
flokinet-006 - Layer 2 Broadcast vs Unicast

Nicholas Morrison nick@nanocat.net



Broadcast vs Unicast

Broadcast Frames

- sent to every host
- processed by every host
- MAC address destination FF:FF:FF:FF:FF:FF
- eg: ARP, IPv6-NDP, OSPF, DHCP
- all about discovery

Unicast Frames

- sent to a single host
 - ignored by other hosts
 - MAC address destination of the host
 - eg: ping, ssh, http, ftp, traceroute
 - transferring data from host to host
-

Broadcast Domain


- area of the network in which a broadcast frame is repeated
- broadcast frames do not leave their own broadcast domain
- usually this will be:
 - a LAN, or
 - a VLAN
- but also could be:
 - VXLAN
 - MPLS-L2VPN
 - ZeroTier
 - EVPN-L2VPN

Broadcast example: ARP

- ARP: Address Resolution Protocol, Layer 2, ethertype 0x0806
- maps MAC addresses (layer 2 addresses) to IP addresses (layer 3 addresses)
- example flow:
 - Host A at 192.168.10.55/24 wants to talk to Host B at 192.168.10.10/24
 - Host A and B are in the same network (layer 3)
 - and the same broadcast domain (layer 2)
 - Host A needs to know the MAC address of Host B to send it traffic
 - Host A sends a Broadcast to the network
 - * ARP WHO-HAS 192.168.10.10 TELL 192.168.10.55
 - Host B hears the Broadcast
 - Host B replies directly to Host A from its own MAC address
 - * 192.168.10.10 IS-AT xx:xx:xx:xx

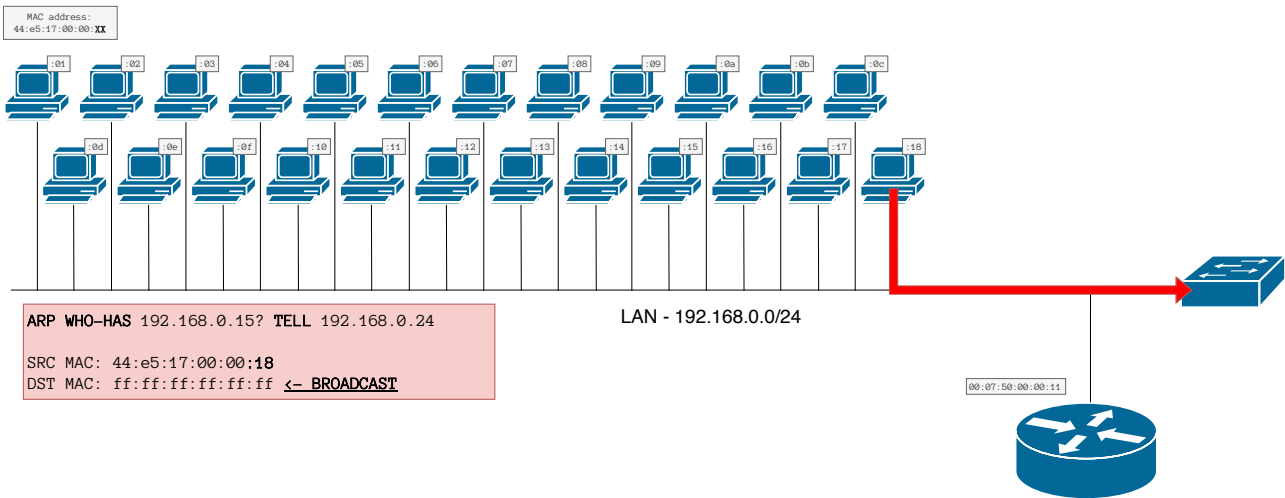
Broadcast example: ARP

A single broadcast domain shared by 24 hosts.



