

Advanced features on Ethernet networks

Fulvio Riso

Politecnico di Torino

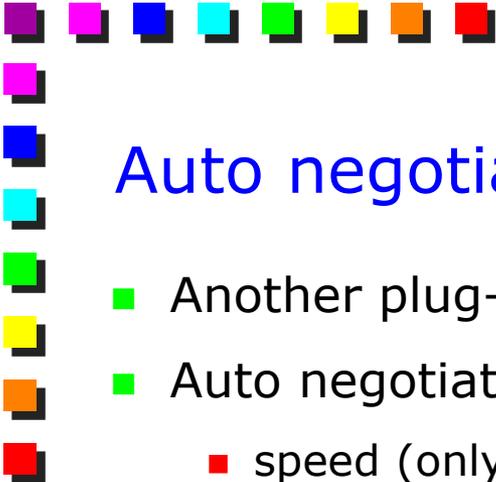




Outline

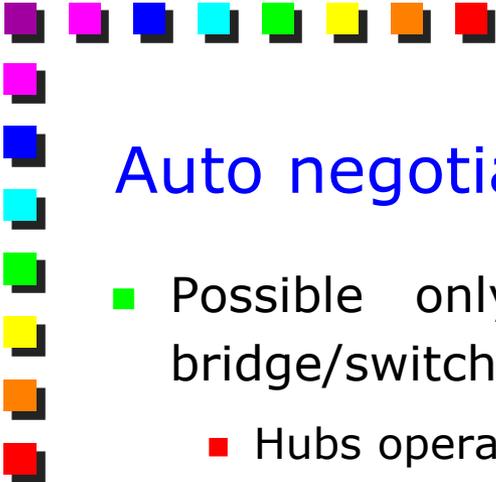
- Autonegotiation
- Jumbo frames
- Power Over Ethernet (PoE)





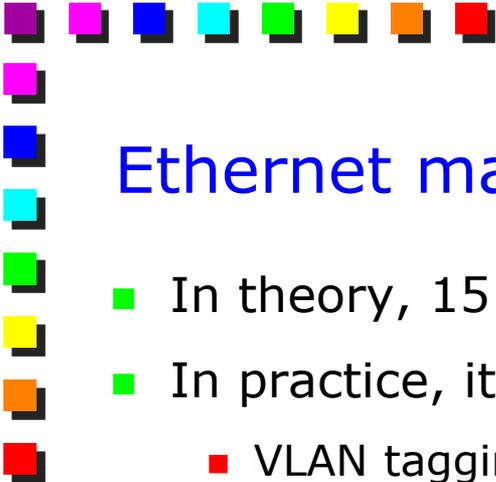
Auto negotiation (1)

- Another plug-and-play oriented feature
 - Auto negotiation possibilities:
 - speed (only over copper)
 - half/full duplex (over copper and fiber optic)
 - Negotiation sequence:
 - 1 Gb/s full-duplex
 - 1 Gb/s half-duplex
 - 100 Mb/s full-duplex
 - 100 Mb/s half-duplex
 - 10 Mb/s full-duplex
 - 10 Mb/s half-duplex
- 

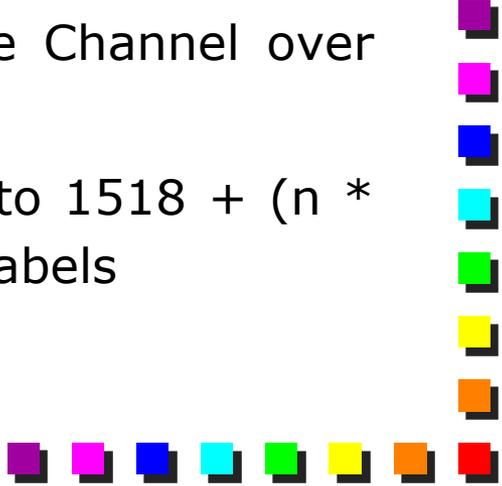


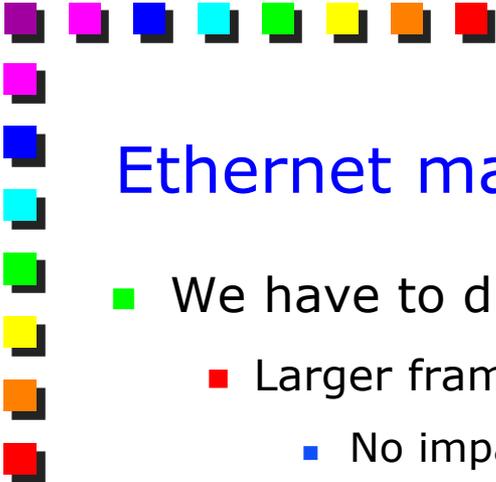
Auto negotiation (2)

