DANTE CERTIFICATION PROGRAM

LEVEL 3 – ADVANCED DANTE NETWORKING

Dante by Audinate

ADVANCED
DANTE NETWORKING

DANTE CERTIFICATION PROGRAM

LEVEL 3

Dante by Audinate

ADVANCED DANTE NETWORKING: INTRODUCTION

DANTE CERTIFICATION PROGRAM

Training program from Audinate

- Official certification lets your customers know that you have the knowledge and skills to implement Dante networks
- Ensures a consistent set of methods and knowledge

Dante by Audinate

DANTE CERTIFICATION PROGRAM

With Dante Certification, you receive:

- Use of the Level 1, Level 2 and Level 3 “Dante Certified” logos
- A certificate of completion for each level passed.

Dante by Audinate

DANTE CERTIFICATION PROGRAM

Level 1: Introduction to Dante

- 100% online delivery
- Background
- Basic signal routing
- Setting up Dante in simple systems (approximately 6 devices, 1 switch)

Dante by Audinate

DANTE CERTIFICATION PROGRAM

Level 2: Intermediate
Dante Concepts

- Delivered in-person and online
- Larger systems (approx. 12 devices)
- Clocking options
- Understanding unicast & multicast
- Latency
- Redundancy
- Dante Virtual Soundcard and Dante Via

Dante by Audinate

DANTE CERTIFICATION PROGRAM

Required steps:

- **Level 1:** Pass Level 1 **online** exam
- **Level 2:** Pass Level 2 **online** exam PLUS skills test at event or online

Dante by Audinate

UNDERSTANDING SCALABLE DANTE NETWORKS

Dante Topics

- Motivation
  - Who are Audinate
  - Why use Dante
  - Where is Dante Used
  - What is Dante
  - How to work with Dante
- Architectural principles
- Assessing and optimizing Network components
- Troubleshooting
- Dante in a converged network

Dante by Audinate

ABOUT AUDINATE

Headquartered in Sydney, Australia

Network engineers first

Develop Dante as **100% interoperable solution** for all audio manufacturers

WHAT WE MAKE

Dante technology
(all of it)

Hardware modules

Development tools

Software products:

- Dante Controller
- Dante Virtual Soundcard
- Dante Via
- Dante Domain Manger

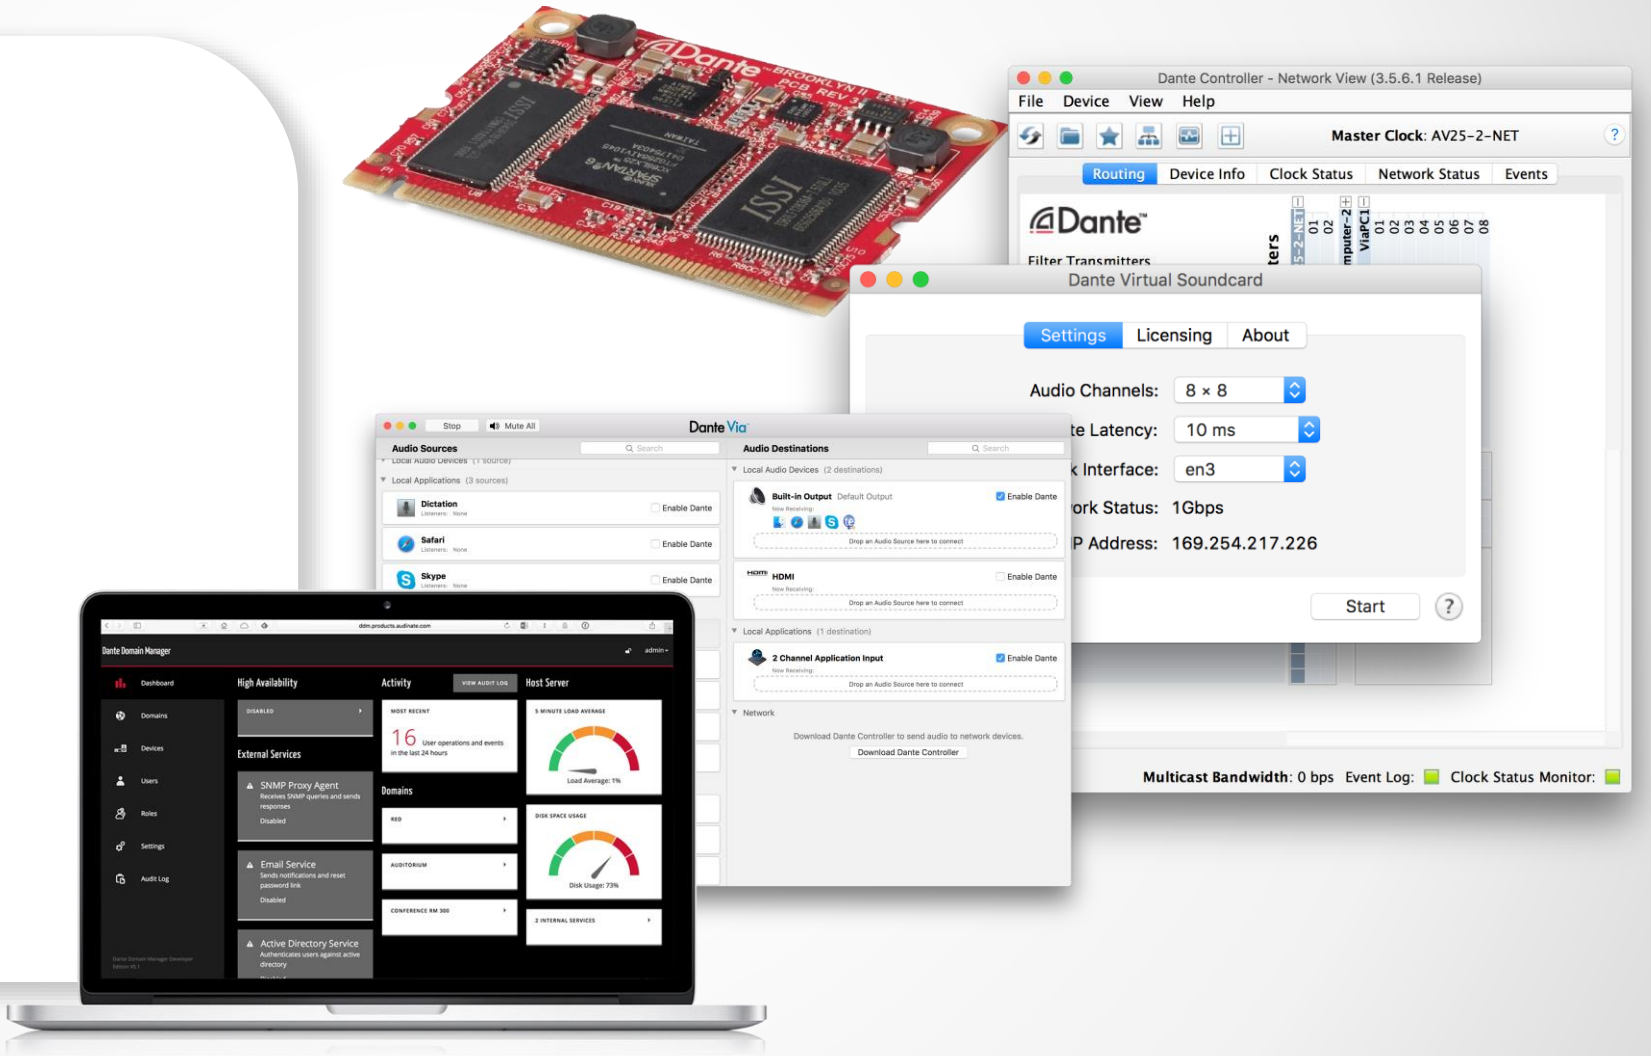

WHY USE DANTE?

WHY USE DANTE?

A network is a group of things that connect

- **Interoperability** is key
- **More than 1000** Dante-enabled products in the market
- **More than 350** OEM manufacturers
- **More than 1 million** Dante-Enabled endpoints have shipped

WHY USE DANTE?

If it has the Dante logo it will **connect to any** other device with a Dante logo

Dante is a commercially developed and supported solution – **improvements and features are added continually**

Most widely adopted audio networking solution ever

WHERE IS DANTE USED?

LIVE SOUND

AMERICAN AIRLINES ARENA – MIAMI, FL

COMMERCIAL INSTALLATION

CHESAPEAKE ARENA – OKLAHOMA CITY, OK

BROADCAST

VICTORIA PARLIAMENT HOUSE – MELBOURNE, AUSTRALIA

EDUCATION

GOVERNORS STATE UNIVERSITY – UNIVERSITY PARK, IL

HOUSE OF WORSHIP

WILLOW CREEK COMMUNITY CHURCH – SOUTH BARRINGTON, IL

FEATURE FILM PRODUCTION

LA LA LAND

POST PRODUCTION

GOLDCREST FILMS – LONDON, UK

RECORDING

SYNCHRON STAGE – VIENNA, AUSTRIA

WHAT IS DANTE?

WHAT IS DANTE?

- Dante is a networking technology
  - Hardware solutions provided to OEM manufacturers
  - Software solutions – DVS, Dante VIA, Dante Controller
  - Network API's
- Dante technology provides:
  - Tightly synchronized (better than 1µs) media playout at every endpoint in an IP network
  - Uncompressed Digital Audio at all professional sample rates
  - Simple plug and play discovery and routing across both a local area network and a routed IP network
- Dante can be deployed on COTS (commercial off the shelf) Network infrastructure

Dante by Audinate

ADVANCED DANTE NETWORKING: SECTION 1

IN THIS SECTION…

Networking

- What is a Network?
- Comparing ways of connecting Audio devices
- Comparing:
  - TDM networking techniques (time centric)
  - Packet Switched (address centric) networking techniques