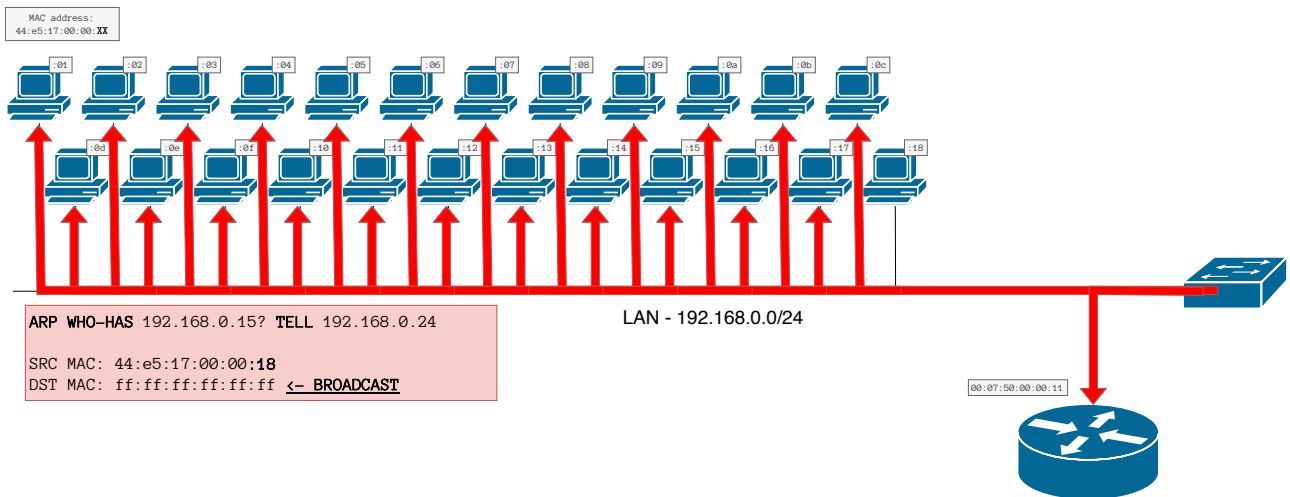
Broadcast example: ARP

192.168.0.24 (MAC 44:e5:17:00:00:18) wants to send a packet to 192.168.0.15, so it needs to know 192.168.0.15's MAC address. It sends a single layer 2 broadcast to the network, addressed to ff:ff:ff:ff:ff:ff.



Broadcast example: ARP

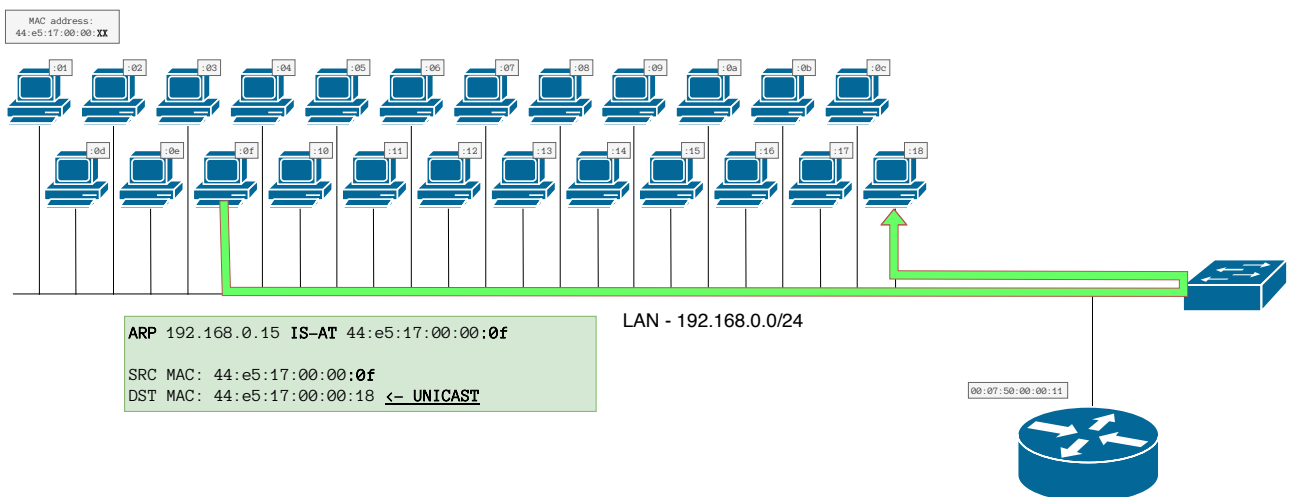
The switch sees the broadcast coming in on port 24, and copies that broadcast to every port in the same VLAN, **except port 24**. Every host must send this packet to its CPU and take a look at it.