- Possible only if connected to another host, or to a bridge/switch
 - Hubs operate at fixed speed; hence, cannot negotiate anything!
- If, during the procedure, the other party does not respond, the negotiating station assumes it is connected to an hub
 - Fixed setting on one side may lead to unexpected errors
- Example
 - One side: fixed 100Mbps Full Duplex
 - The other party does not receive any message and it will assume it is connected to an hub
 - It will configure the interface in 100Mbps Half Duplex
 - There may be tons of false collisions on that host

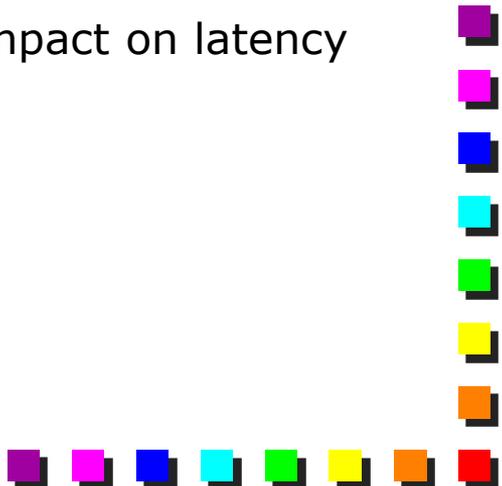


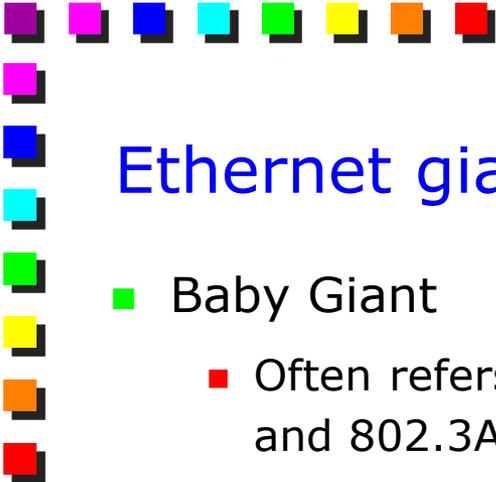
Ethernet max frame (1)

- In theory, 1518 bytes (the original specs)
 - In practice, it has been enlarged several times
 - VLAN tagging (IEEE 802.1Q), + 4 bytes
 - Provider Bridge (802.1ad) (also known as 802QinQ), + 8 bytes
 - Ethernet Frame Expansion (802.3as) proposed a new size of 2000 bytes
 - Size of the data portion of the frame (46-1500 octets) does not change (MTU still 1500 bytes)
 - T11 adopted an MTU of 2500 bytes for Fibre Channel over Ethernet (FCoE) frames
 - MPLS increases the max Ethernet frame size to $1518 + (n * 4 \text{ bytes})$, where "n" is the number of stacked labels
- 



Ethernet max frame (2)

- We have to distinguish
 - Larger frames because we need to add more headers
 - No impact on the hosts
 - Required by new Ethernet extensions
 - Larger frames because we need to transport more data
 - MTU impact
 - Desirable for different reasons (see later)
 - Ethernet 1500 bytes MTU for historical reasons
 - Can enlarge the max frame without significant impact on latency
 - Backward-compatibility issues
 - In fact, GbE has frame bursting
- 



Ethernet giants

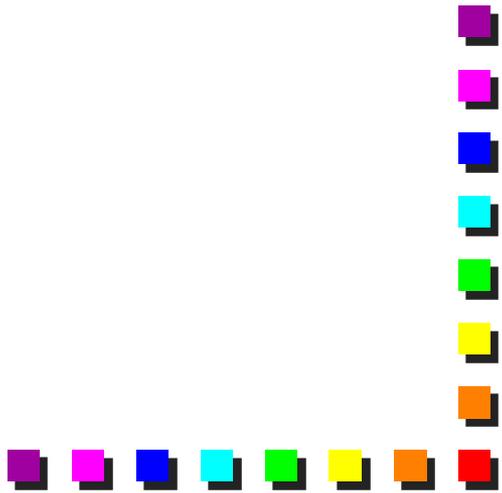
- Baby Giant

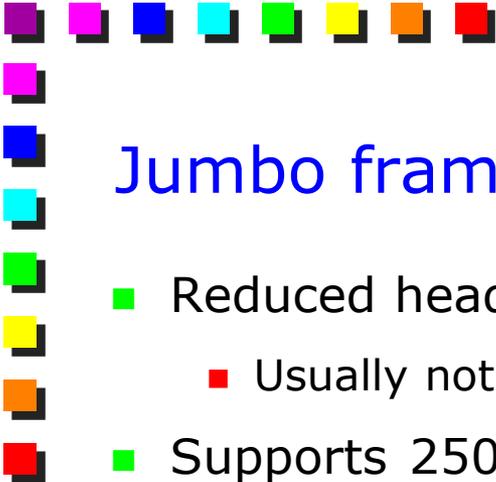
- Often refers to the frame type used with MPLS, 802.1Q, 802.1ad and 802.3AE
- Just headers

