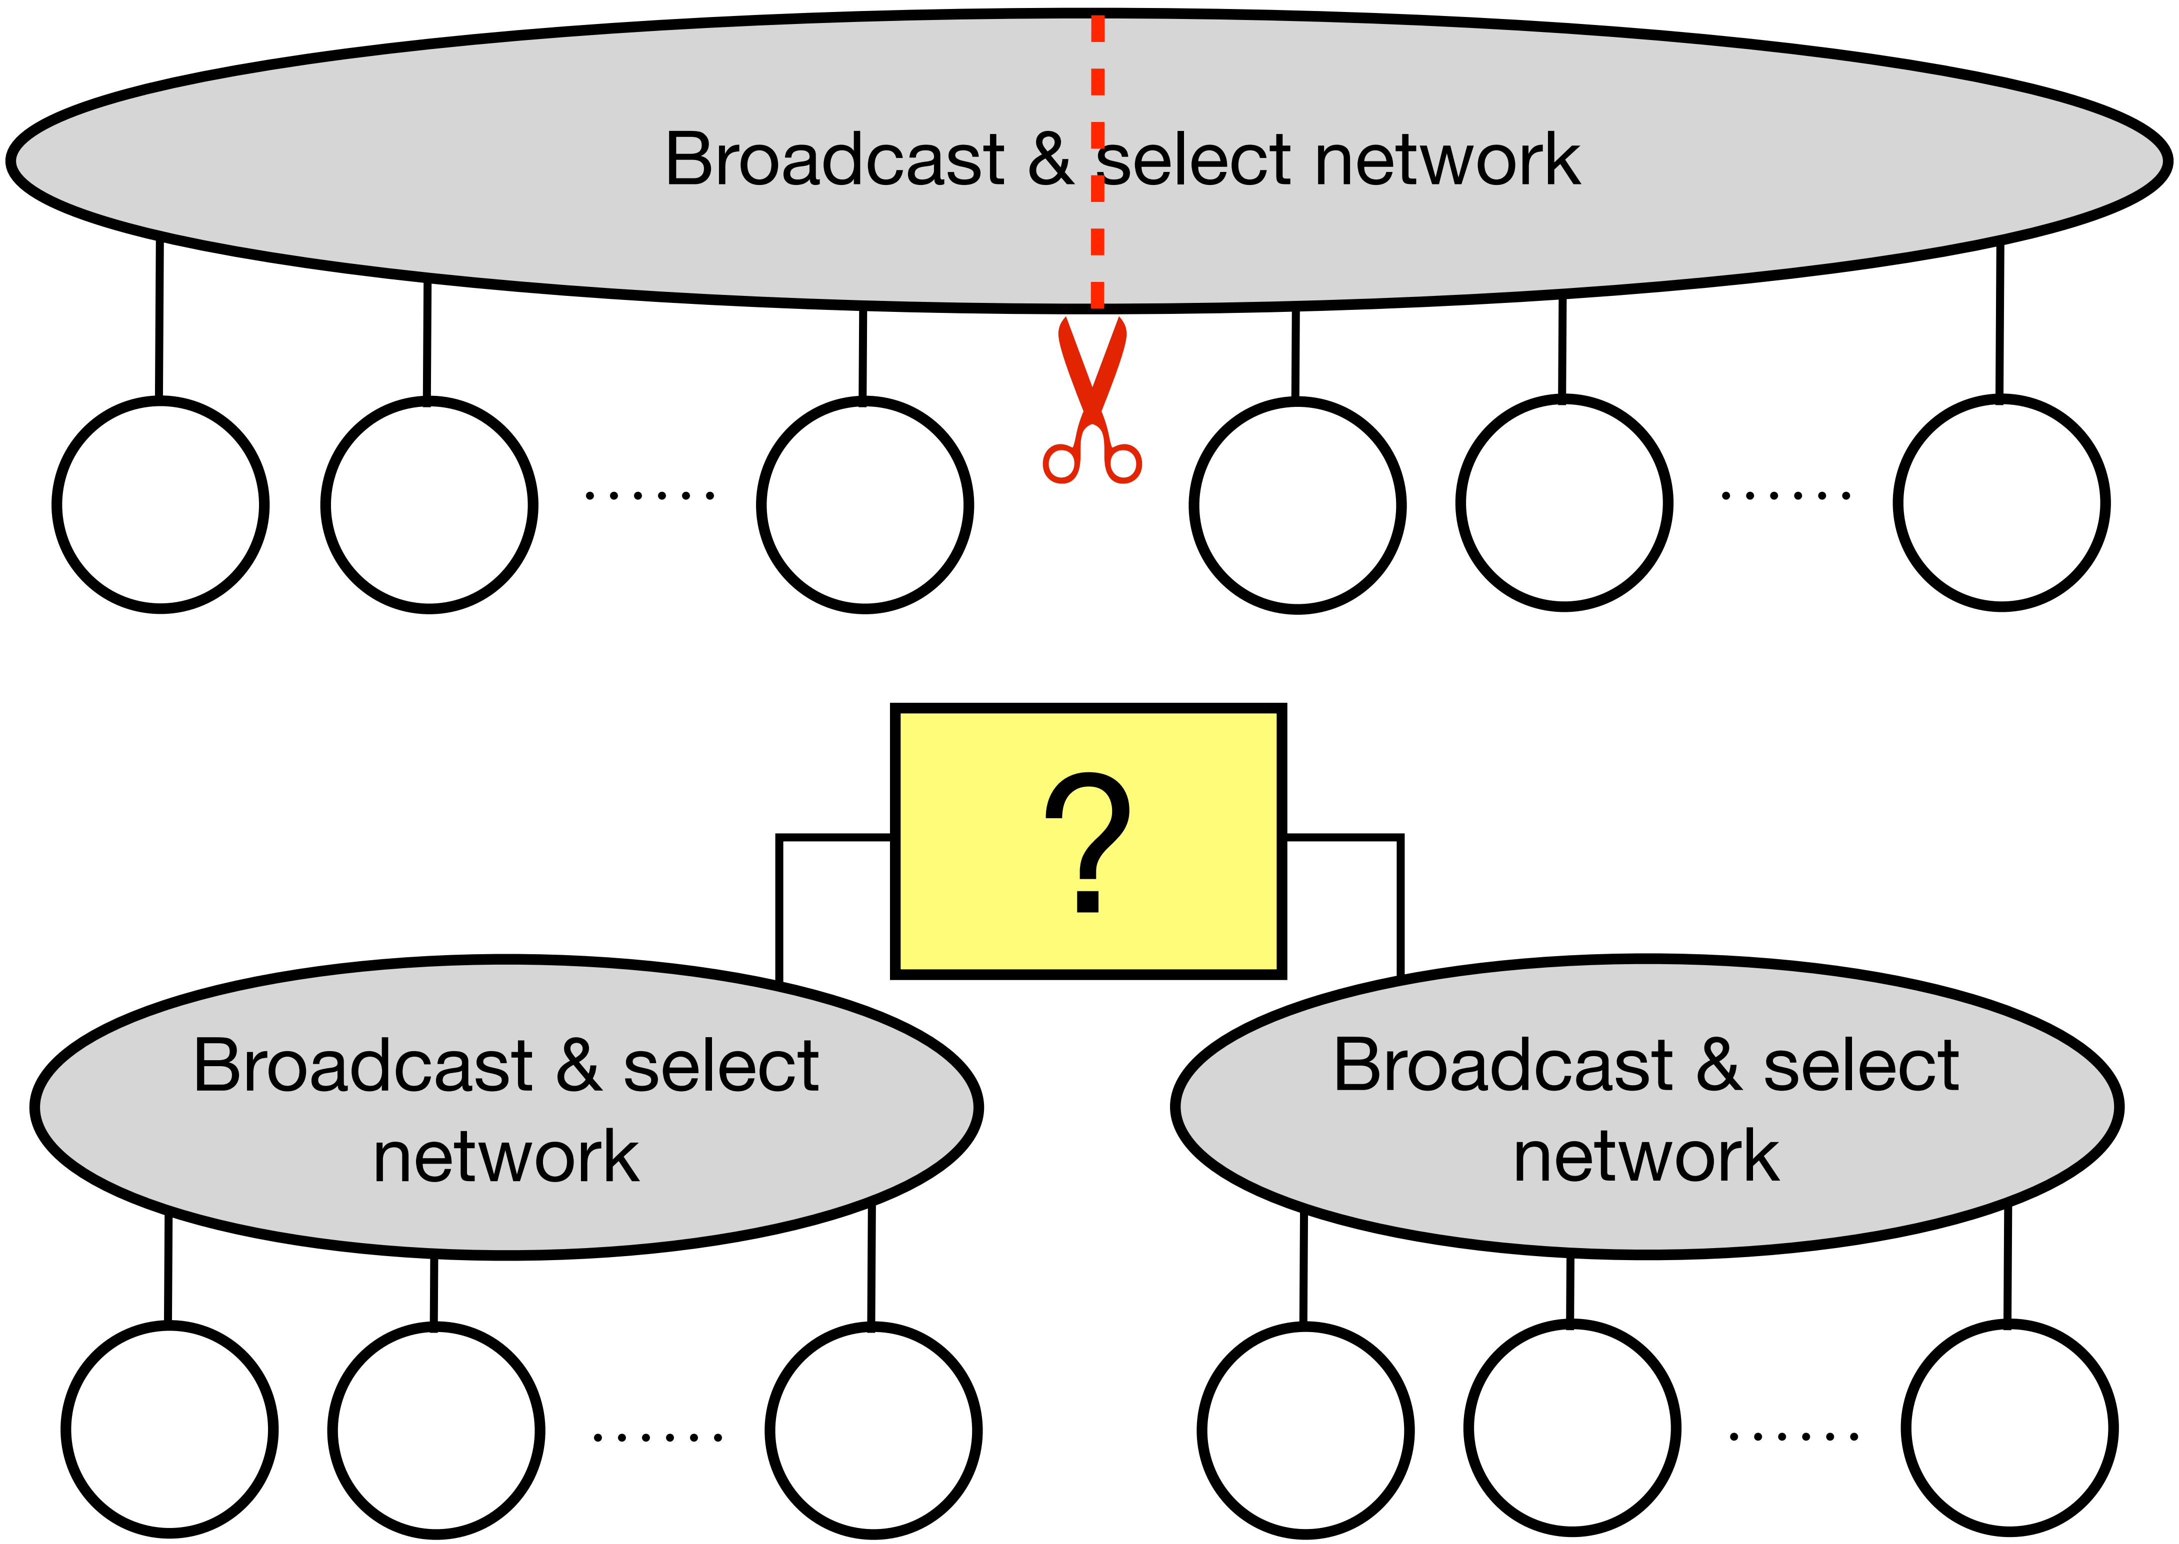

Link layer bridging