WIRING UP AN "OLD" DIGITAL AUDIO SYSTEM

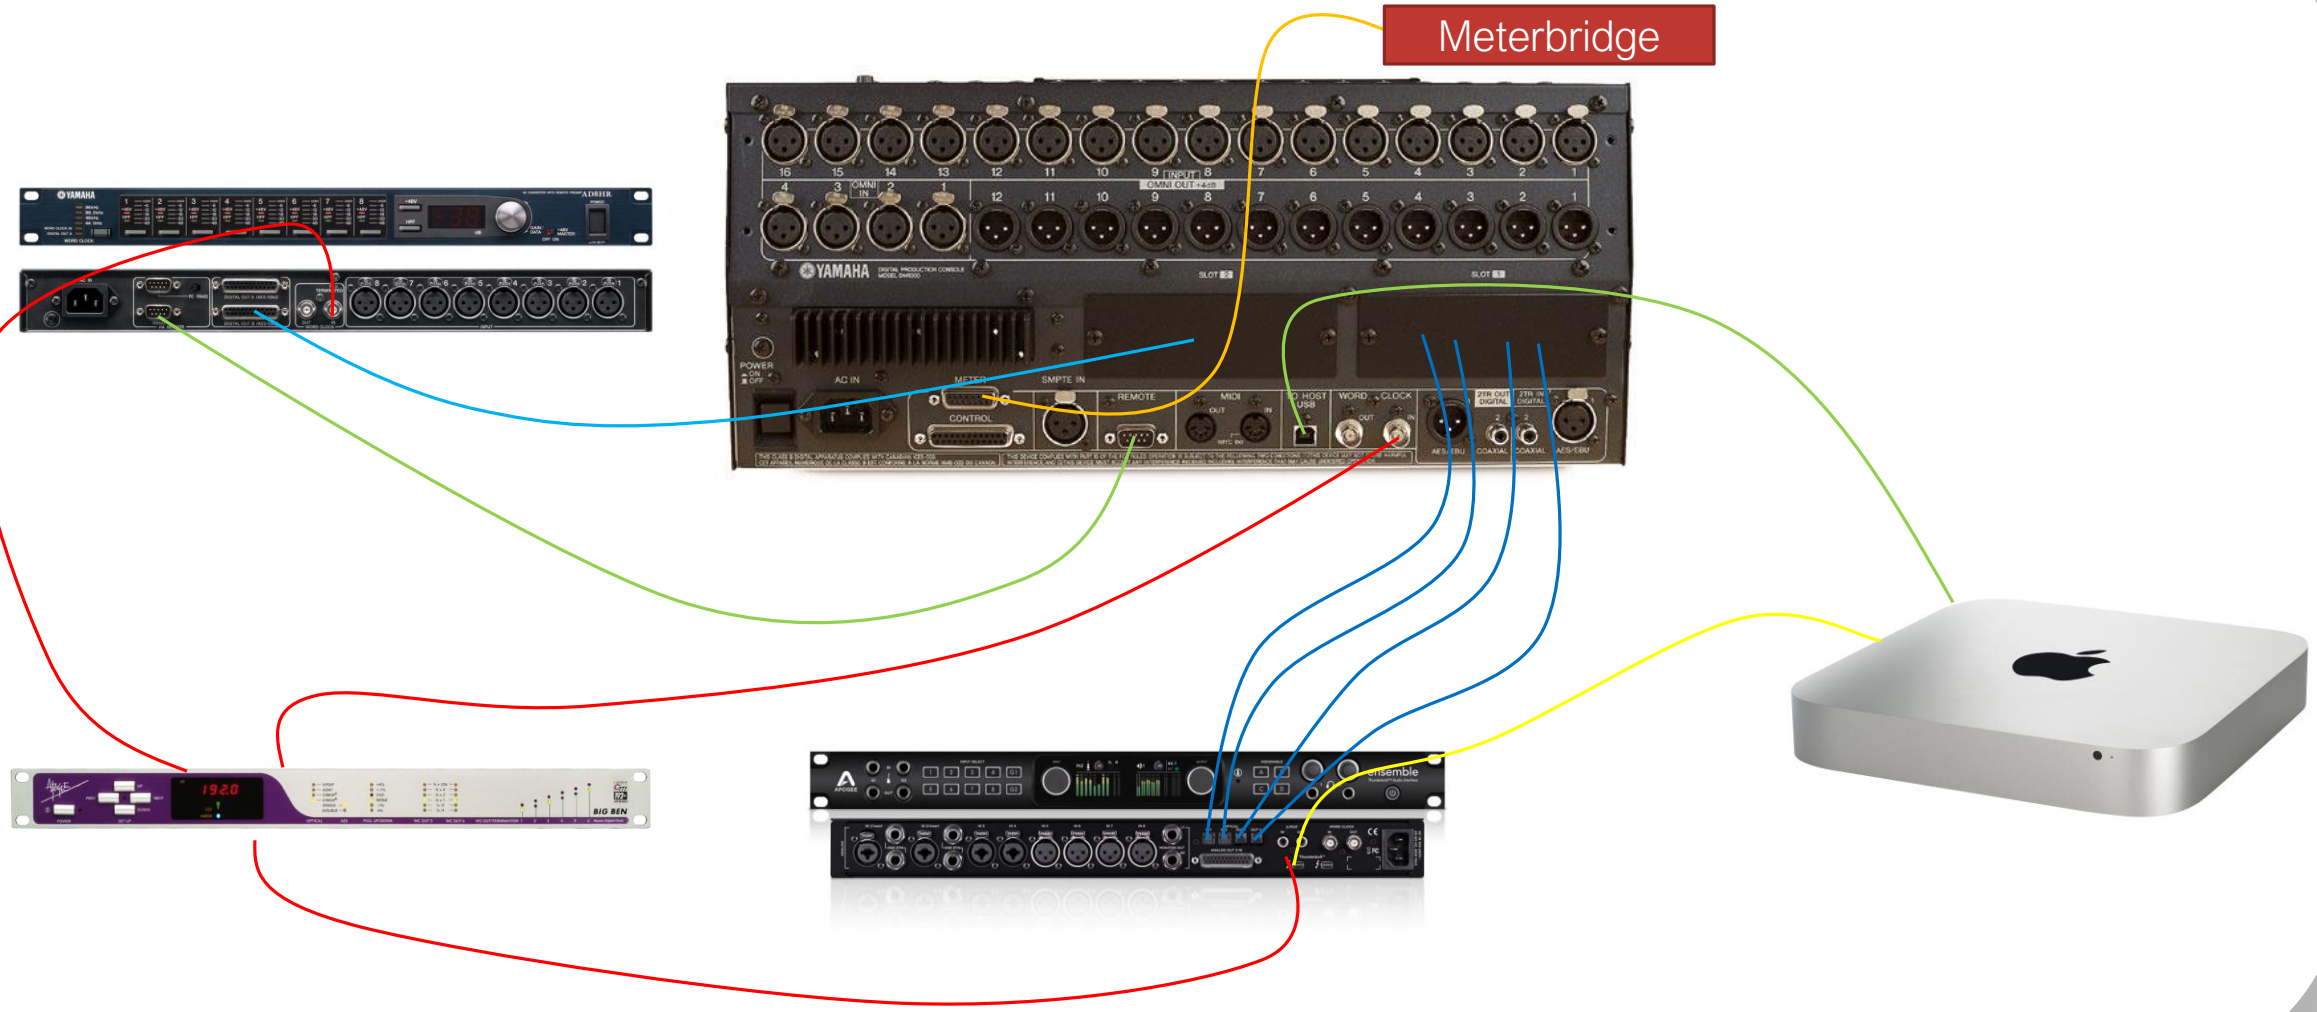

WITH DANTE

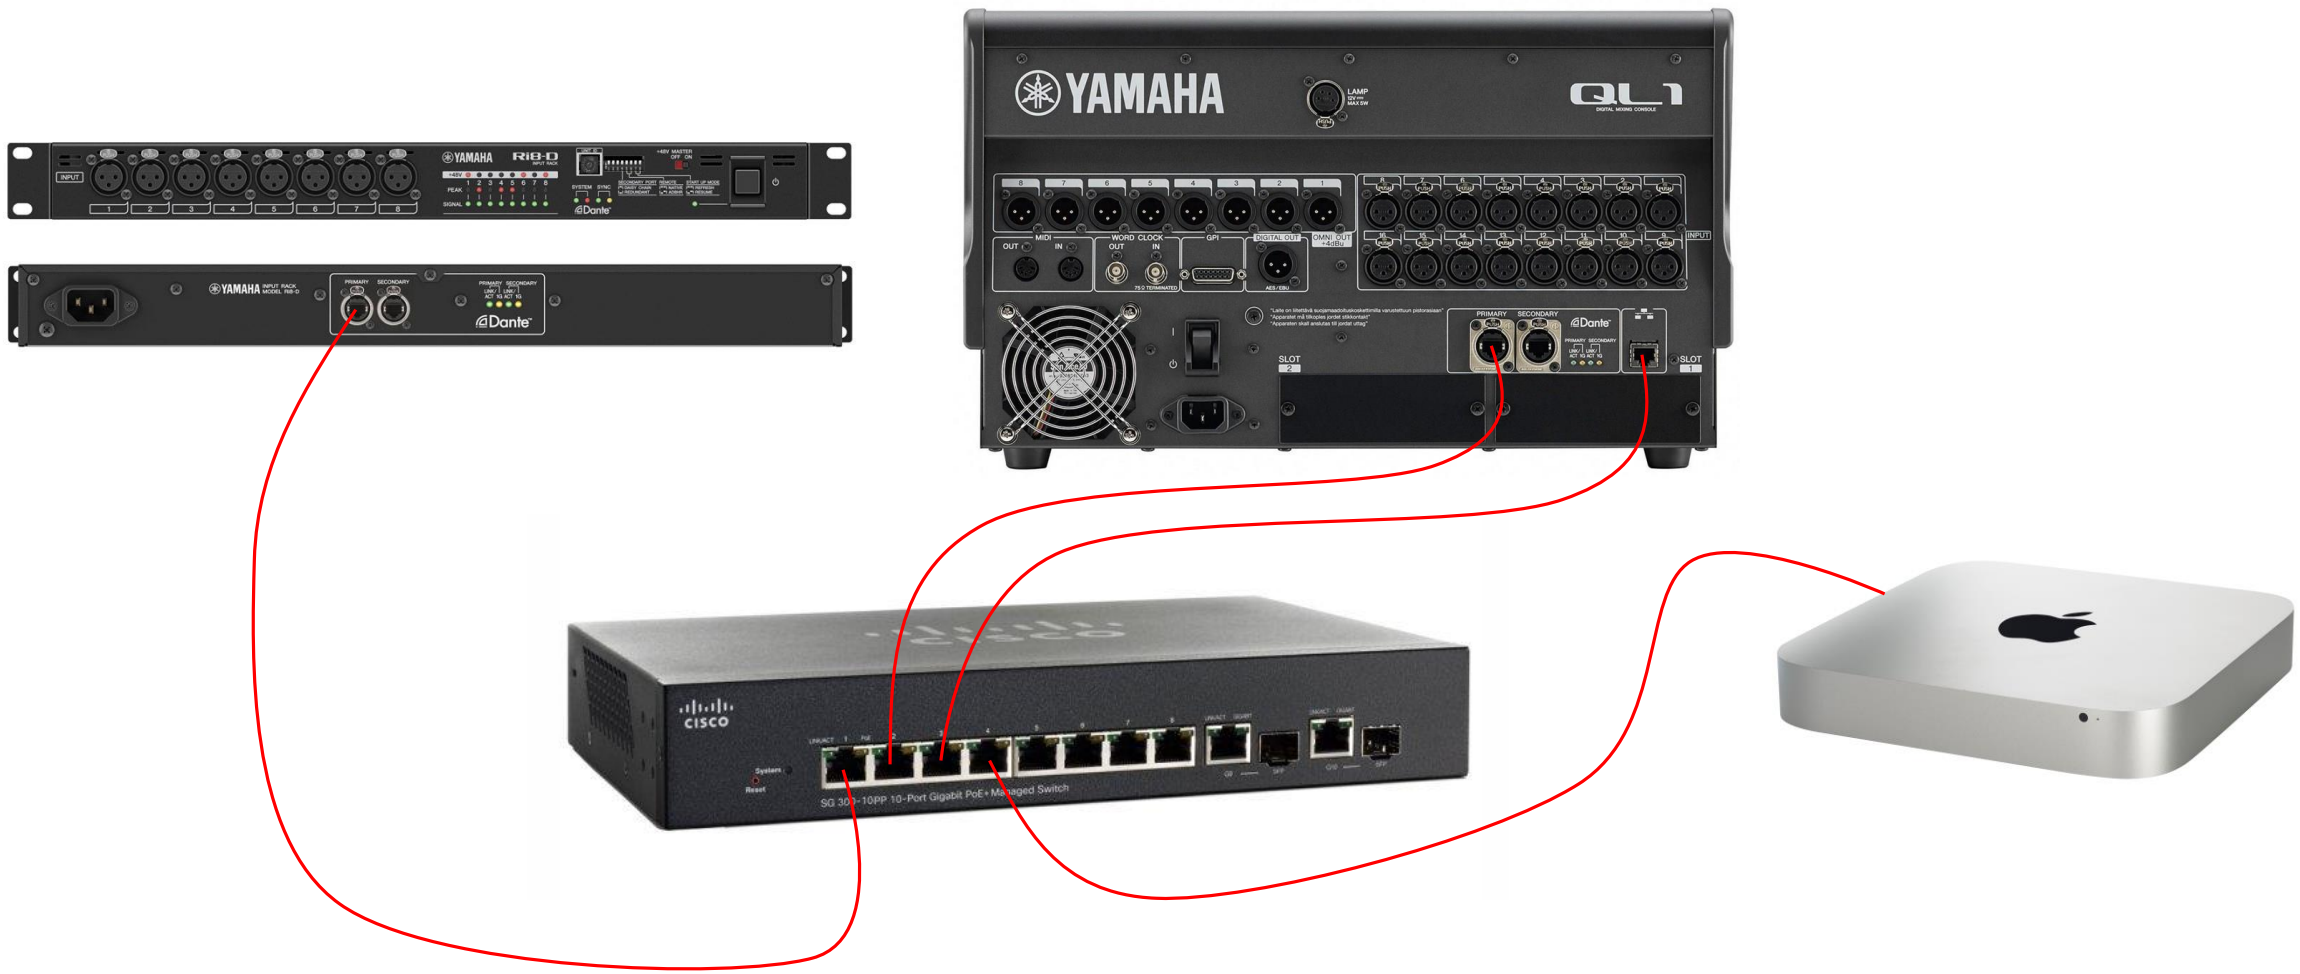

CABLING REQUIREMENTS

“OLD” DIGITAL SYSTEM

1x 9-pin serial (HA control)

1x Proprietary meterbridge cable

3x 75ohm BNC cables (wordclock)

1x 25pin D-SUB M-M (AES)

1x USB cable

1x Thunderbolt cable

4x TOS-LINK cables

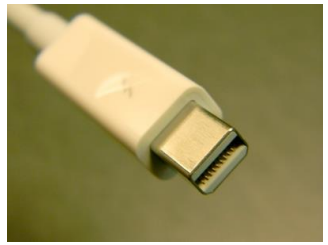

“WITH DANTE”

4x Cat5e cables

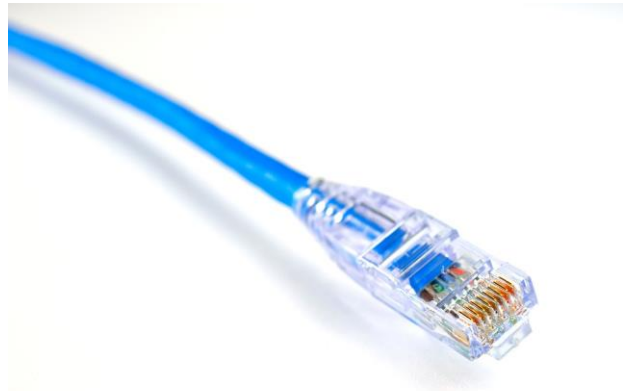

WHAT DO WE MEAN BY A NETWORK?

- Dante network is a **collection of devices that are connected via IP** and can exchange audio information with each other
- Dante does much more than this
  - In order to make a modern network useable, devices are automatically discovered
  - Devices use human readable names – both for the device and its channel labels
  - Diagnostic information is made available from the devices to controller applications
  - Devices and channels can be managed by a systems administrator and a secure ecosystem can be built
- Dante devices exchange information over an **IP – Internet Protocol Network**
  - In all other areas of life the term “network” has become synonymous with IP network
  - The biggest single network in the world is the Public Internet – unsurprisingly this too uses IP

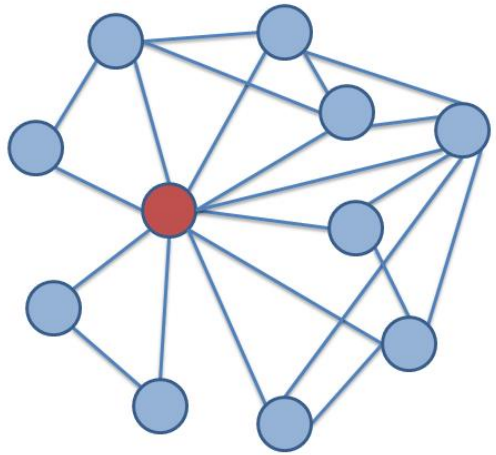

SO…DOES “NETWORKING” MEAN IP?

- IP networks are the most ubiquitous kinds of networks in the world today
- Why has IP become the dominant networking technology?
  - Supported by the most manufacturers (both in terms of software and hardware)
  - Has been applied to all industries (from Banking to Space Rockets to huge science projects)
  - IP is hugely flexible
  - The network is “neutral”
  - IP networks reach true global scale
  - Equipment and cabling takes advantage of “economies of scale” in an unprecedented way
  - The model IP follows is vastly more scalable than alternatives
  - Understood by millions of engineers and technicians world wide
  - Used by billions of end users every hour of every day

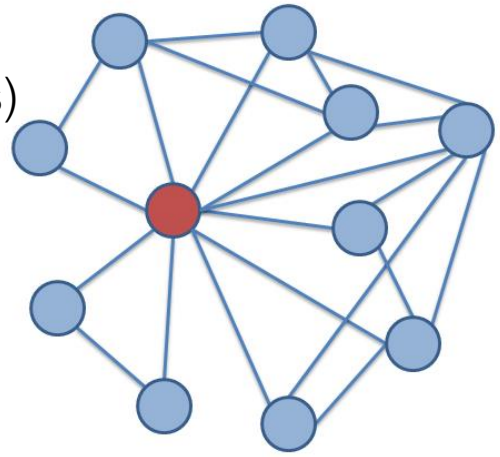

Dante by Audinate

ALTERNATIVES TO IP (HISTORY)

- Before IP networks became popular in telecommunications, **TDM – Time Division Multiplex** networks were common

  - TDM is the method that older digital audio formats used (AES3, AES10 (MADI) AES50)

  - A TDM network runs at a set rate – it is considered a “synchronous” network

  -“Timeslots” are filled with data (or not) and as long as the transmitting end puts the “right” data in the “correct” timeslot (as far as the receiver is concerned) then the outgoing resultant data matches that which went in

  - The capacity of the network depends upon the clock frequency, and by extrapolation the number of “timeslots” that are implied by this

ALTERNATIVES TO IP (HISTORY)

- TDM networks are not “neutral”
  - In telecommunications, as long as traffic was only voice, this was simple
  - Once Telcos started to provide other services, the TDM model showed its inefficiencies

COMPARE THIS WITH TDM

TDM Systems are only efficient at 100% useful capacity

Number of Channels 24bit	“Boring” bits Dante unicast	“Boring” bits 64 channel MADI	Bandwidth % on 100mbps link MADI	Bandwidth % on 100mbps link Dante
1	104	1512	100%	6%
2	80	1488	100%	6%
4	32	1440	100%	6%
64	128	0	100%	88%
65	104	1512*	200%	94%

WHY ARE "PACKET SWITCHED" NETWORKS USED ON COMPUTERS?