Broadcast example: ARP

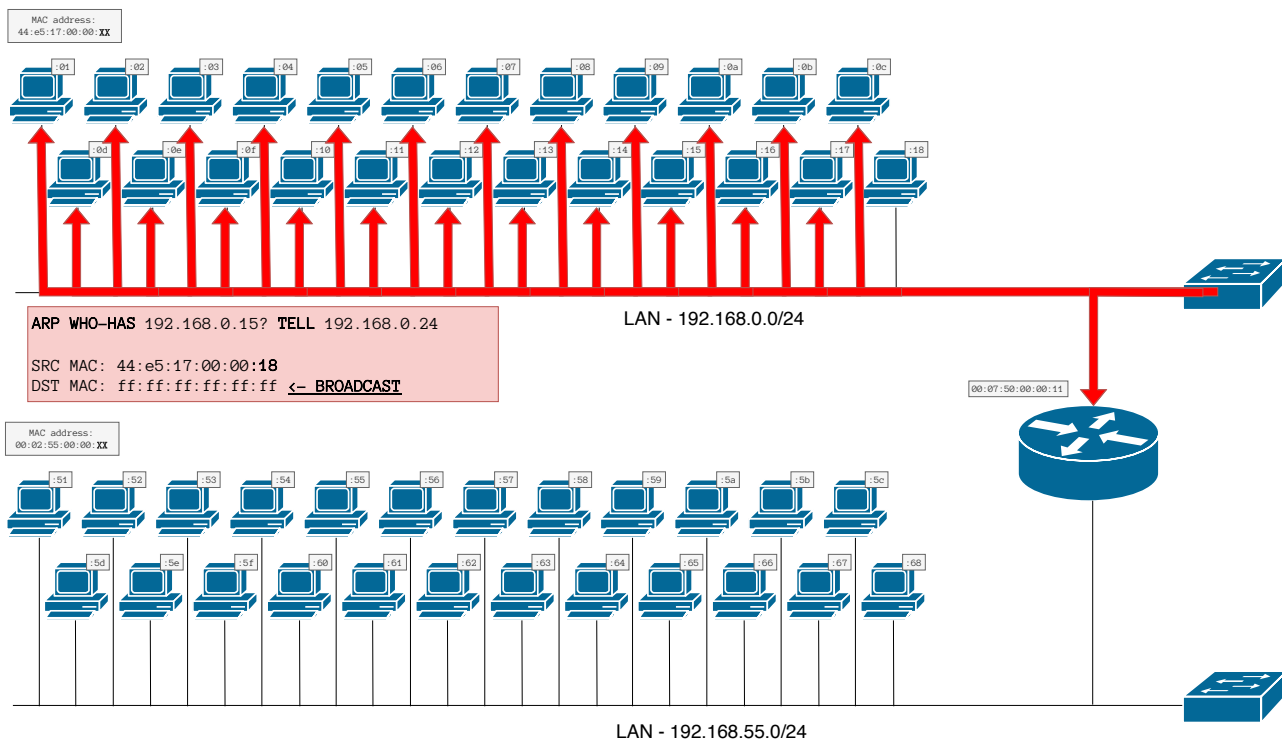
The host that HAS the address 192.168.0.15 realises that this question is for it, and replies quietly (with a UNICAST frame) to the host who asked. All other hosts stay silent on the matter.

Now the two hosts can communicate directly with one another!



Broadcast example: ARP

Broadcast packets are not forwarded by routers into other broadcast domains!