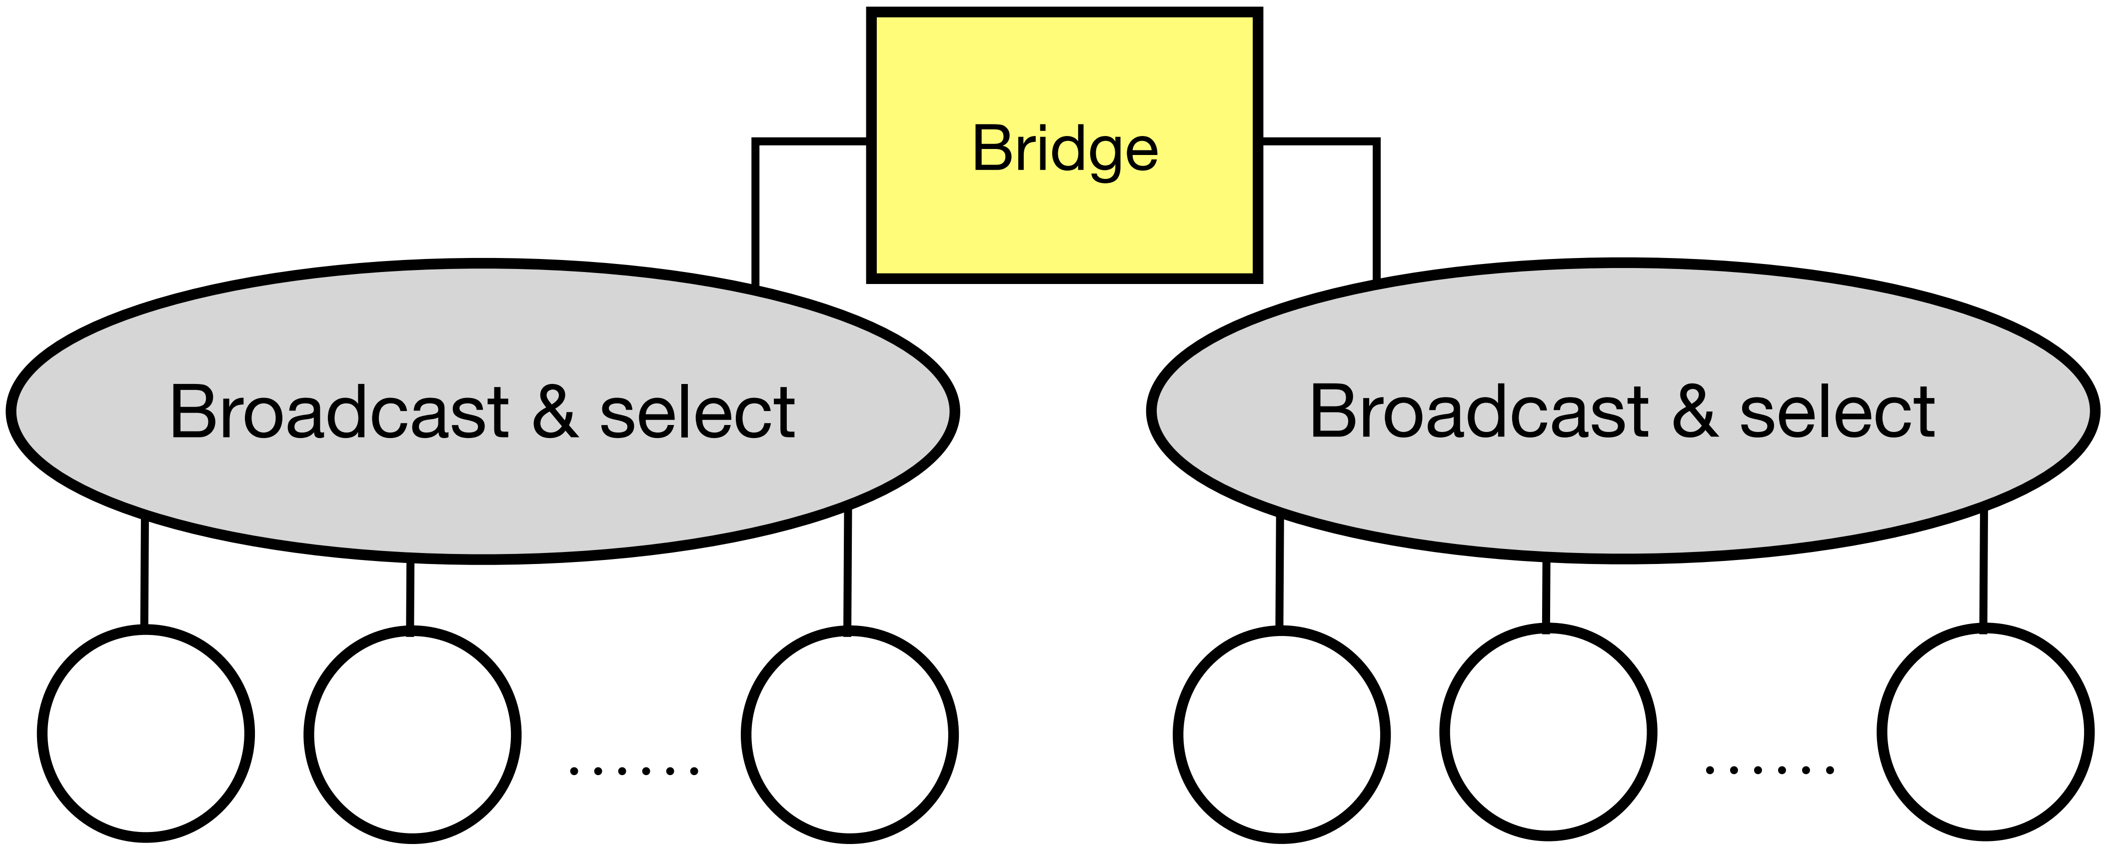

- Bridge “opens” for non-local traffic and broadcasts
- Bridge learns node locations from passing traffic and stores them in its bridging table