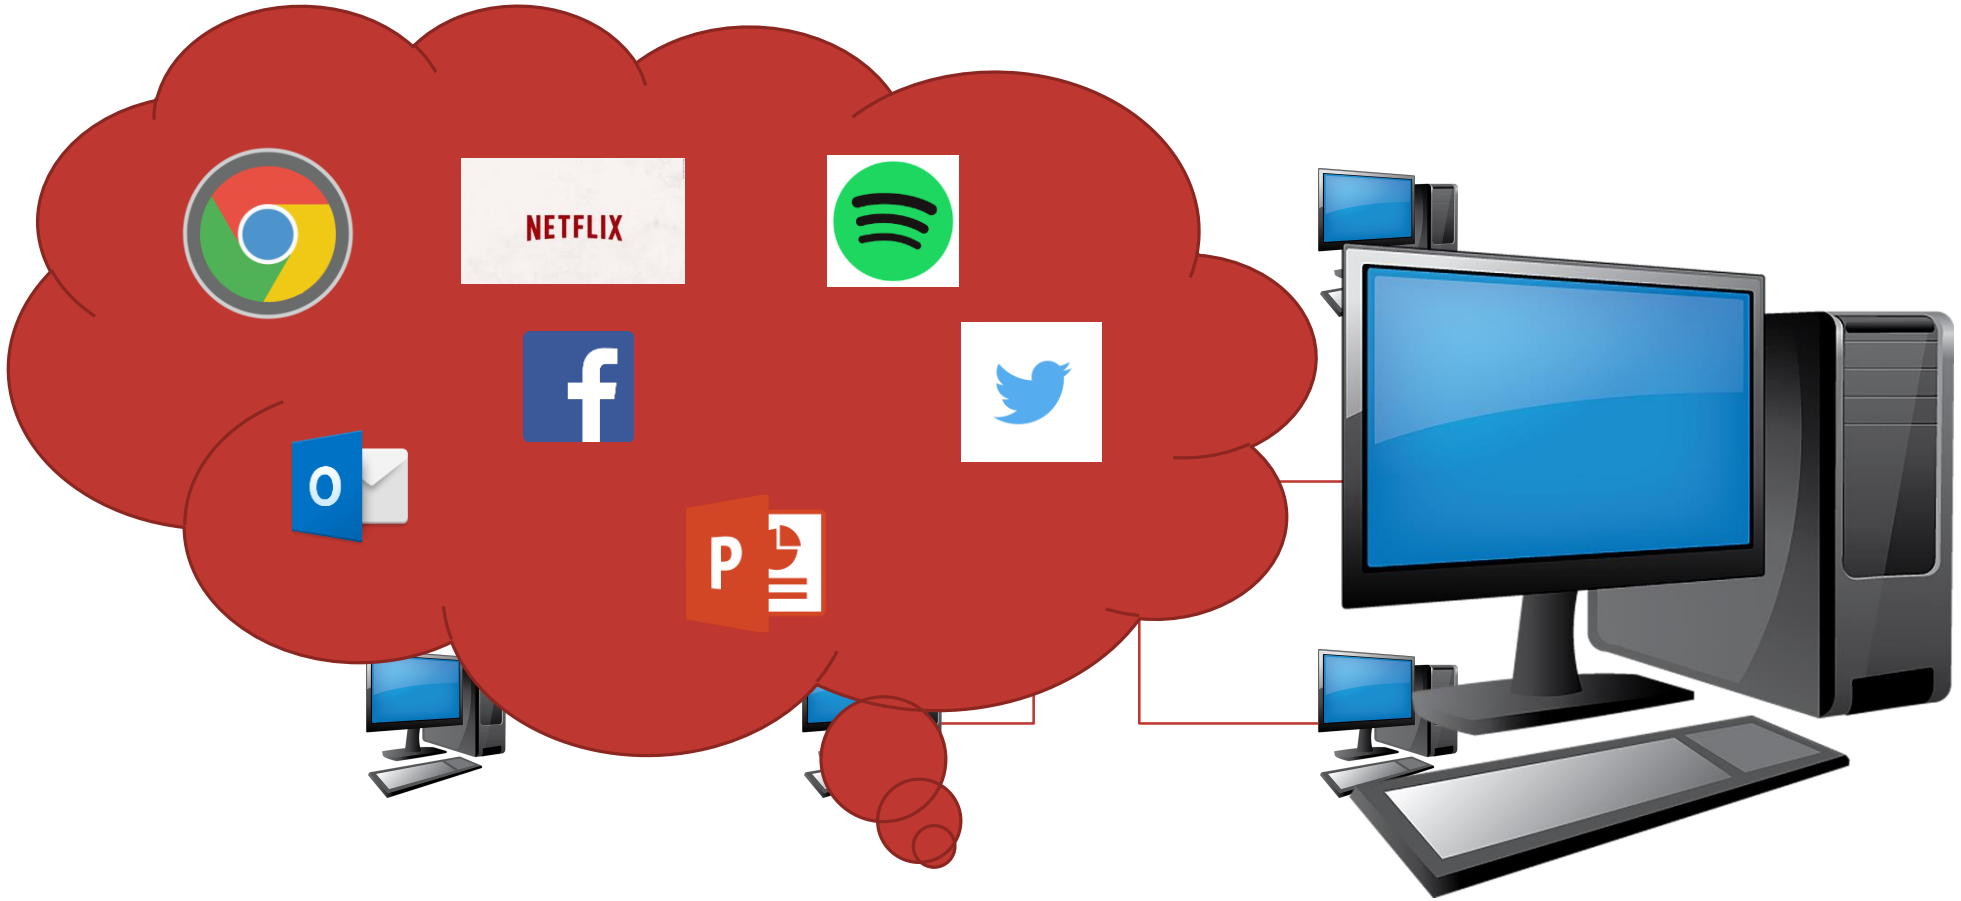

APPLICATION ADDRESSES

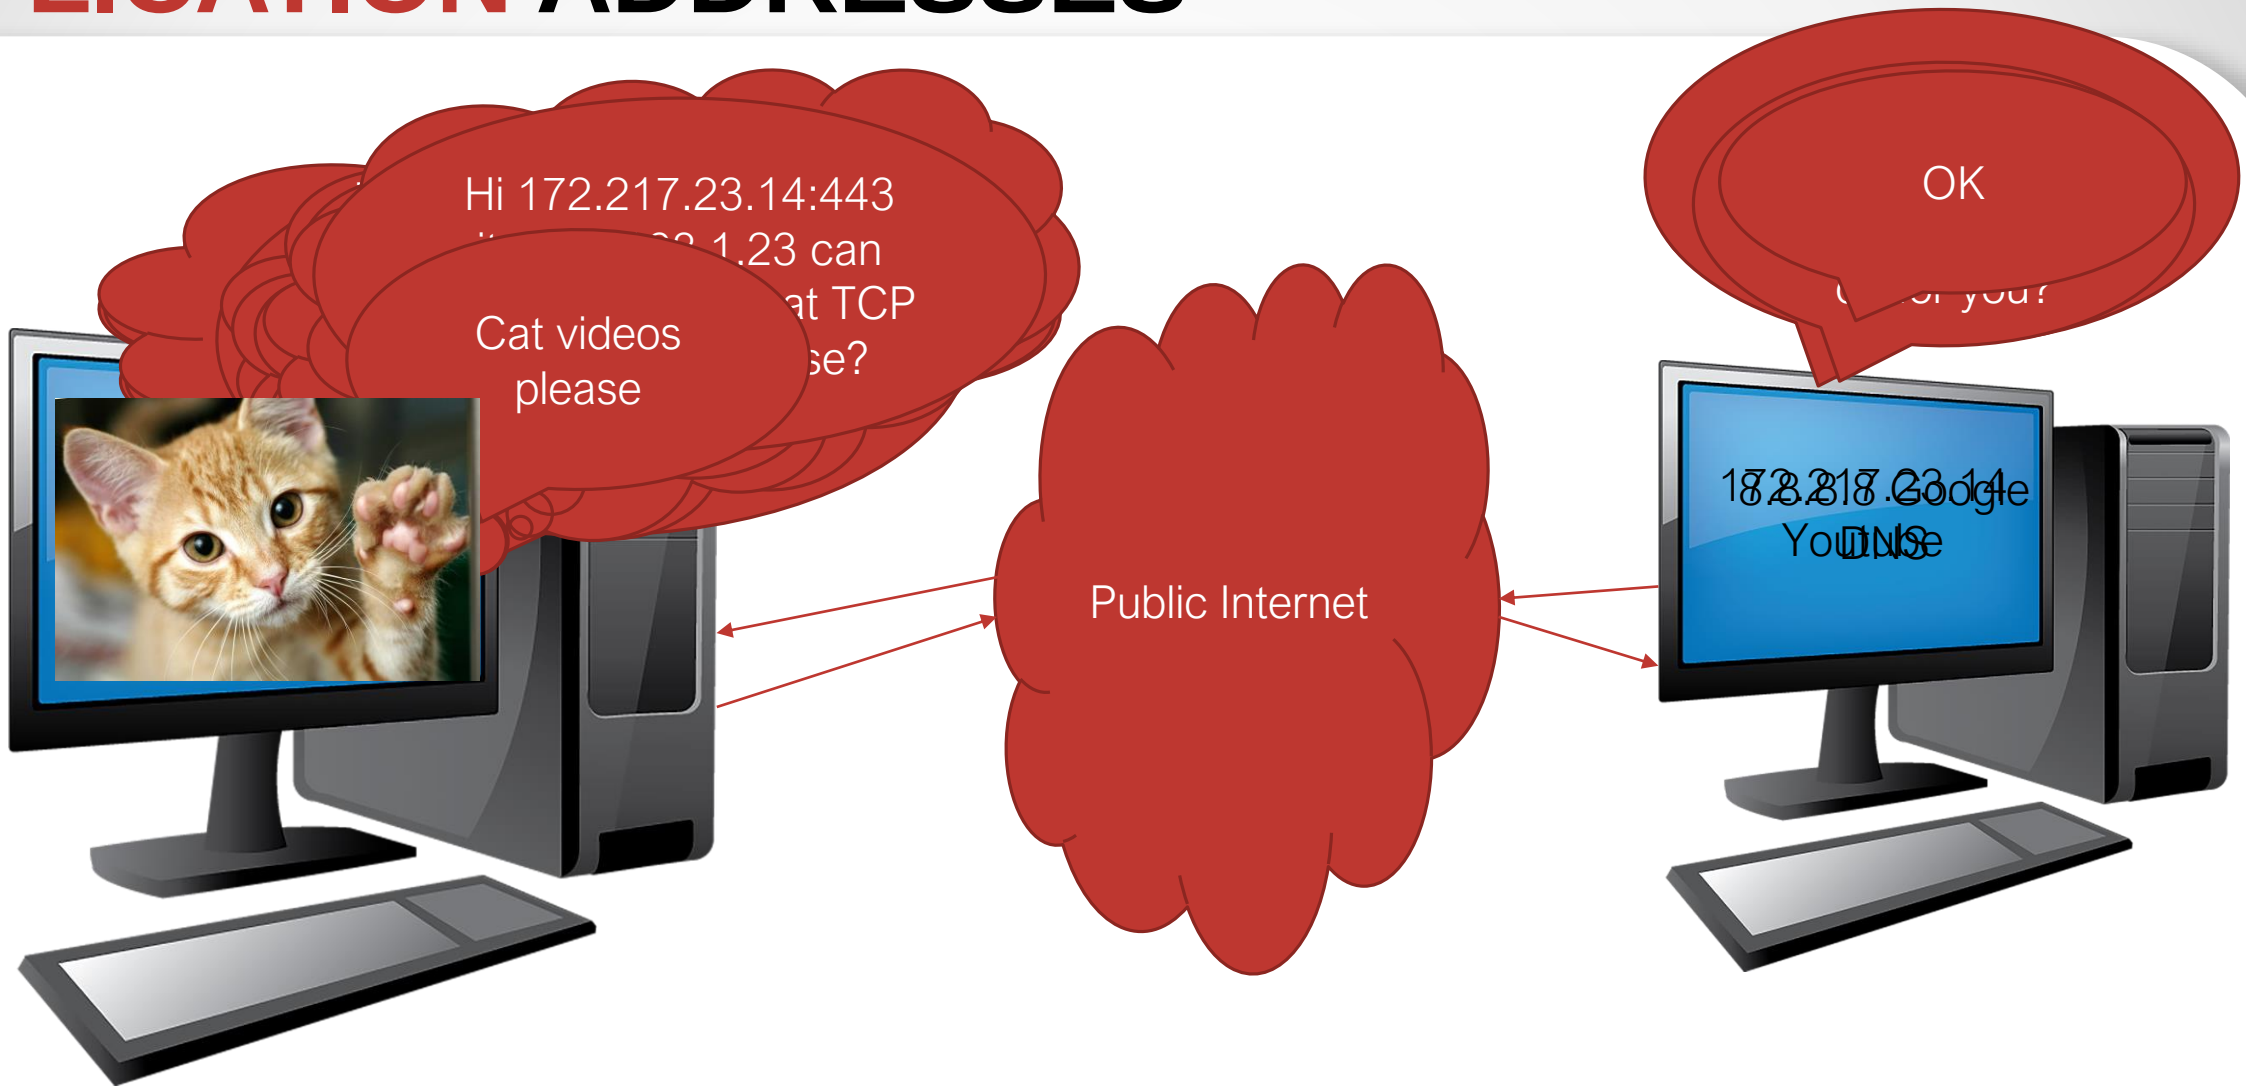

APPLICATION ADDRESSES

- The same process repeats for every application
- Each application has its own unique Internal (port) address
- In an Audio device this could be:

Application	Local Port	Remote IP	Remote Port
PTP	UDP	224.0.1.129	UDP 319
Audio Flow	UDP 14340	192.168.1.56	UDP 14390
Audio Flow	UDP 14350	192.168.1.60	UDP 14367
Gain control	UDP 50135	192.168.1.56	UDP 50231

HOW DO WE PUT IT TOGETHER… BUT KEEP IT SEPARATE?

IP Network Address Book

VS

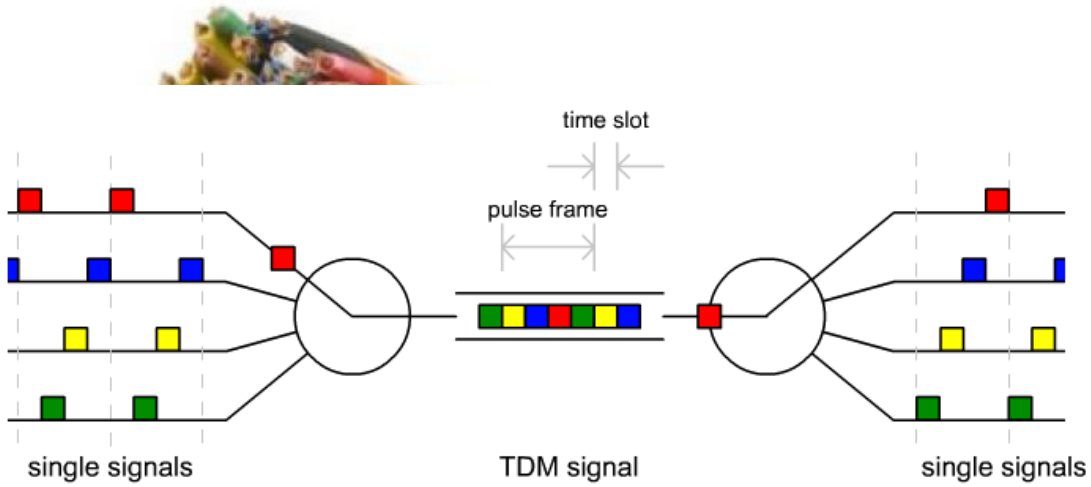

Dedicated TDM “Pairs”

ADVANCED DANTE NETWORKING: SECTION 2

IN THIS SECTION…

Layered model of a network

- **Encapsulation**
  - How we organize data to be sent
  - How we create address centric circuits all the way to the application
  - Understanding Network Layers – the postal model

LAYERED MODELS

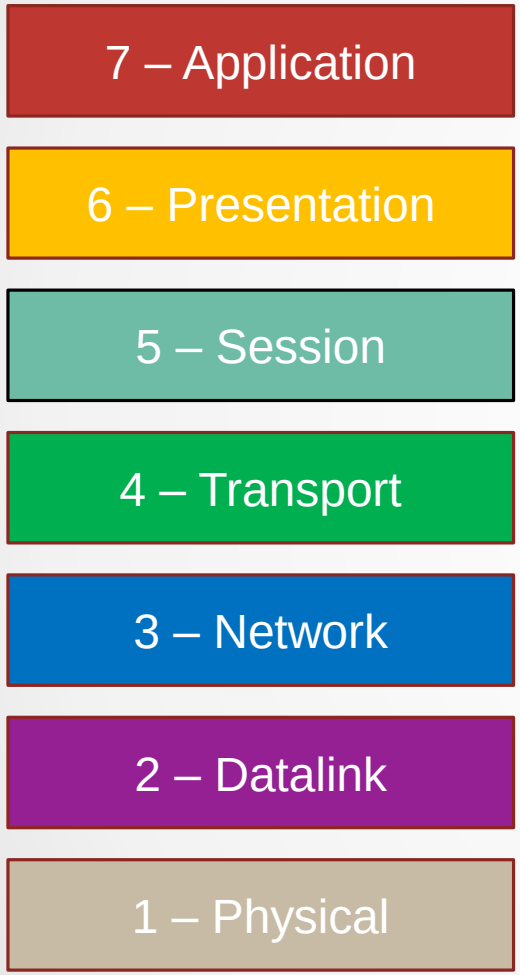

OSI Model

- Older
- Good at defining This part
- Less relevant here

TCP/IP Model

- Newer
- Good at defining This part
- Lacks detail here

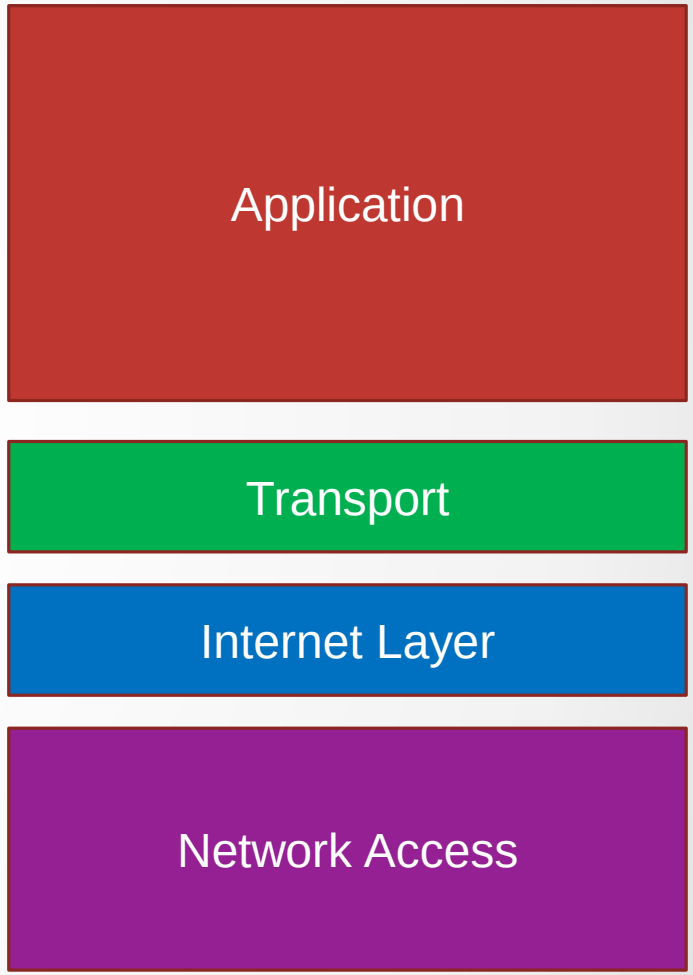

ENCAPSULATION – THE "NETWORK STACK"

- The OSI model is huge and contains scope for building many different communications methods
- At a practical level the available toolkit this implies is very large
- A simplified (and accurate) model of what happens in the network is as follows:

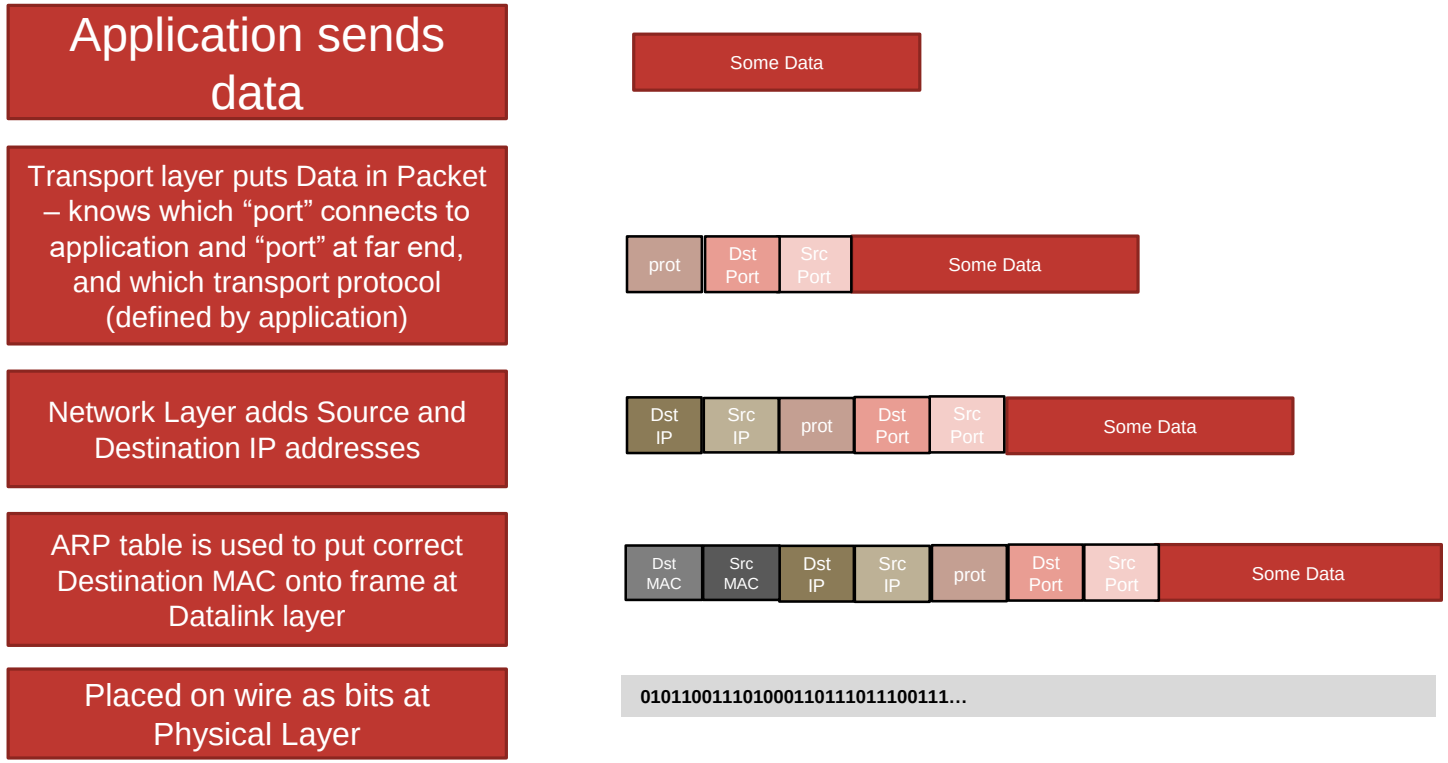

ENCAPSULATION – THE "NETWORK STACK"

Application sends data

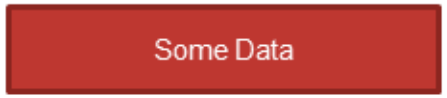

Transport layer puts Data in Packet – knows which "port" connects to application and "port" at far end, and which transport protocol (defined by application)

Network Layer adds Source and Destination IP addresses

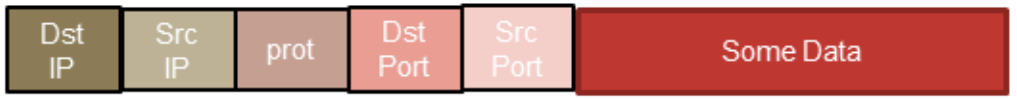

ARP table is used to put correct Destination MAC onto frame at Datalink layer

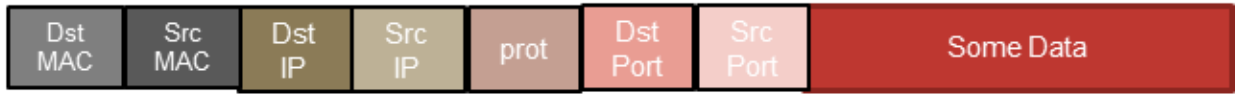

Placed on wire as bits at Physical Layer

ENCAPSULATION – AN ANALOGY

IP stacks are abstract

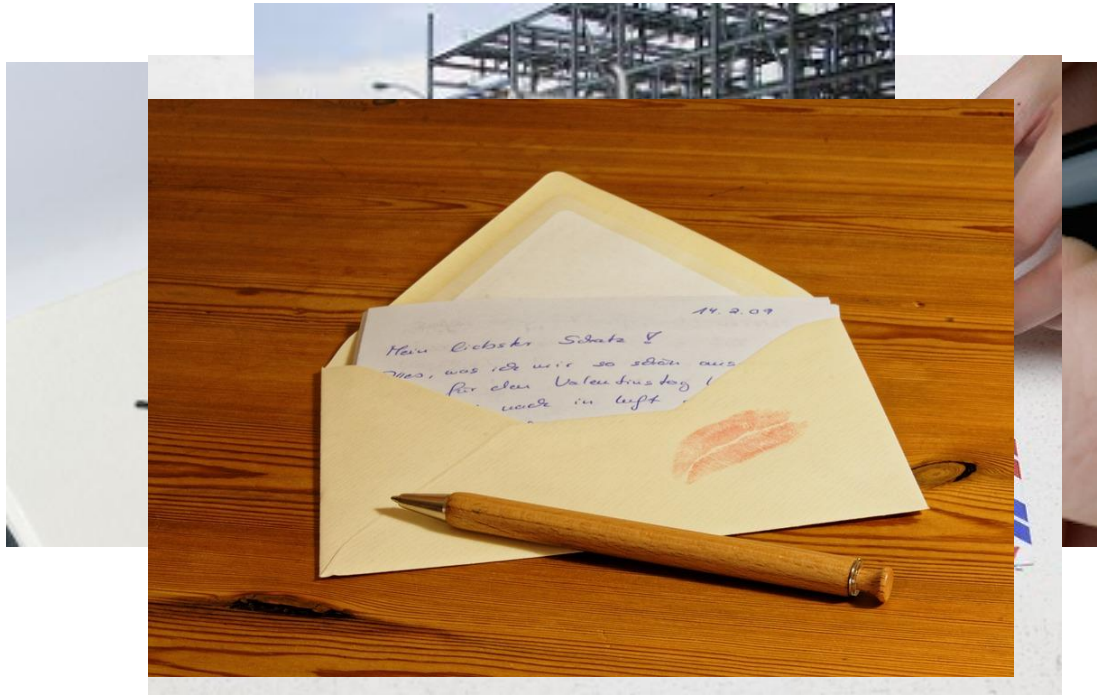

1. I create some data (write a letter)
2. I address the envelope
3. I put the letter in the envelope (packetization)
4. At the post office the envelope goes in a mailbag (Frame (MAC) encapsulation)
5. Mailbag is put in truck and taken to next sort facility (physical layer)

ADVANCED DANTE NETWORKING: SECTION 3

IN THIS SECTION…

Different modes of communication

- Broadcast
- Unicast

Where they become applicable

- Collision Domains
- Broadcast Domains

Routers and Switches

- Segmenting the collision domain - switching
- Segmenting the Broadcast domain
- ARP
- VLANs
- IP Subnets
- Routers

TWO (OR 3) KINDS OF MESSAGES

- "Addressed" messages give huge flexibility advantages

- In a TDM network there is no requirement to use any form of address
  - Destinations are interpreted by their timeslot position nothing more

- In a Packet Switched network (just like the mail system) the addressee is unique
  - It is also possible to send "unsolicited" mail – from a sender to many or all addresses
  - This is both very useful, and potentially irritating/bad at the same time

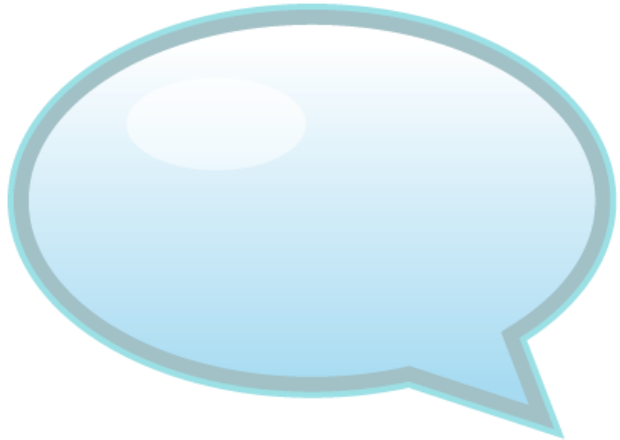

TWO (OR 3) KINDS OF MESSAGES

- IP networks are very good at dealing with limiting “unsolicited mail” through management techniques

- The three types of messages in an IP network are:
  - Unicast – one to one communication
  - Multicast – one to many communication
  - Broadcast – one to all communication

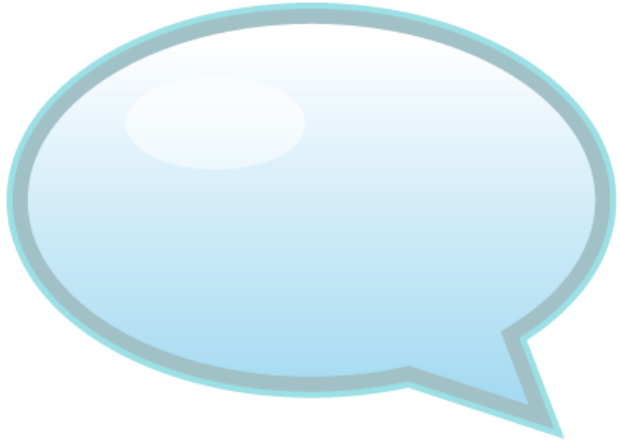

- By contrast a TDM network is only really unicast (but through “wiring tricks” can be analogous to Broadcast) – either way there is no strict distinction

HOW DIFFERENT MESSAGE TYPES WORK

- Old IP networks used “hubs” or a single cable to connect computers together
  - “layer 1” based network
  - all messages are effectively “broadcast” messages within the network
  - we can very precisely address a message
  - Using a hub or a single “token ring” connection, we effectively remove the “sorting office” from our postal system
  - All computers receive every message, and have to discard anything they are not interested in themselves

- Packet switched networks use a specific IP address for IP broadcast messages,
  - This maps to a specific broadcast destination MAC address
- In order for this to be useful we need something in the network to “sort” the mail