Broadcast and (Unknown-)Unicast and Multicast and Switches

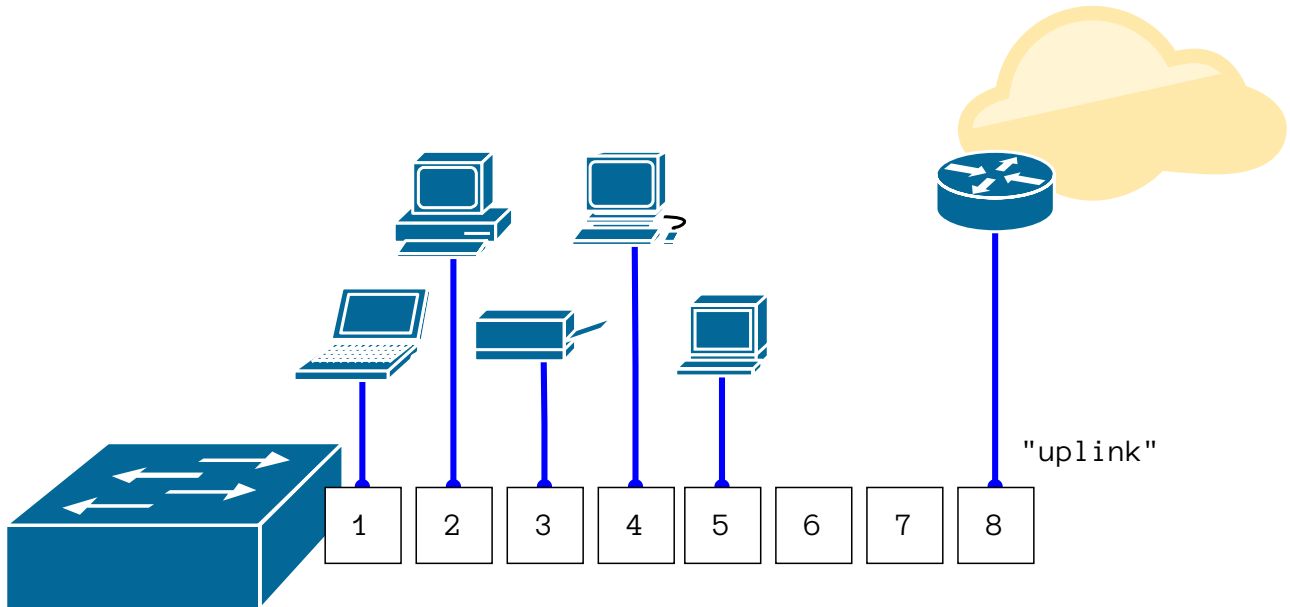
- Layer 2 switches optimise traffic flow by:
 - remembering which MAC address is connected to which switch port (show mac address-table)
 - only copying frames to ports when it's necessary
- Broadcast packets **must always** be sent to **every port** in the broadcast domain!
 - **except** the port from which the broadcast packet was received (split-horizon)
- **Unknown-Unicast** packets are **flooded** to every port until the switch learns which port has the MAC address attached.
 - Once they're learned, they are just Unicast packets.
- Multicast packets **must always** be sent to **every port** in the broadcast domain
 - **except** in the case that the switch is doing IGMP-snooping
 - IGMP-snooping can "prune" ports from the Multicast group
- Broadcast, Unknown-Unicast and Multicast packets are known collectively as **BUM** packets

Layer 2 Loops and Broadcast Storms


- Creating a loop between switches can break your network
 - broadcast packets will loop between switches as fast as the switches can go
 - each broadcast packet will be copied to every port (except the receiving port)
 - every device must process the broadcast packet

- ... raising CPU on every device in the broadcast domain
- The Spanning Tree Protocols (STP, MSTP, RSTP, PVSTP et al) are designed to break such loops

Happy network with no loop




Happy network with two switches and no loop




Sad network a loop between two switches


Somebody disabled spanning tree.