- Mini Jumbo

- Often used to refer to an MTU size of 2500 bytes and has become specific to the frame size used by FCoE
- Bigger data

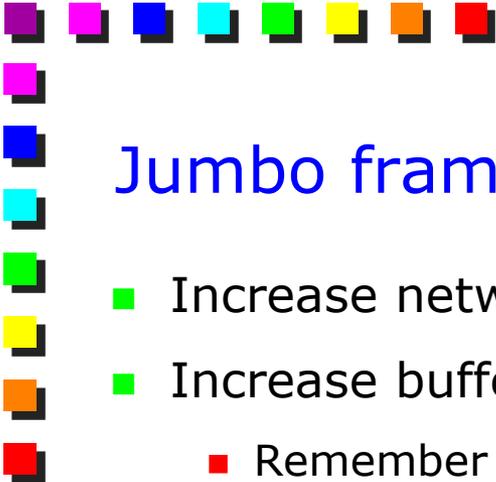
- Jumbo (or “Giants” / “Giant Frames”)

- Bigger data
 - Often frames go up to 9KB
 - Non standard, but widely supported
- 



Jumbo frames: advantages

- Reduced header overhead
 - Usually not very important in practice
 - Supports 2500bytes FCoE frames, 8192 NFS data blocks, 8K iSCSI blocks (derived from the typical TCP Windows size)
 - Reduced operating system overhead (interrupt handling at each new packet)
 - Very significant cost for network-intensive servers
 - Possibility to transport traffic coming from networks with larger MTUs
 - Used long back in the past (e.g. to transport Token Ring frames on corporate backbones)
- 



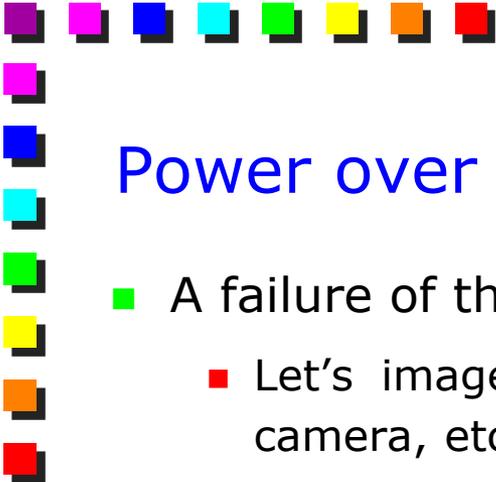
Jumbo frames: problems

- Increase network latency
 - Increase buffer pressure on switches
 - Remember that usually switches have very small buffers
 - Network stacks (and operating systems) often tuned for 1500bytes
 - Internal buffers dimensioning
 - Fragmentation (at the IP level)
 - If a jumbo frame has to be delivered to a traditional station
 - A new CRC calculation had to be defined
 - The original algorithm was not robust enough with those sizes
- 

Power over Ethernet

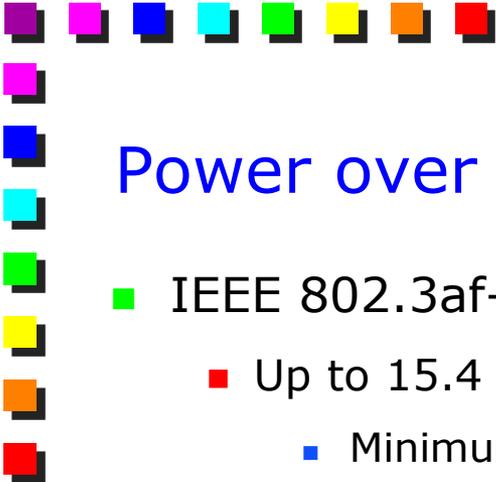
- Distributes electrical power on the Ethernet cable
 - Twisted pair only (no fiber)
- Useful to connect users with moderate power needs
 - VoIP phones
 - WiFi access points
 - Surveillance camera
 - Etc.
- Avoids additional cabling due to the power cord
- More power and longer distances than USB





Power over Ethernet: problems

- A failure of the PoE switch may be catastrophic
 - Let's image a building in which network, phones, surveillance camera, etc. stop working
 - Panic!
 - Hence, we need to install switches with proper redundant power supplies
- More power consumed by the switch
 - 48 ports at 25W = 1.2 KW



Power over Ethernet: standards

- IEEE 802.3af-2003
 - Up to 15.4 W of DC power to each device
 - Minimum 44V DC and 350mA
 - Only 12.95 W is assured to be available at the powered device as some power is dissipated in the cable
 - Updated with IEEE 802.3at-2009
 - Known as PoE+ or PoE plus
 - Up to 25.5 W of power
 - Products available that provide even more power (e.g., 51W) over a single cable by utilizing all four pairs in the Cat.5 cable
 - Compatible with non-PoE stations
- 