COLLISION DOMAINS

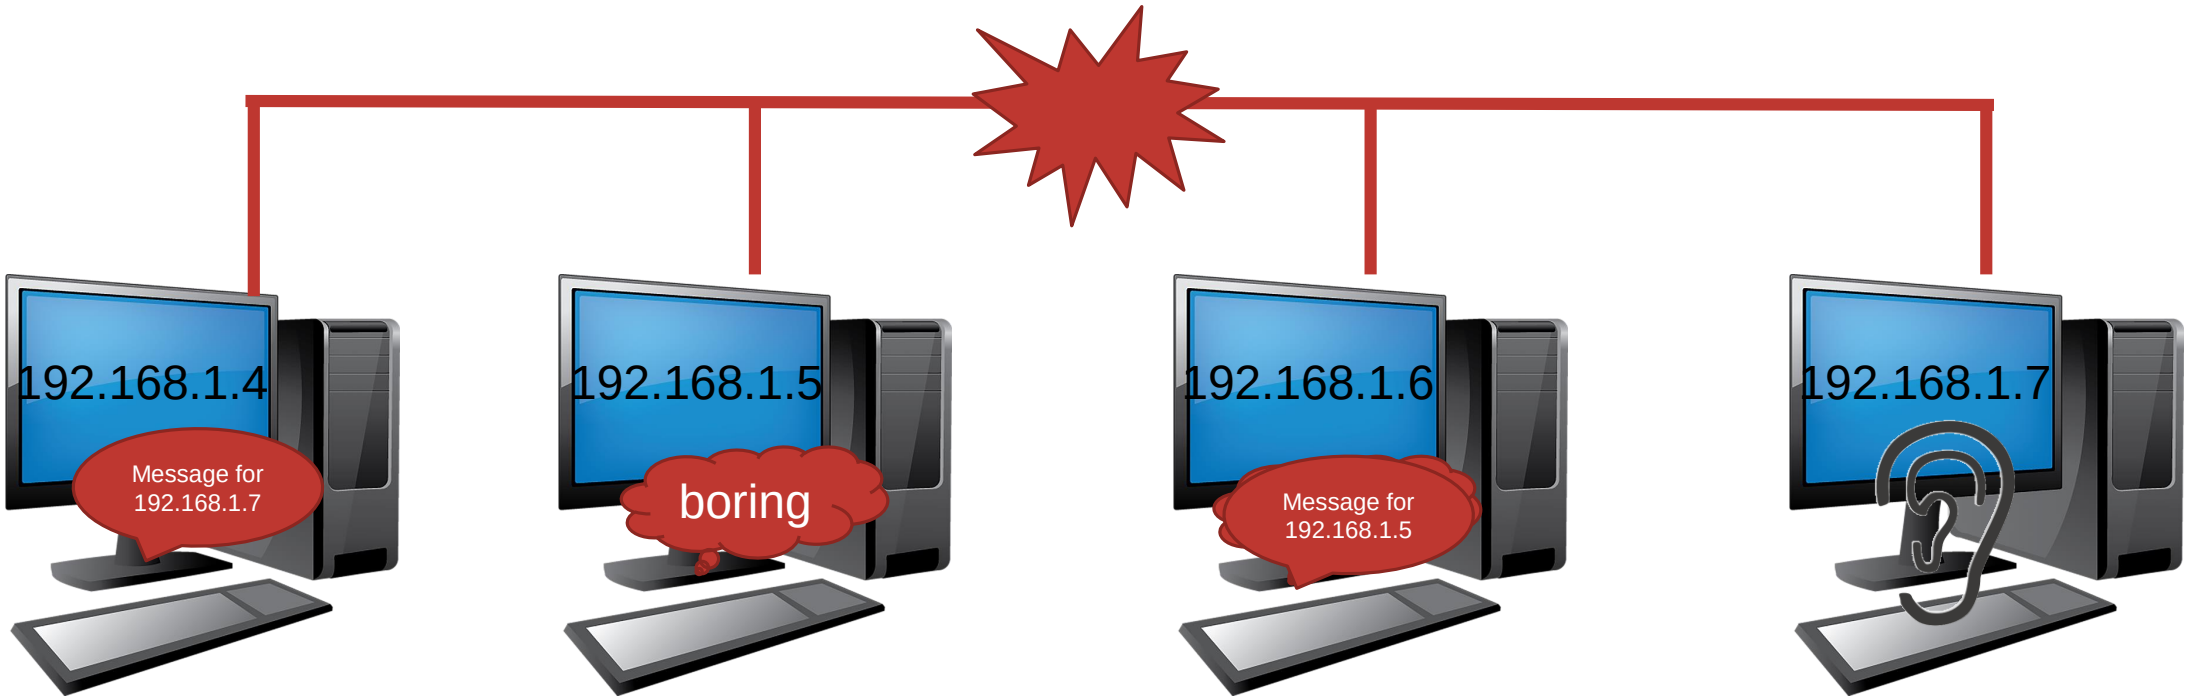

COLLISION DOMAINS - SEGMENTING

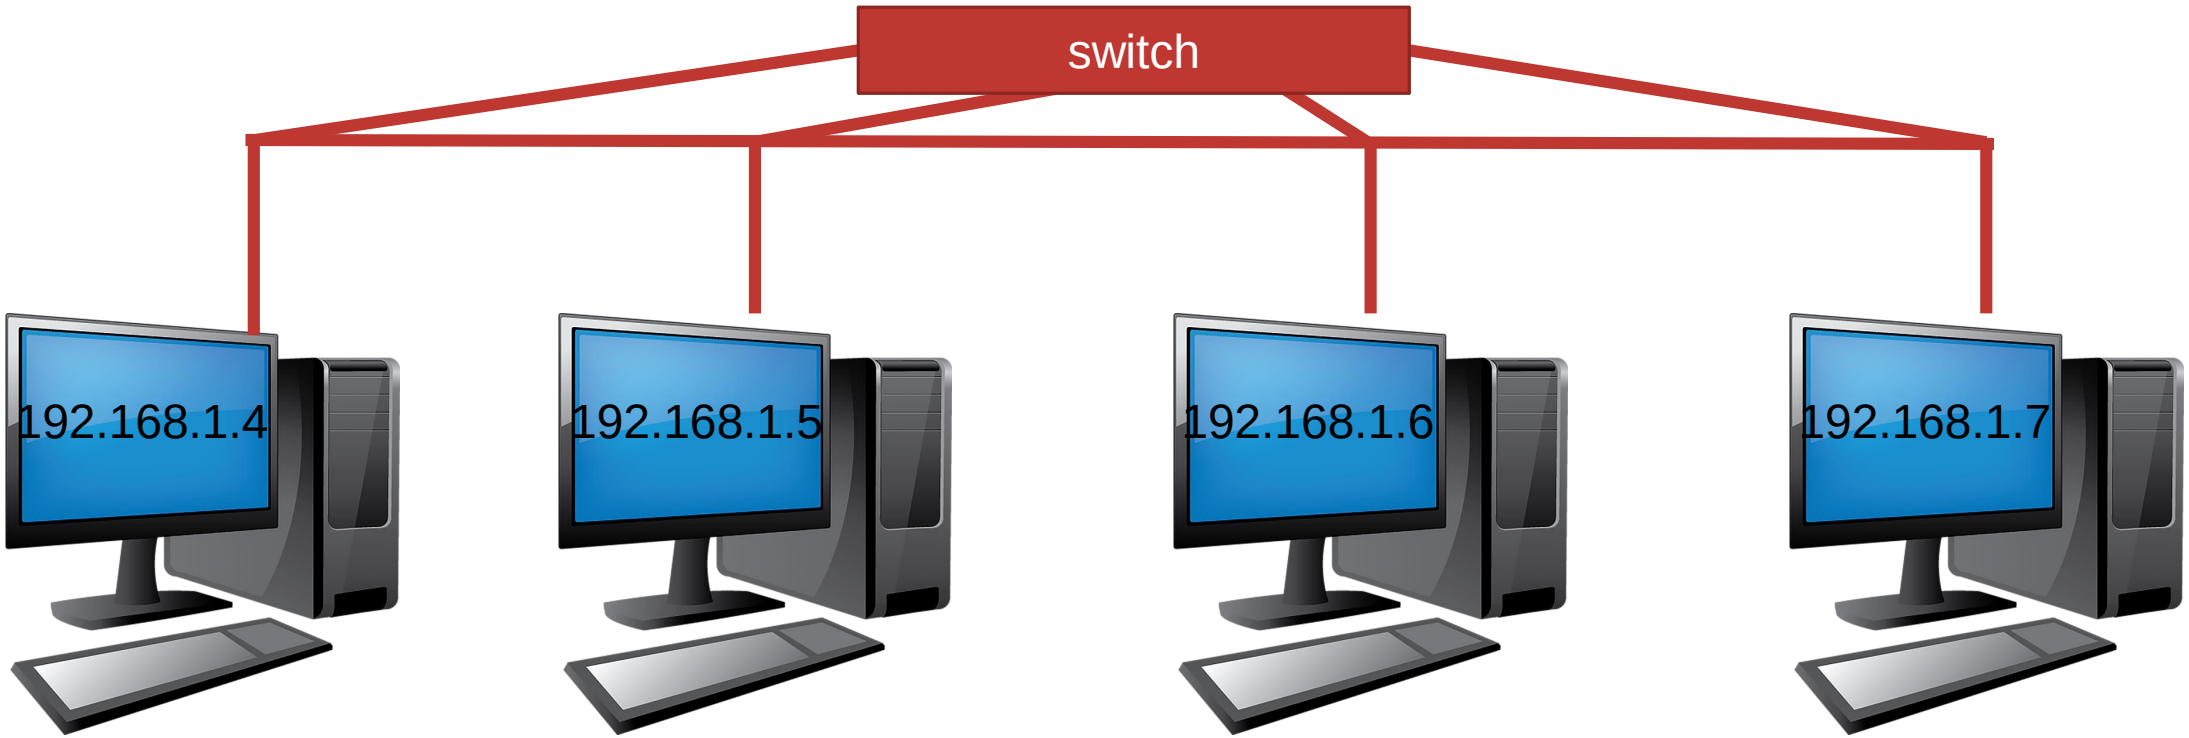

UNICAST TRANSMISSION

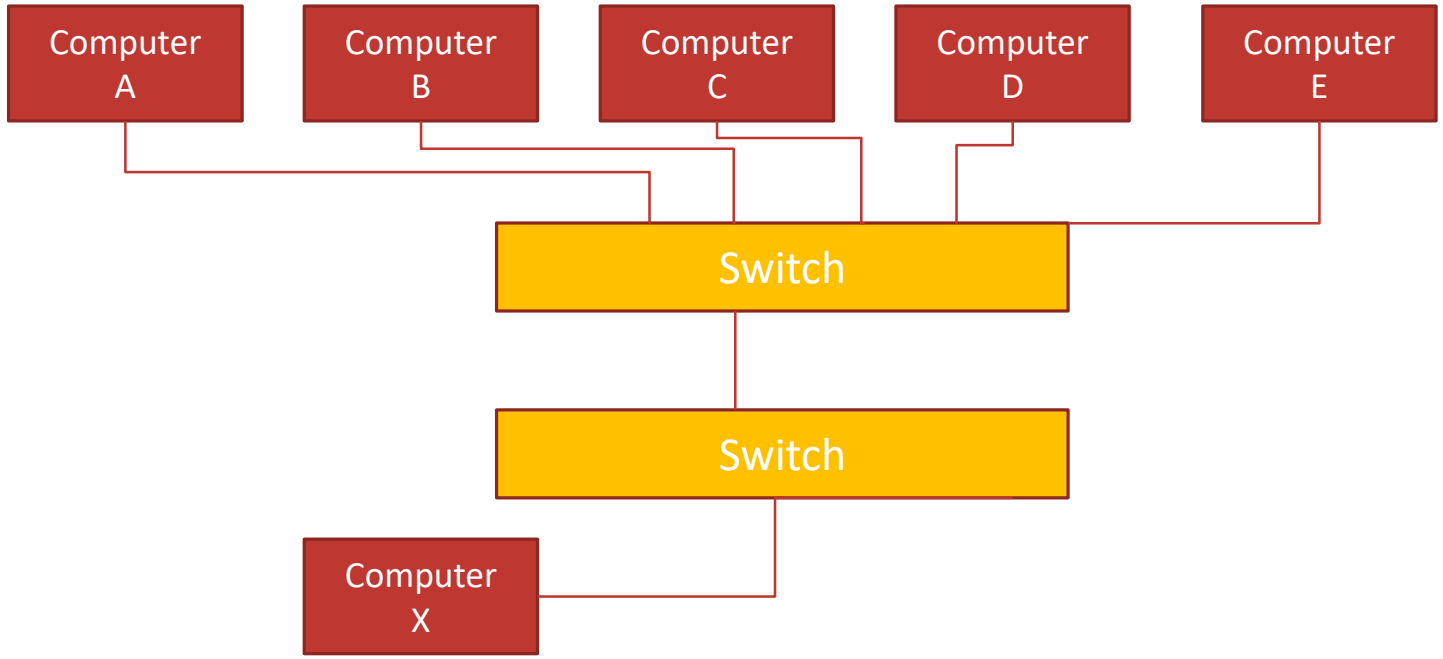

1. Network is connected as shown
2. Computer A sends message to Computer E – session starts
3. Session completes
4. Computer B sends message to Computer D – session starts
5. Computer C sends message to computer E – session starts
6. Session between Computer B and D completes
7. Network is expanded
8. Computer X sends message to computer A
9. Session between computer C and E completes
10. Session Between Computer X and A completes

“LAYER 2” NETWORKS

Switches are only concerned with MAC addresses

There are 3 types of MAC address

- Host (the MAC address at Layer 2 of the device’s IP stack)
- Multicast – a specific group of destination MAC addresses that tells a switch to send out of all interfaces (unless managed)
- Broadcast – a destination MAC address that tells a switch to forward out of all interfaces

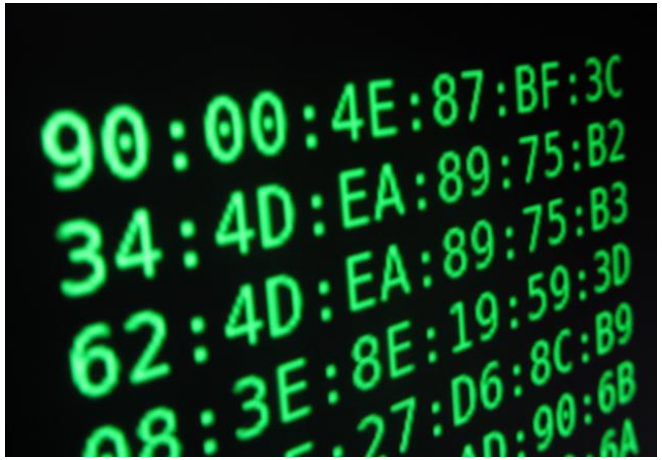

Some network technologies only used Layer 2

- MAC address to MAC address
- This network technique is better than just using a technique that cannot identify a device
- However this approach does not scale as well as a “full stack” implementation
- Operating systems require specific drivers to be written to handle the “missing part” of the stack

IP ADDRESSES IN A LAN? WHY?

- A “standard” network stack in an Operating System on a computer connects ports to the application software running on it – through the IP address

- It is assumed that a consistent method is used to communicate with a neighboring device and a device in the other side of the world – the “detail” is the preserve of the network infrastructure.

- **Important note:** it is as important to understand what sticks the “layers” of a network together as it is to understand the layers themselves

- For a device to understand how to communicate with other devices in the network we use a process called **ARP – Address Resolution Protocol** to resolve IP addresses to MAC addresses

SENDING AND RECEIVING UNICAST

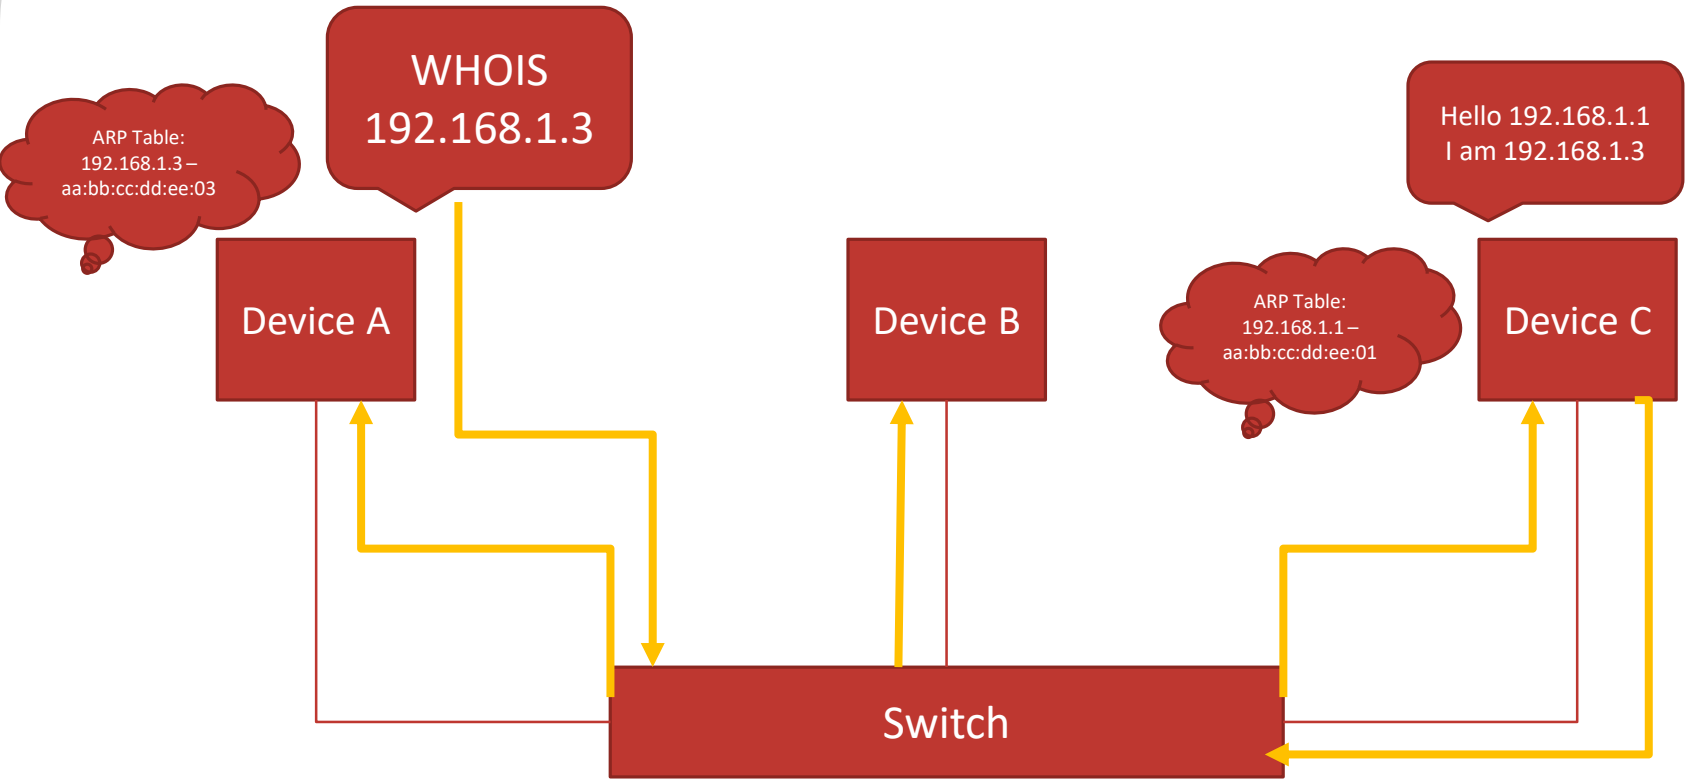

Interface	MAC Address
GE0/1	aa:bb:cc:dd:ee:01
GE0/2	
GE0/3	aa:bb:cc:dd:ee:03

1. An application in Device A decides it needs to send a message to Device C
2. Device A knows the IP address of device C (we are not concerned with how in this example)
3. The network has just been switched on – no communication has taken place yet
4. In order to ”find” device C, device A issues an “ARP” (Address Resolution Protocol) request (broadcast message)
5. This broadcast message goes everywhere
6. The ARP message has some standard useful information around it:
    1. Source MAC address
    2. Source IP address
7. The switch takes a note of the source MAC address, and populates its MAC address table
8. Device B receives and ignores the message
9. Device C receives, and responds to the message (unicast)
10. Device C knows destination MAC address from the frame containing the packet, from which it knows the sending IP address sent from Device A
11. Device C populates its ARP table
12. Switch populates MAC address table from seeing the returning message