Bridging table

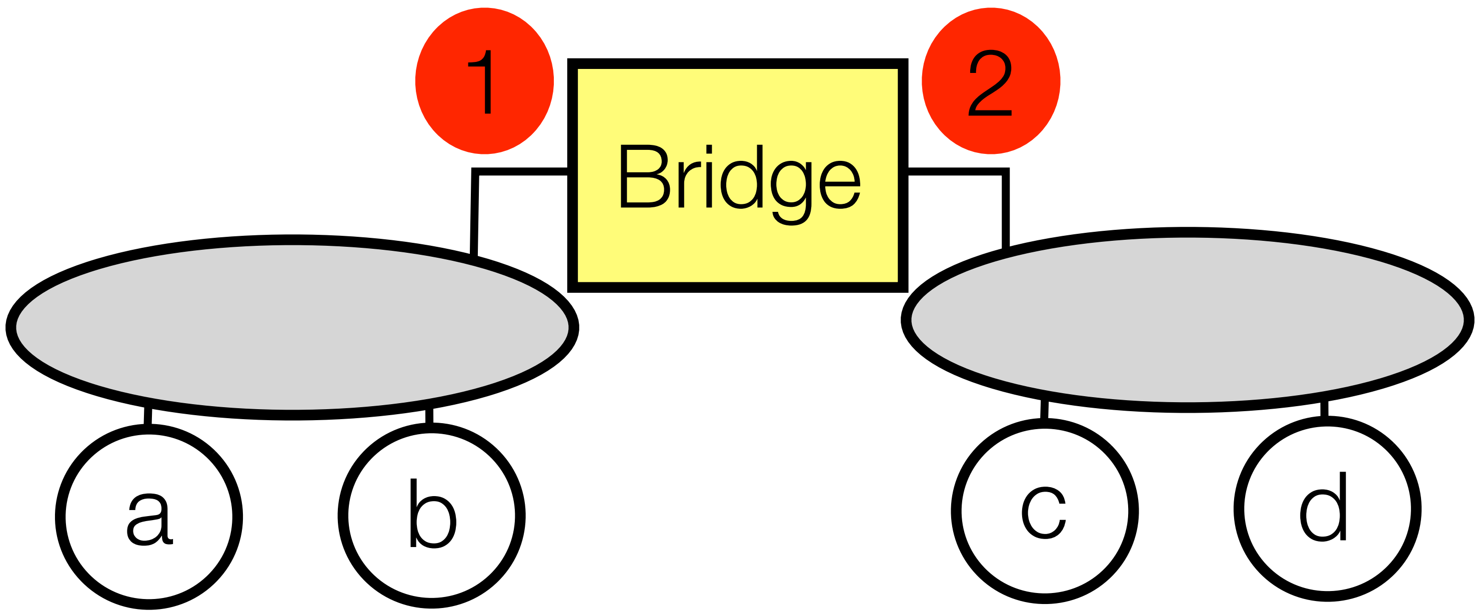

MAC	Interface
a	1
b	1
c	2
d	2

- Initially, the bridging table (BT) is empty
- Broadcast traffic is let to pass, source address recorded in the BT
- Traffic to an unknown destination is let to pass through the bridge, source address address is recorded in the BT
- Non-local traffic (to a known destination that is associated