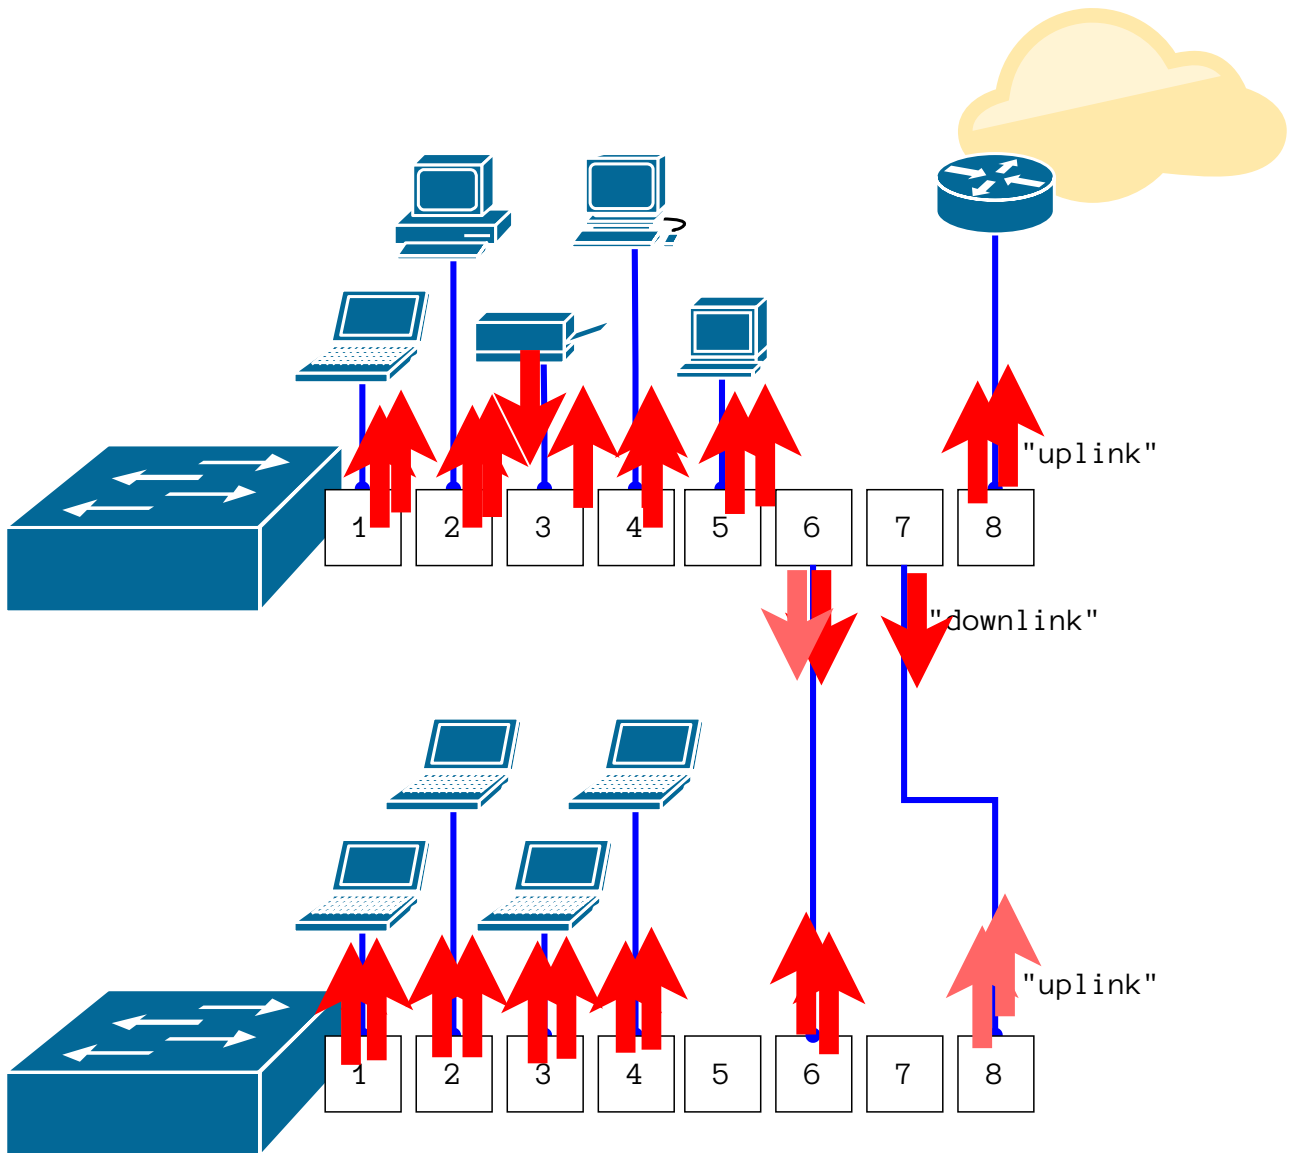
The beginning of the storm - an innocent "arp who-has"




The broadcast who-has is repeated by the switches



.. and the broadcast storm begins ..



.. and continues until devices run out of capacity



Aftermath

- Since every single device has to process a broadcast packet, **all** devices on the LAN / Broadcast Domain will be affected by the broadcast storm.
- Because it is often consumed before network switching capacity, unexpected failures can occur, such as routing protocols being unable to send their HELLOs quickly enough
- Spanning Tree Protocol is used to break layer 2 loops
- Always* run spanning tree!