BROADCAST DOMAINS

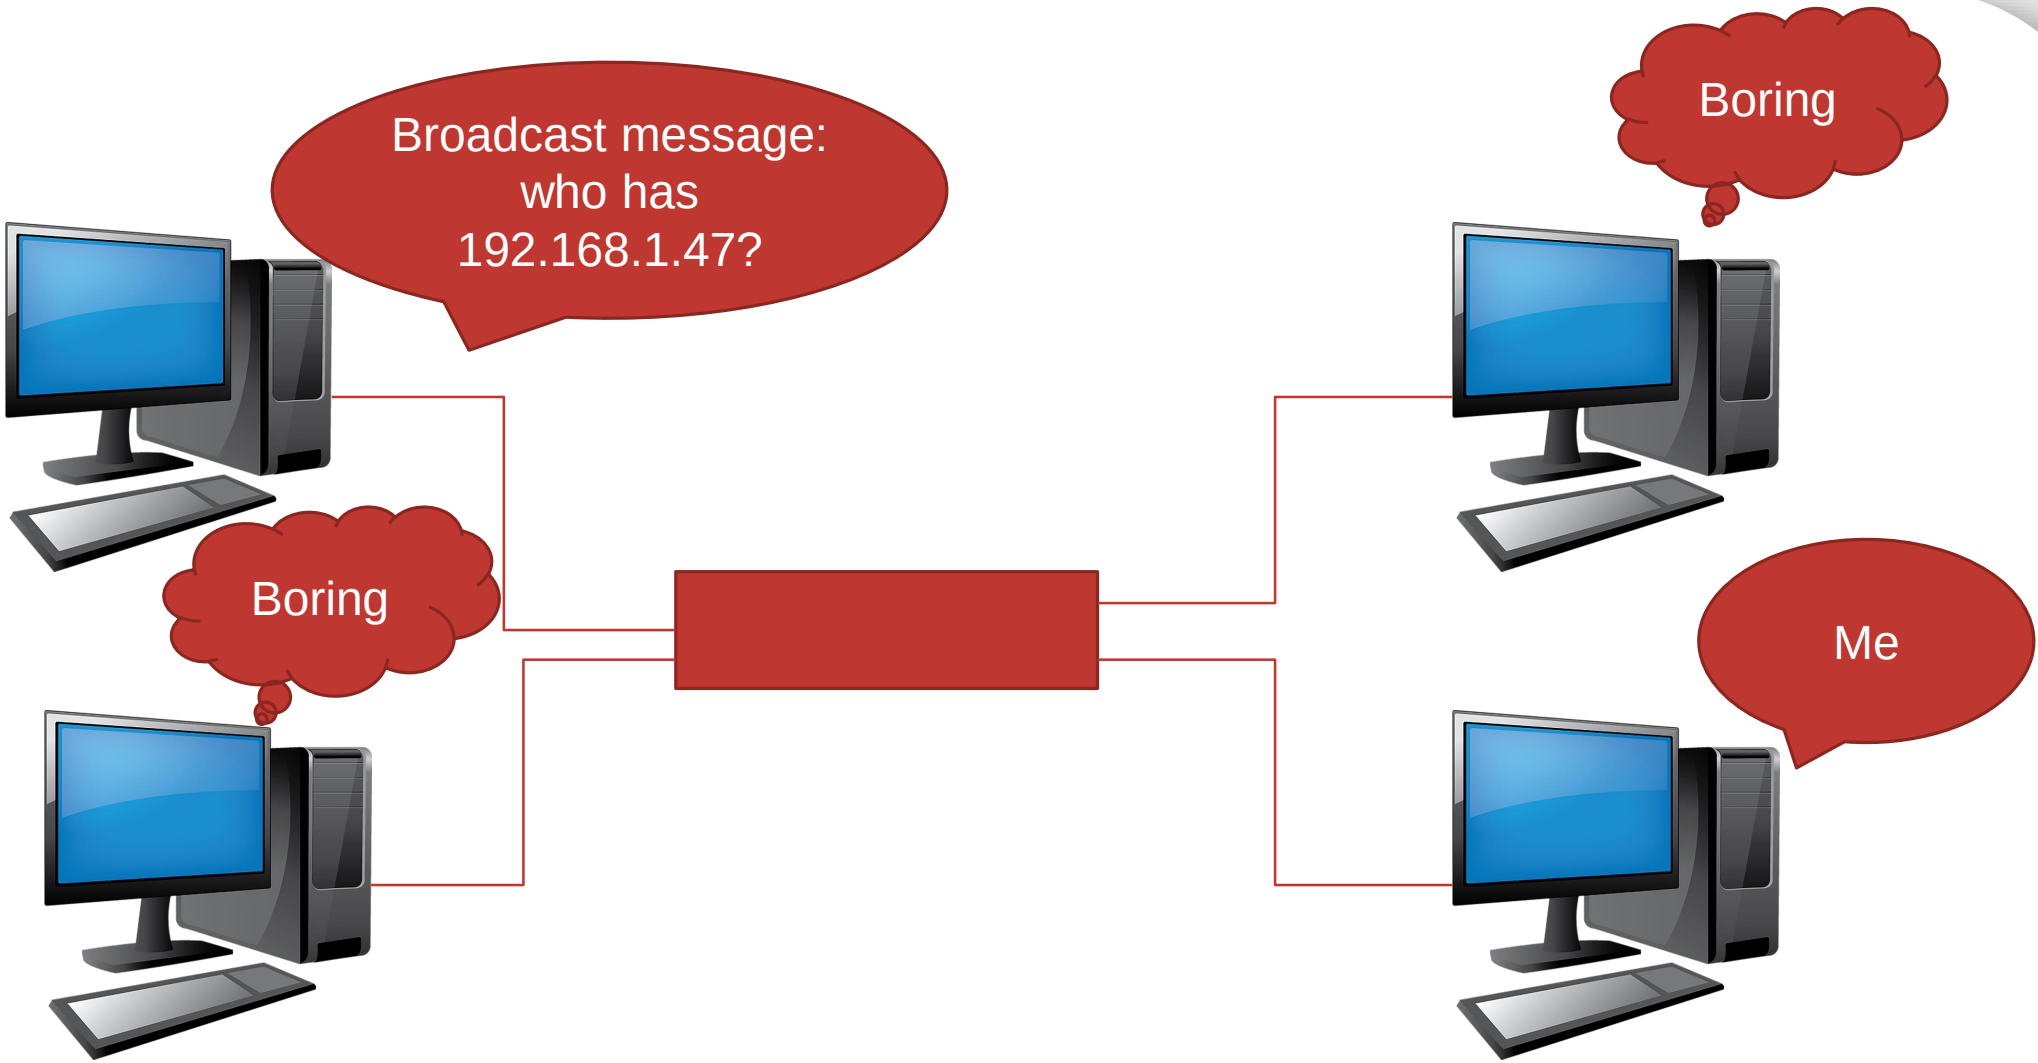

BROADCAST DOMAINS

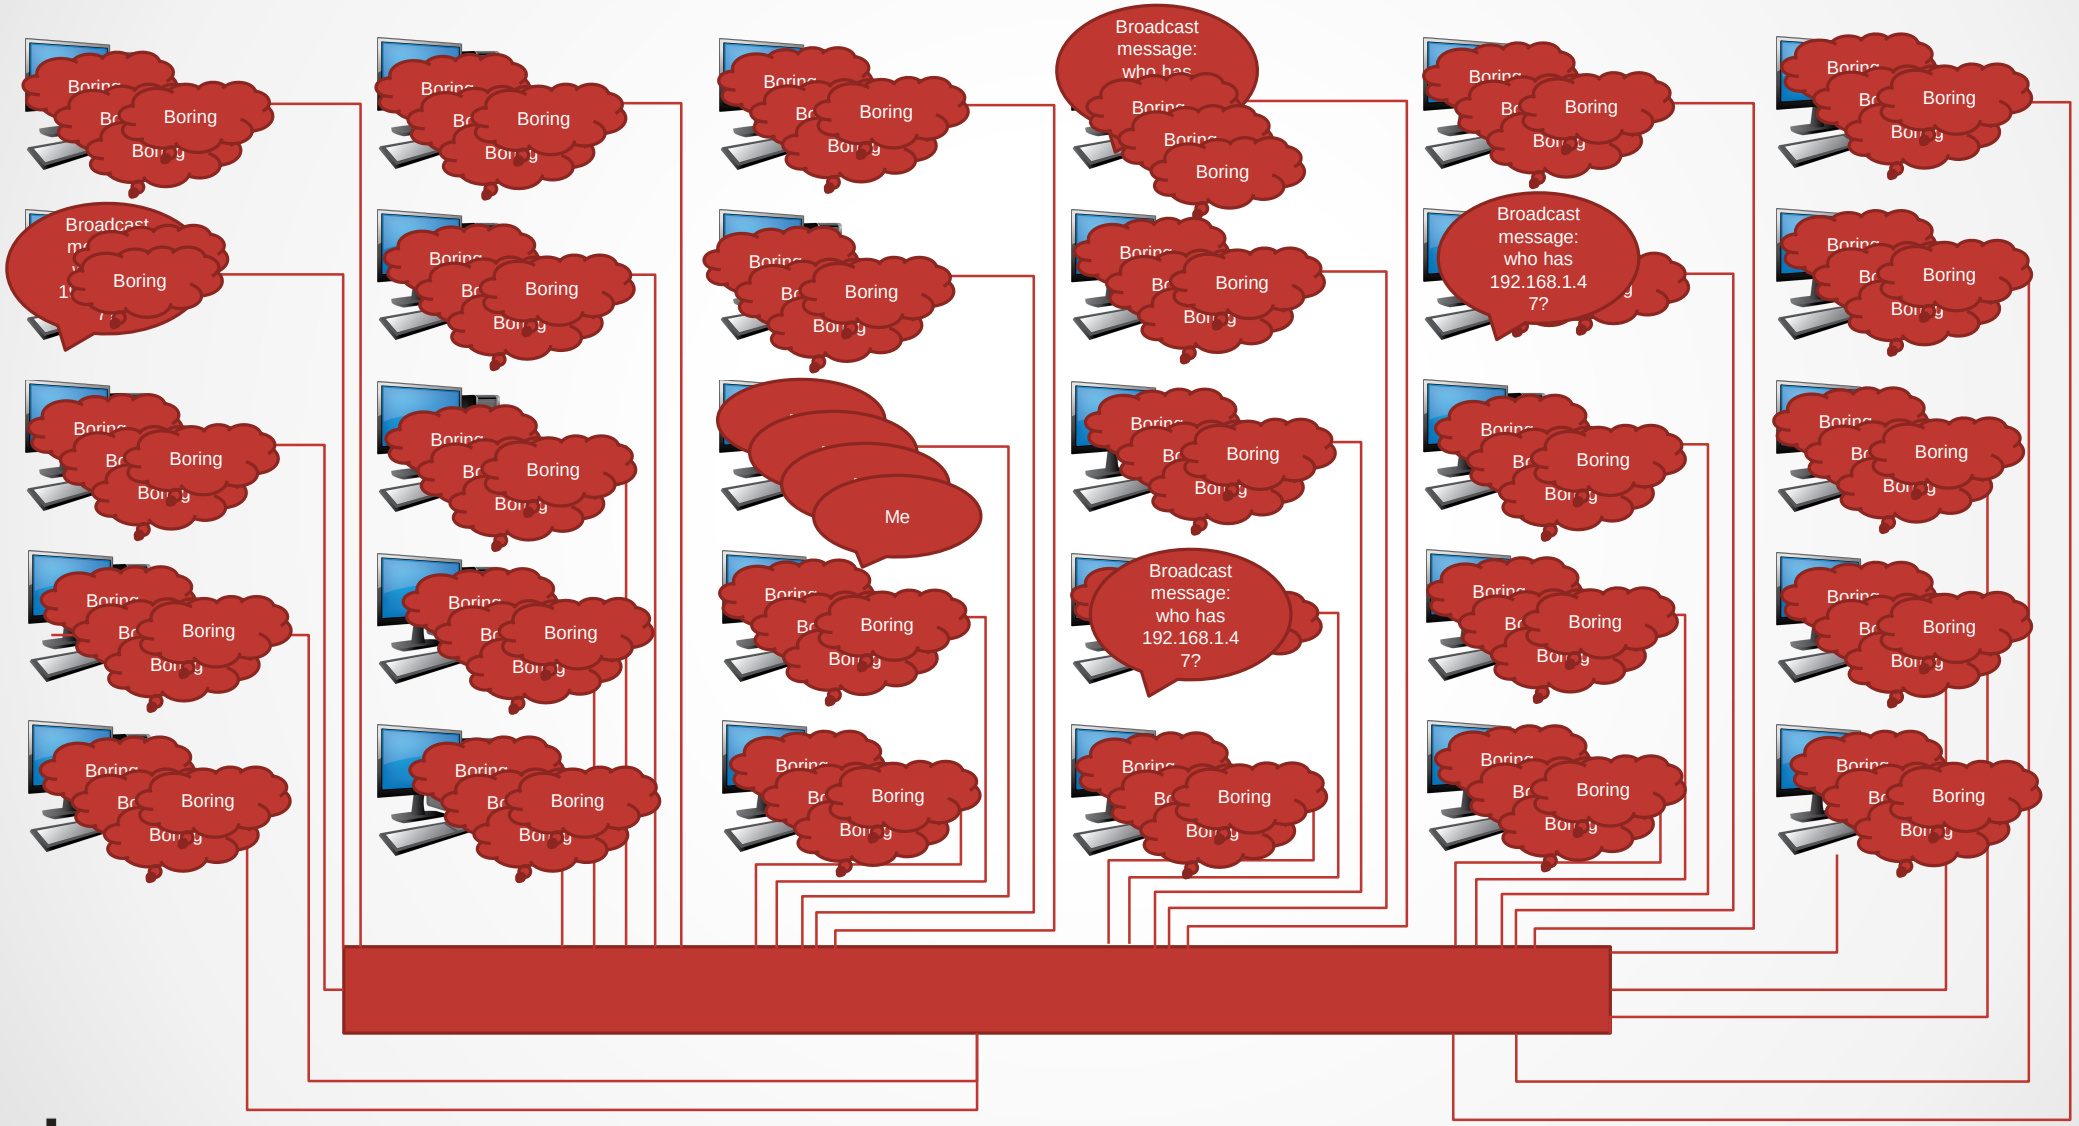

Surely there is a better way to deal with this?