BRIDGING - EXAMPLE

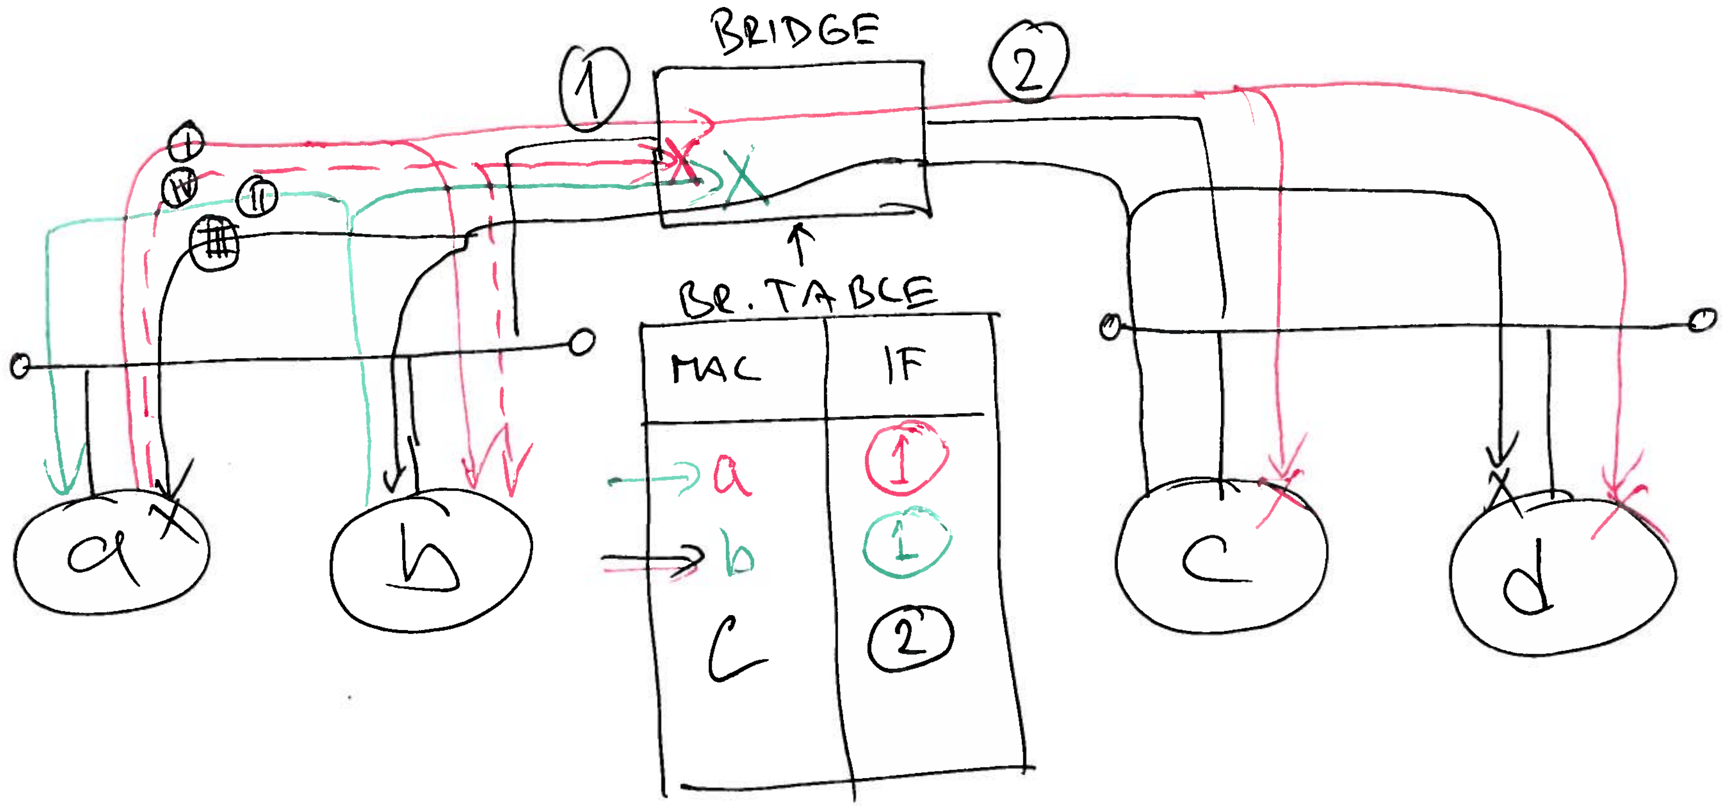

1. (I) a sends pkt to b
2. (II) b sends response a
3. (III) c sends to b
4. (IV) a sends to b ---

Scaling Bridged Networks

- Spanning Tree Protocols
- VLAN - virtual LAN (virtual L2 overlay network)

Spanning Tree Protocols

- Loops in the network topology help to increase resiliency of the network
  - but introduce problems when populating bridging tables
- **Solution**: Spanning Tree Protocols (STP)
  - temporarily disable links to break loops
  - monitor health of active links and re-enable links if network partitioning is detected
  - tradeoff: link health monitoring overhead vs repair latency

SP. TREE PROTOCOL (STP)

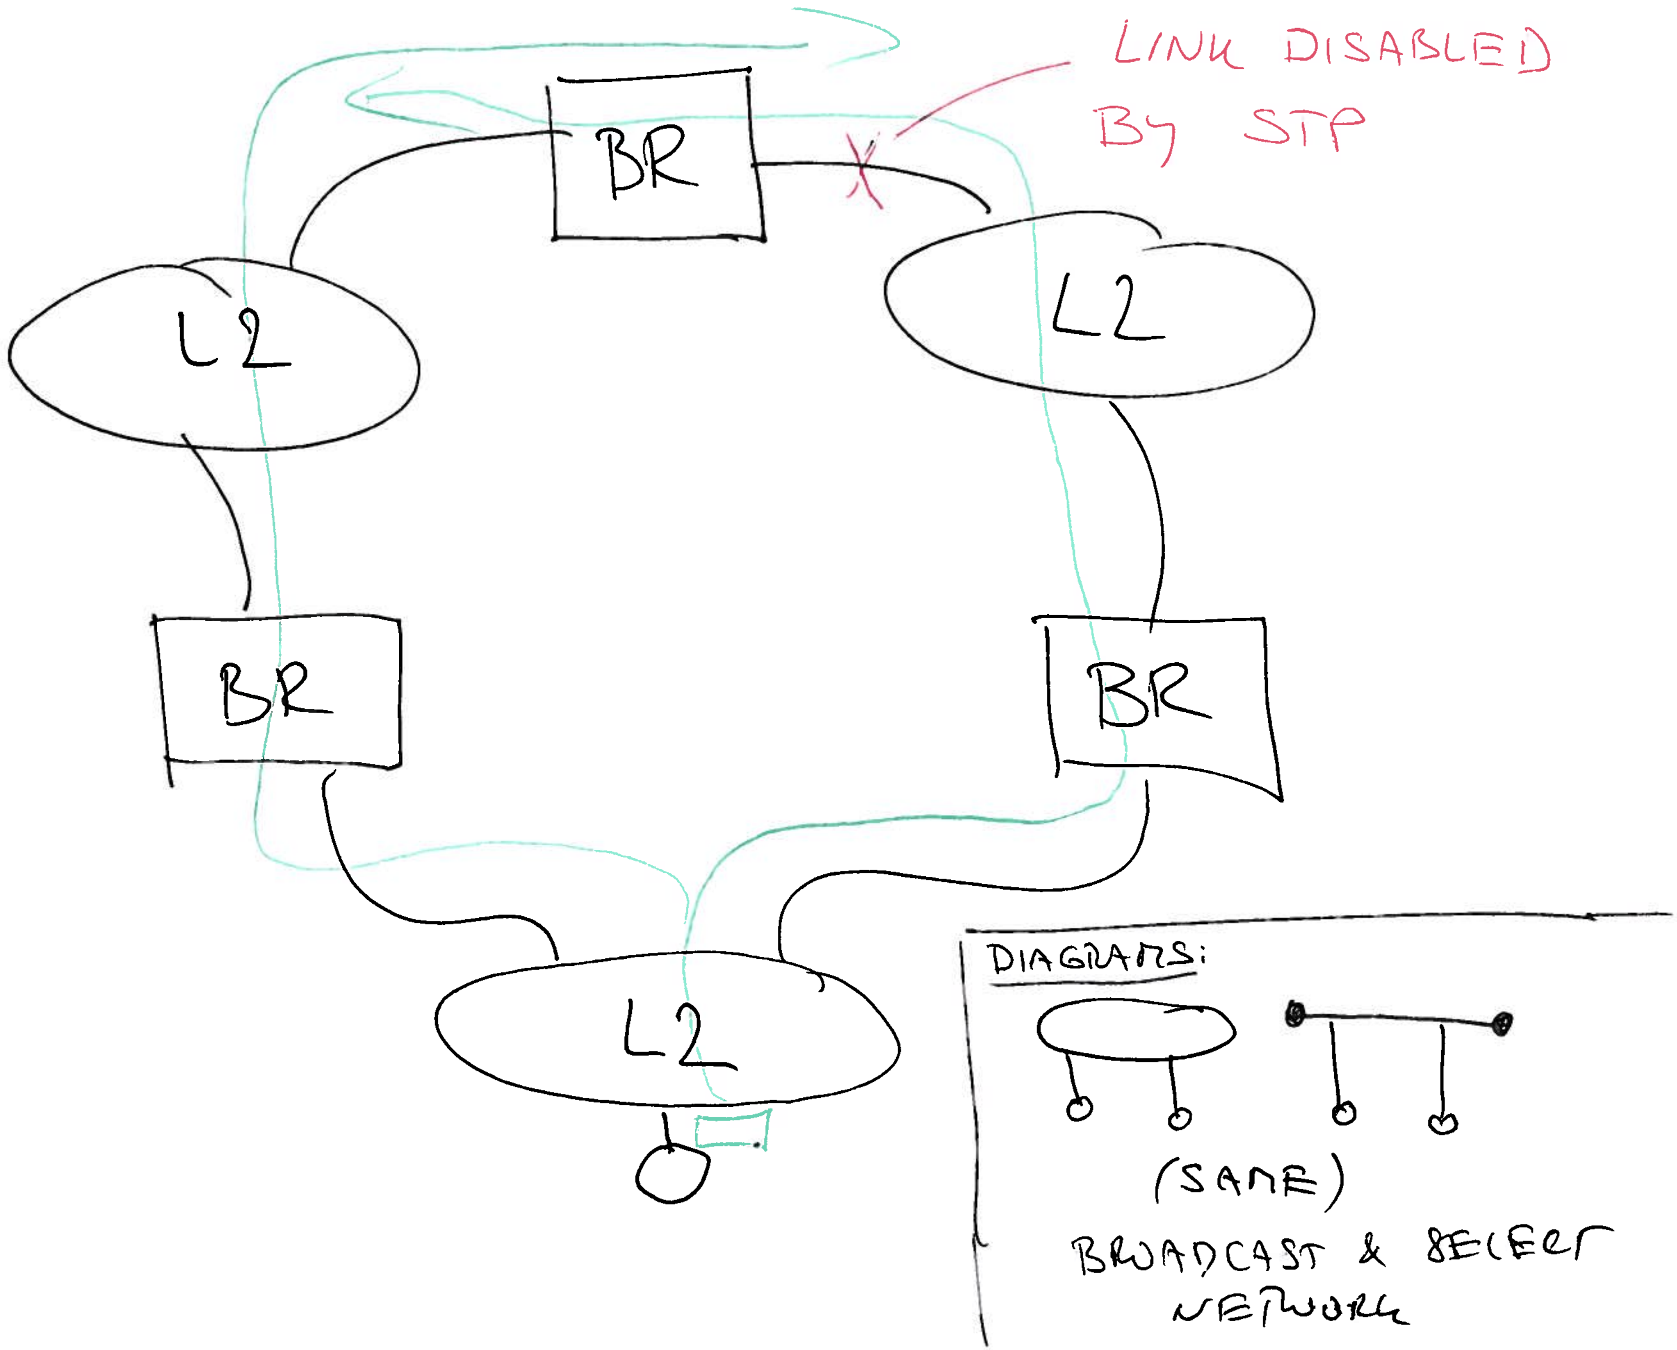

Virtual LANs (VLANs)

- Bridging/switching eliminates delivering unicast traffic that is not destined to the node
- Does not work for L2 broadcast traffic (still has to be delivered to all nodes)
- **Solution**: Virtual LANs (VLANs)
  - broad approach: decouple logical and physical topology: virtual networks, overlay networks, ...
  - specific approach: break broadcast domains into smaller ones
  - other benefits: QoS, security, control, ...