SEGMENTING BROADCAST DOMAINS – GOOD PRACTICE

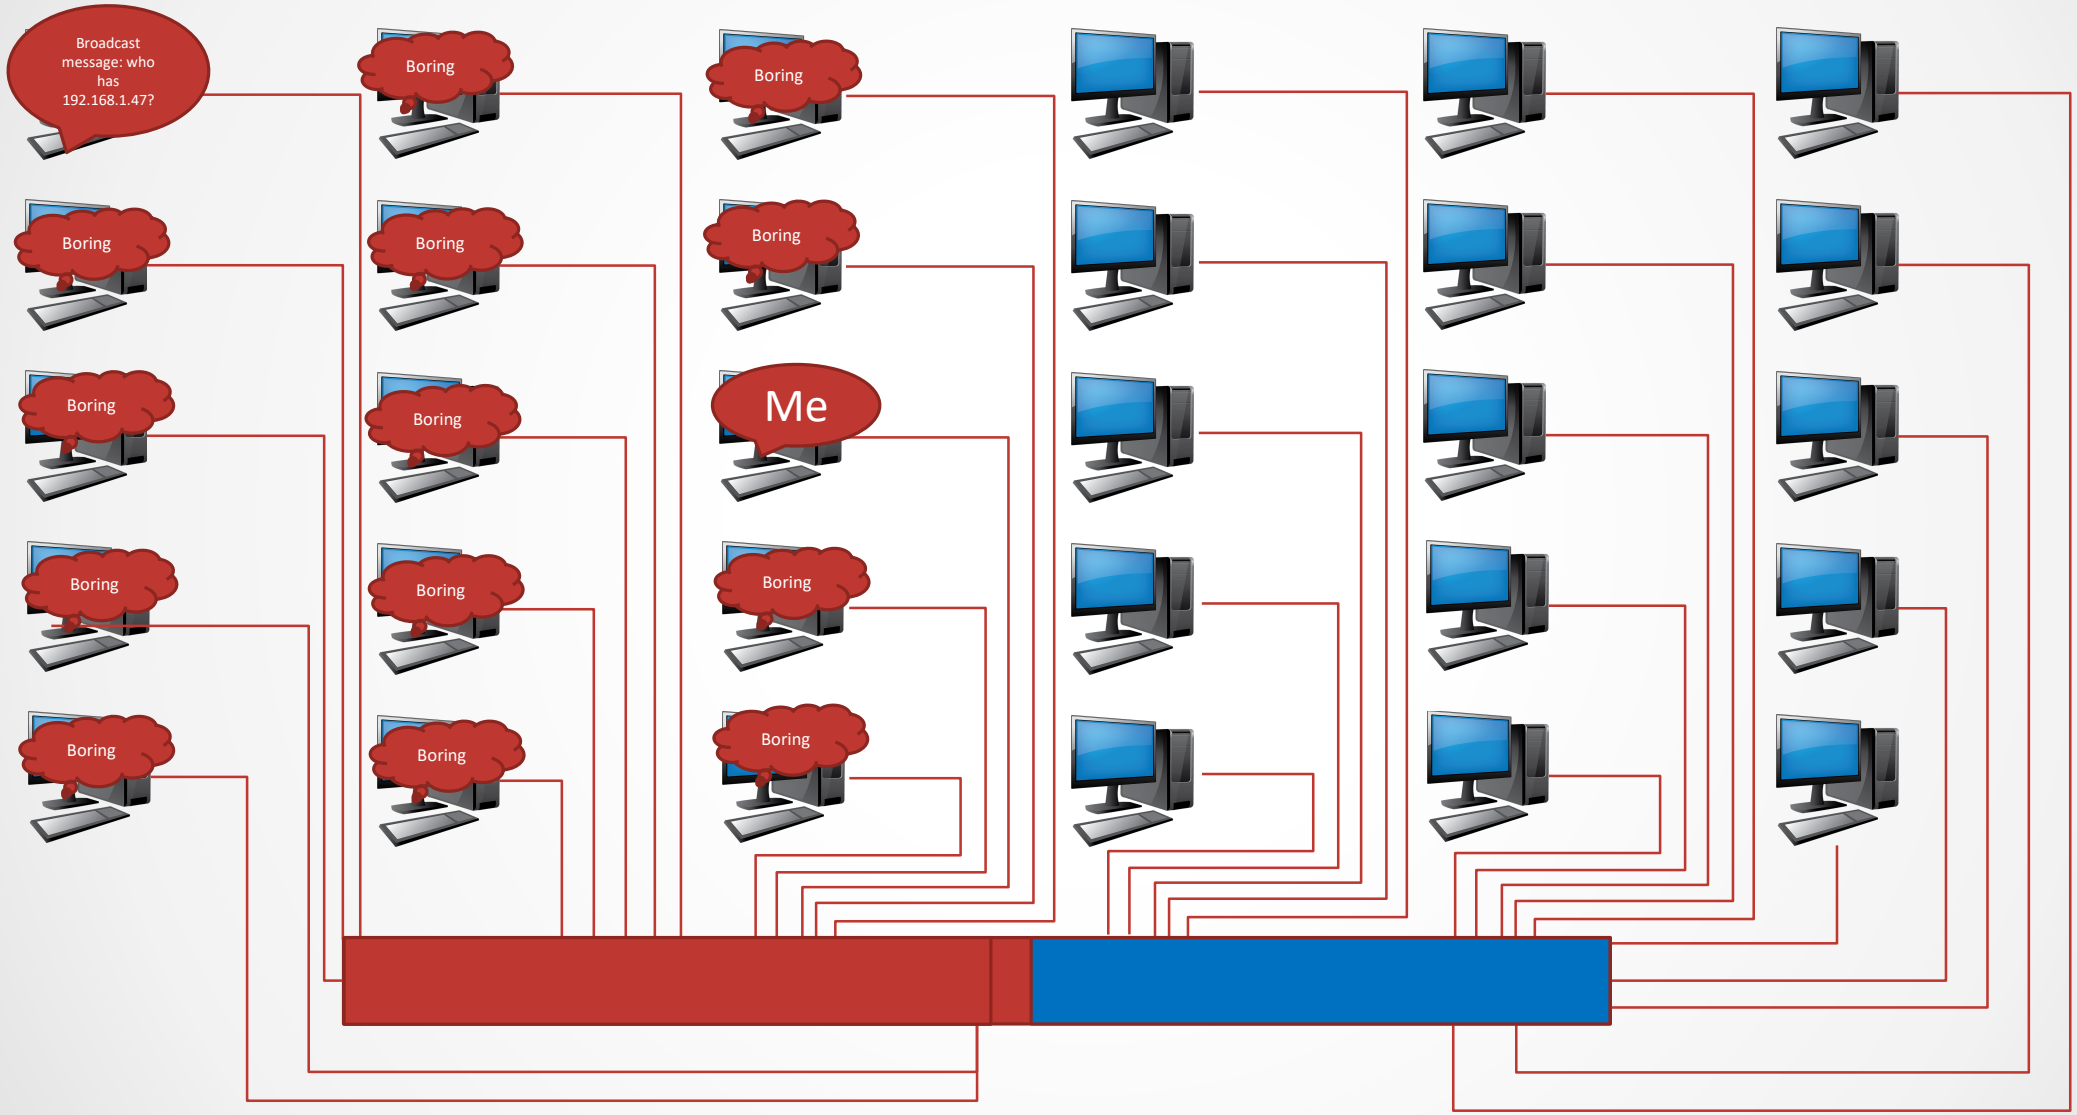

SEGMENTING BROADCAST DOMAINS – BUT STAYING CONNECTED

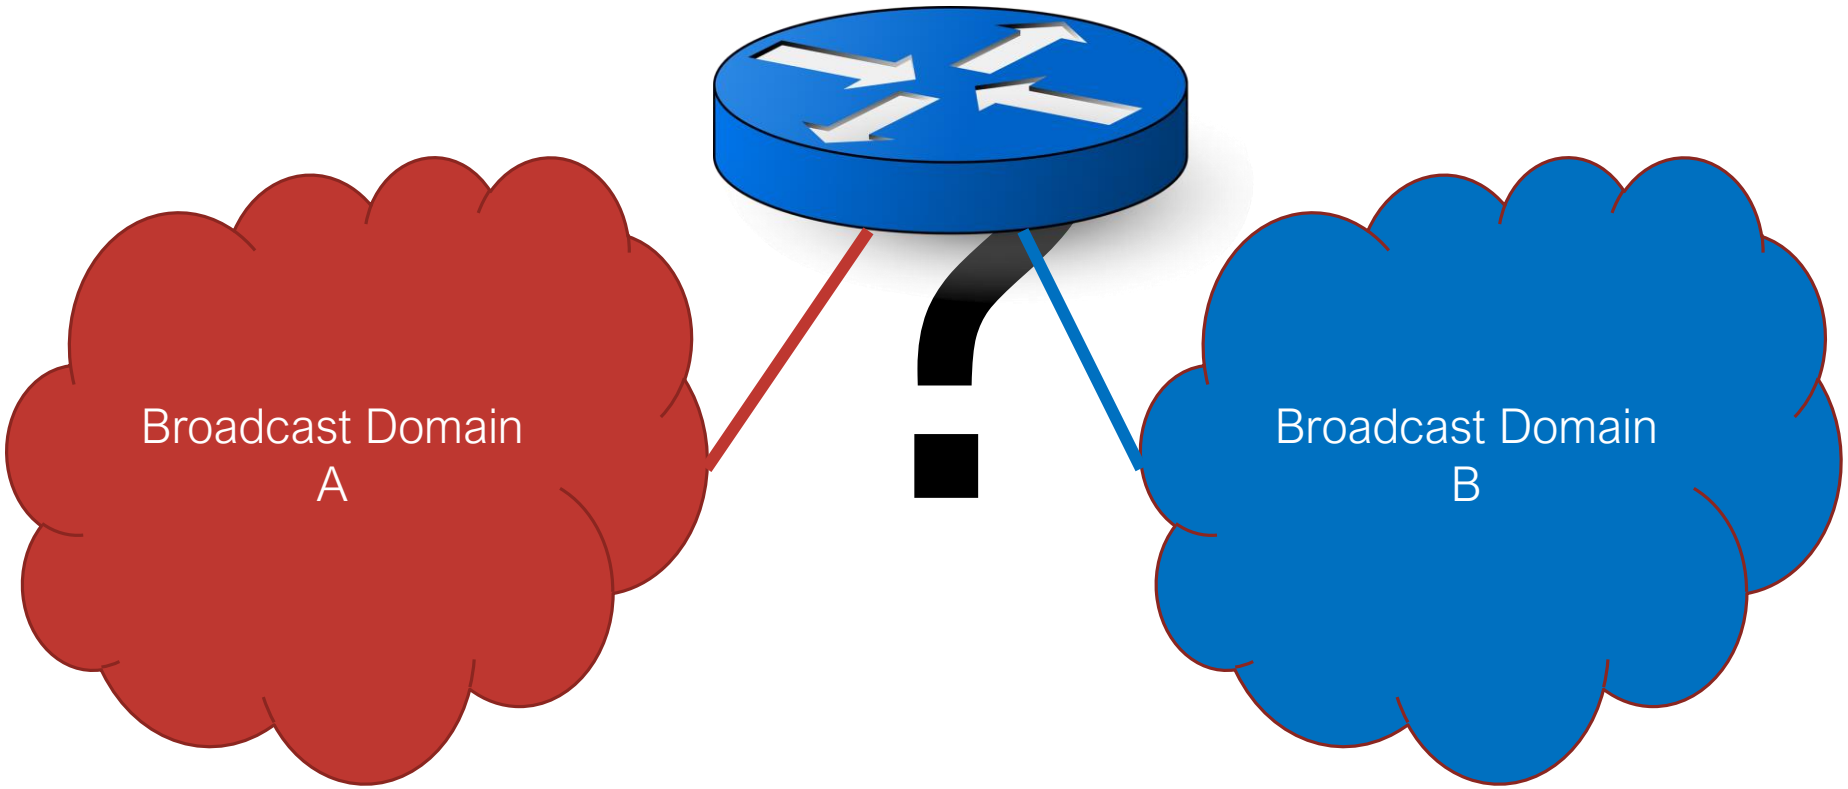

VLANS - HISTORY

Originally when Network managers segmented networks without VLANs a different switch would be used for each “group” of computers

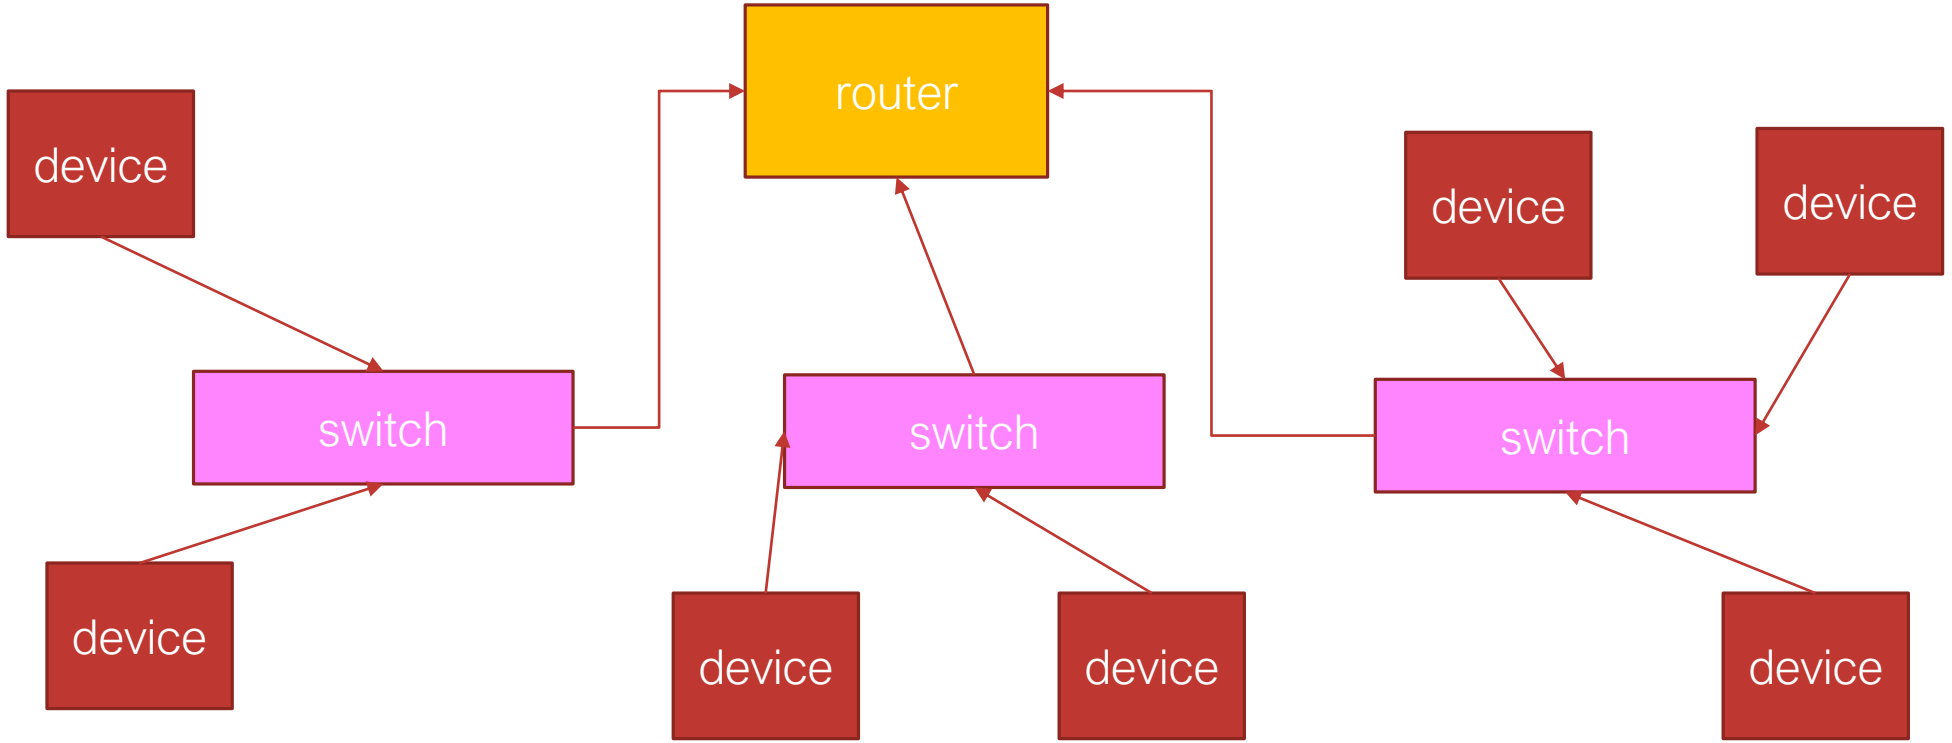

VLANS - HISTORY

This led to “empty” ports on switches, and lots of extra wiring

Empty ports = wasted money!

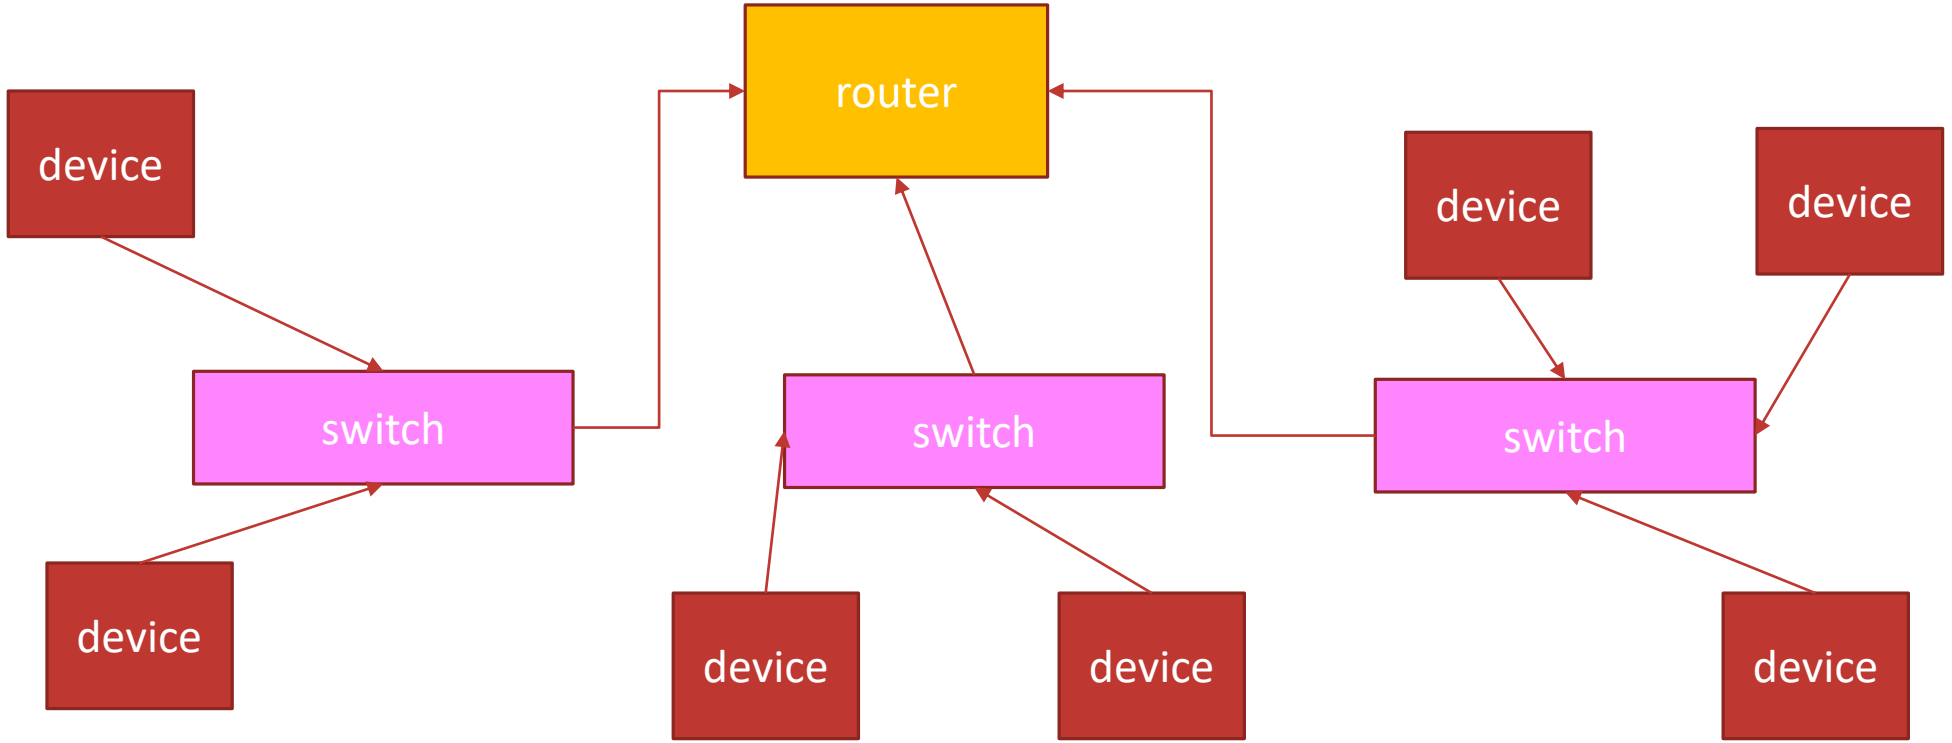

VLANS - HISTORY

Why not put all these physical LANs on one switch?

Saves “wasted” ports – still same cabling

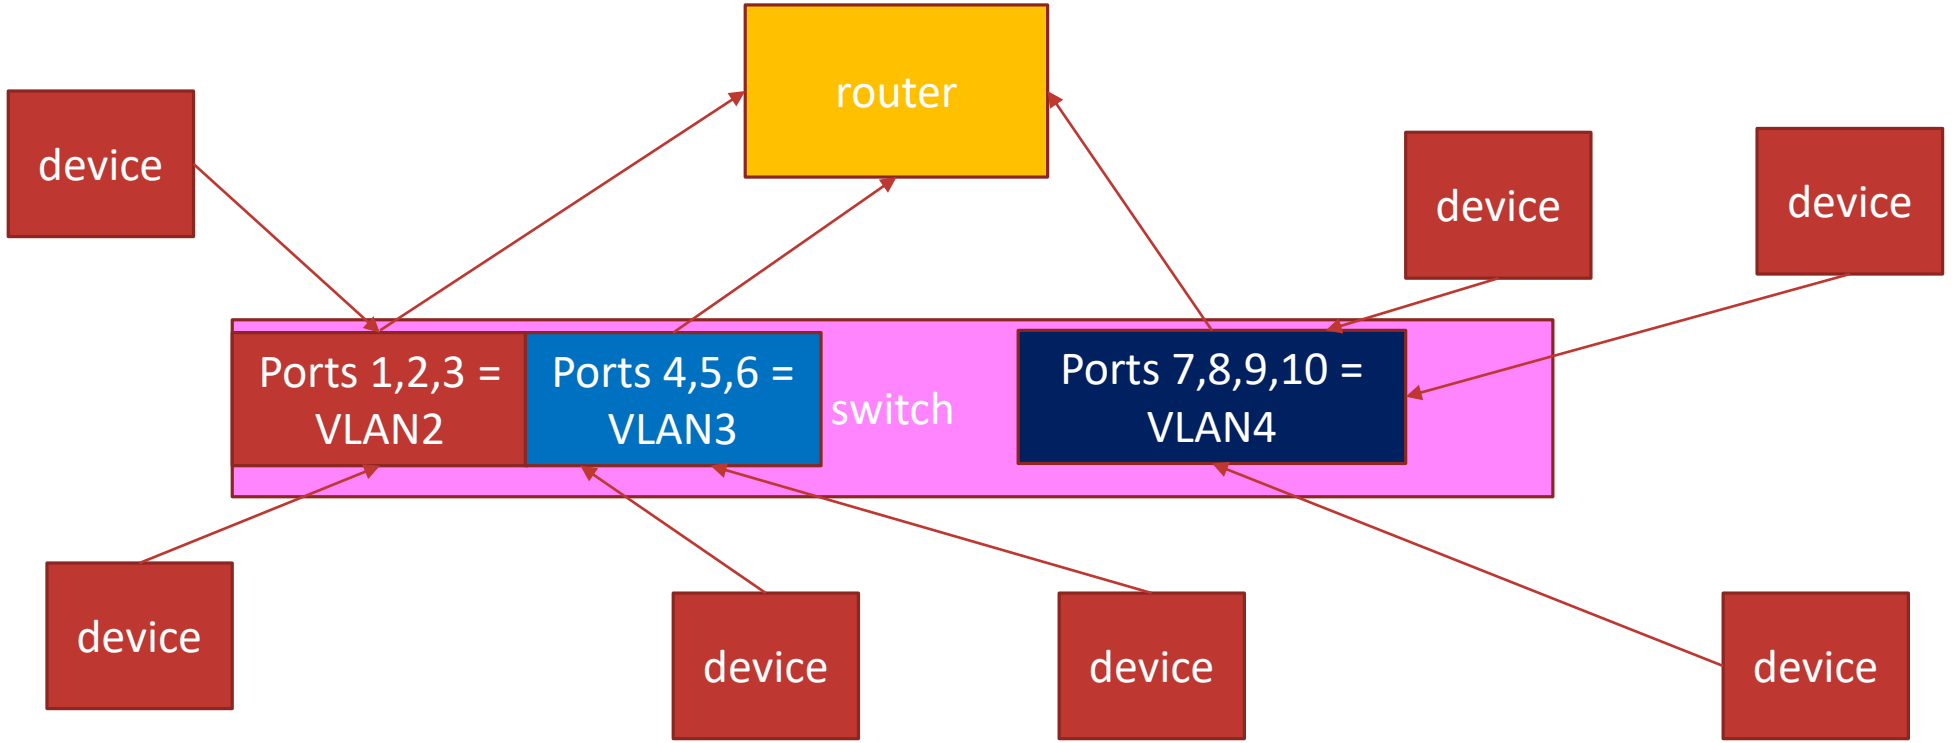

VLANS - HISTORY

Port based VLANs certainly saved money on expensive switch ports

Now 802.1q “tagging” addresses duplicate cabling – tag allows frames to be on same wire

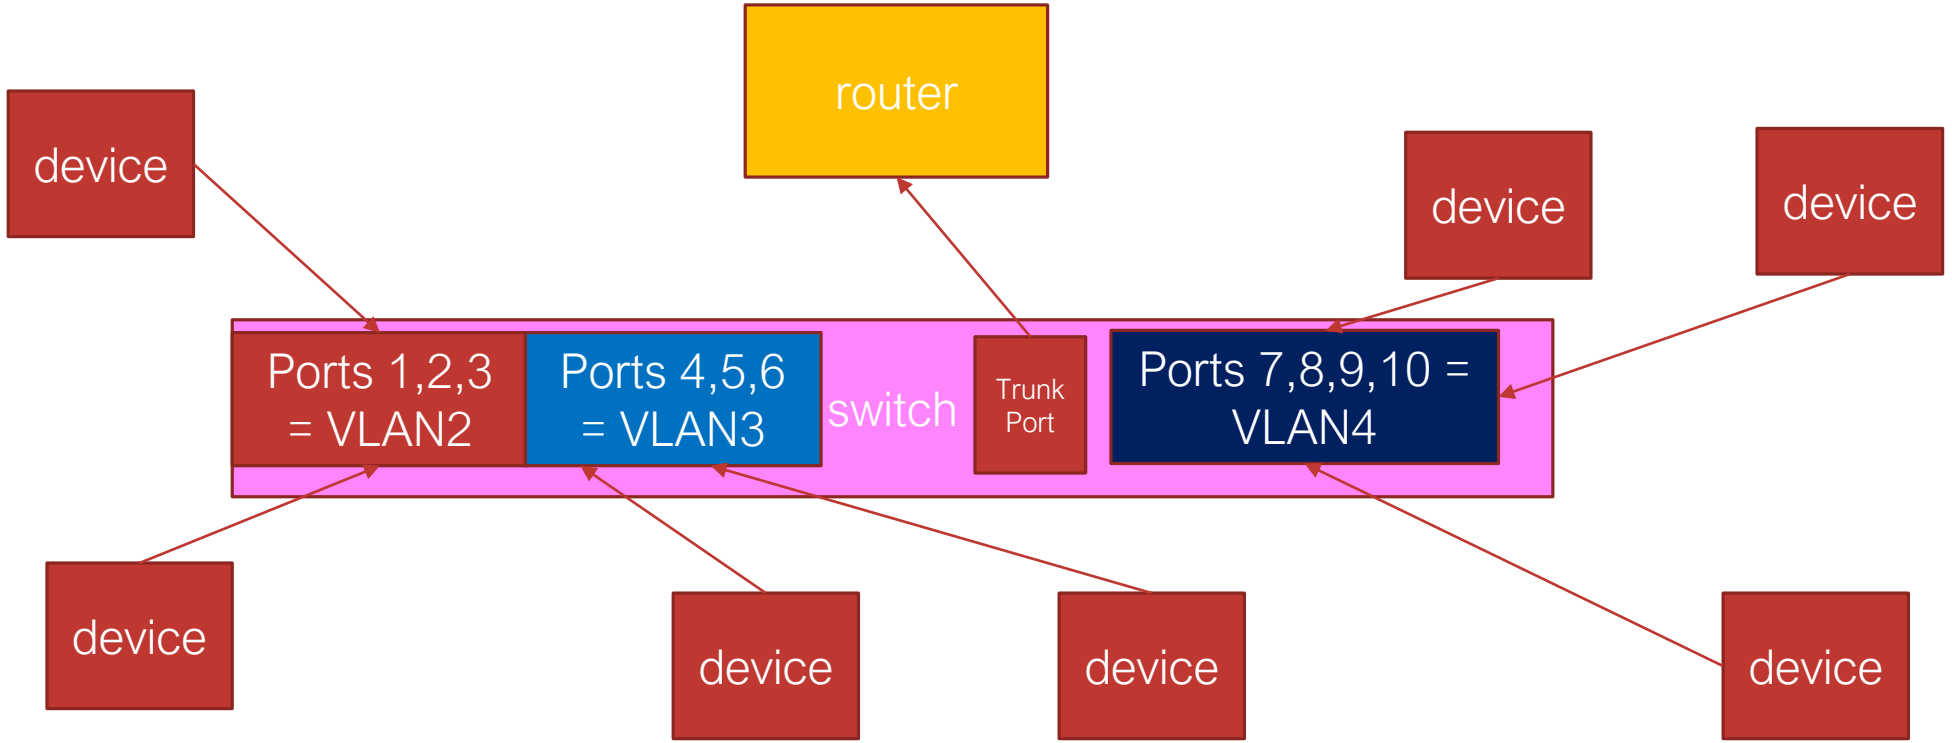

VLAN TAGS – AND PVID (PORT VLAN ID)

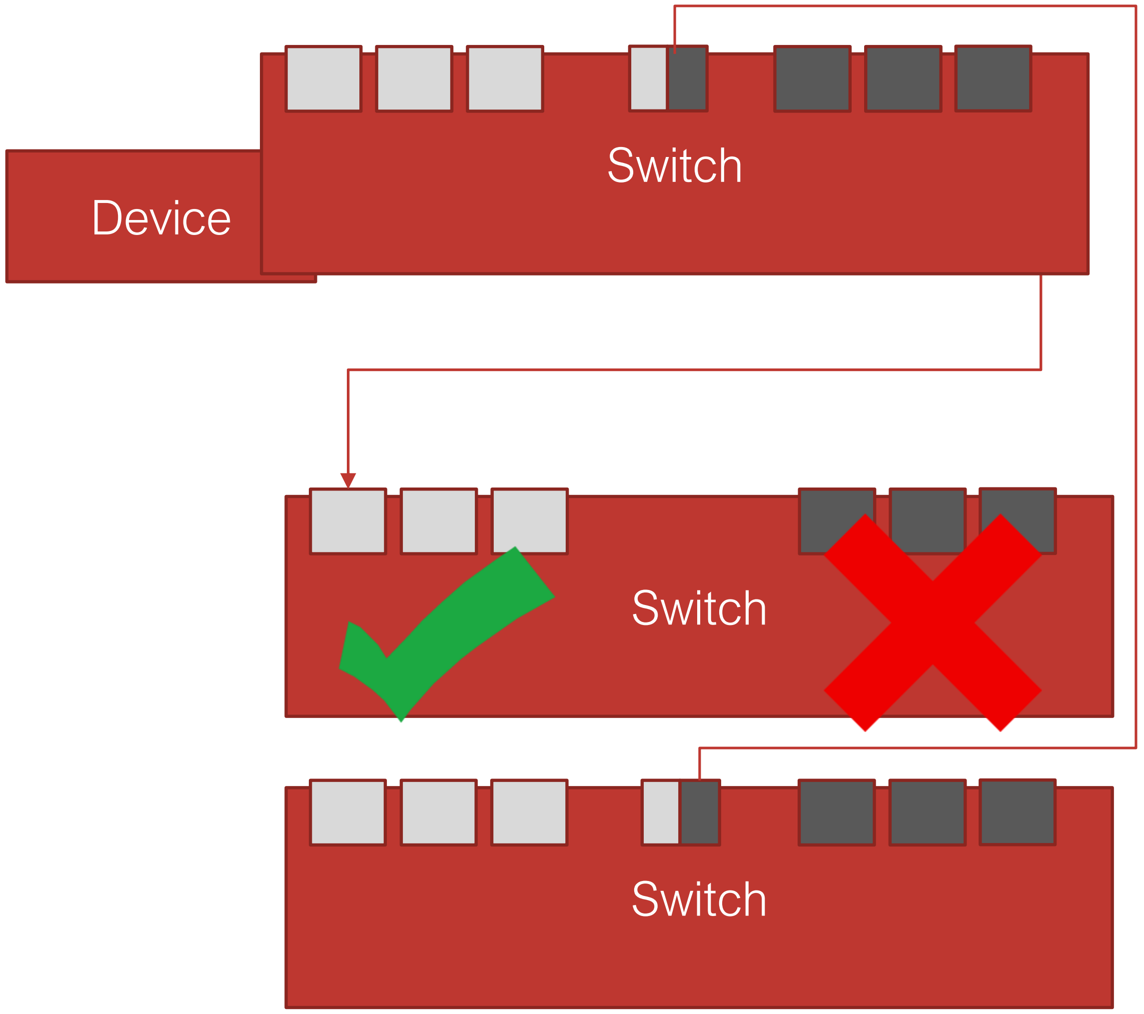

1. [illegible] the switch. The
[illegible] bits (7x 10101010)
[illegible] (10101011) this
[illegible] through any
[illegible] able to leave
[illegible] leaving through any
[illegible] only
[illegible] the IP packet)
[illegible] traffic with different tags can be carried on the same physical connection

IPV4 IP ADDRESSES – HOST AND NETWORK

- Routers segment Broadcast domains – because they will not forward broadcast traffic
  - Routers do not care about MAC addresses (they obscure them – a good security feature)
- VLANs can be used to separate broadcast domains using 802.1q tags
- Routers can join IP subnets together
- There are 2 “kinds” of IP address
  - “Host” addresses – the IP address that a device has
  - “Network” address – the IP address that describes a group of hosts
- The size of an IP subnet is defined by the Subnet Mask. The subnet mask:
  - Is a 4 octet number that is supplied via DHCP or an essential part of Static IP configuration
  - Allows devices to know which IP addresses are “local” (same subnet) or “remote”
  - Local IP addresses cause the device to issue an ARP
  - Remote IP addresses are forwarded to the MAC address of the “default gateway” (if one is defined)
  - If there is no defined “default gateway” the message will just be dropped

IP SUBNET MASKS

- There are 32 bits worth of IPv4 addresses – this could be 1 network with 4.3 billion devices in it
- Imagine a Broadcast domain that size!
- This is why we have IP subnets (to divide up that broadcast domain)
- The subnet mask is the “boundary” between the “network” address section of an IP address and the “host address” section of an IP address
- A Valid Subnet mask is always a continuous line of binary 1s until the mask is reached, then it becomes 0s

Subnet Mask (decimal)	Subnet Mask (binary)	Subnet Mask CIDR	Number of hosts per network	Number of possible networks
255.255.255.0	11111111.11111111.11111111.00000000	/24	254	16.7 million
255.0.0.0	11111111.00000000.00000000.00000000	/8	16.7 million	256
255.255.255.252	11111111.11111111.11111111.11111000	/29	6	537 million

VLANS AND SUBNETS

- A single VLAN is a single Broadcast domain
- An IP subnet is also a broadcast domain
- VLANs can therefore be described to a router as an IP subnet

- WARNING! – static IP addressing two IP subnets in the same VLAN is not wise!
  - When an IP broadcast packet is created in a VLAN – it is STILL mapped to the broadcast MAC address
  - At Layer 2 a device cannot distinguish between a broadcast frame from another IP subnet
  - Therefore an IP broadcast in a different IP subnet in the same VLAN will still go to devices in the “other” IP subnet – This is a potential security risk
- Tip – Make each VLAN its own IP subnet

ROUTER "VIRTUAL INTERFACE" "VIRTUAL LAN"

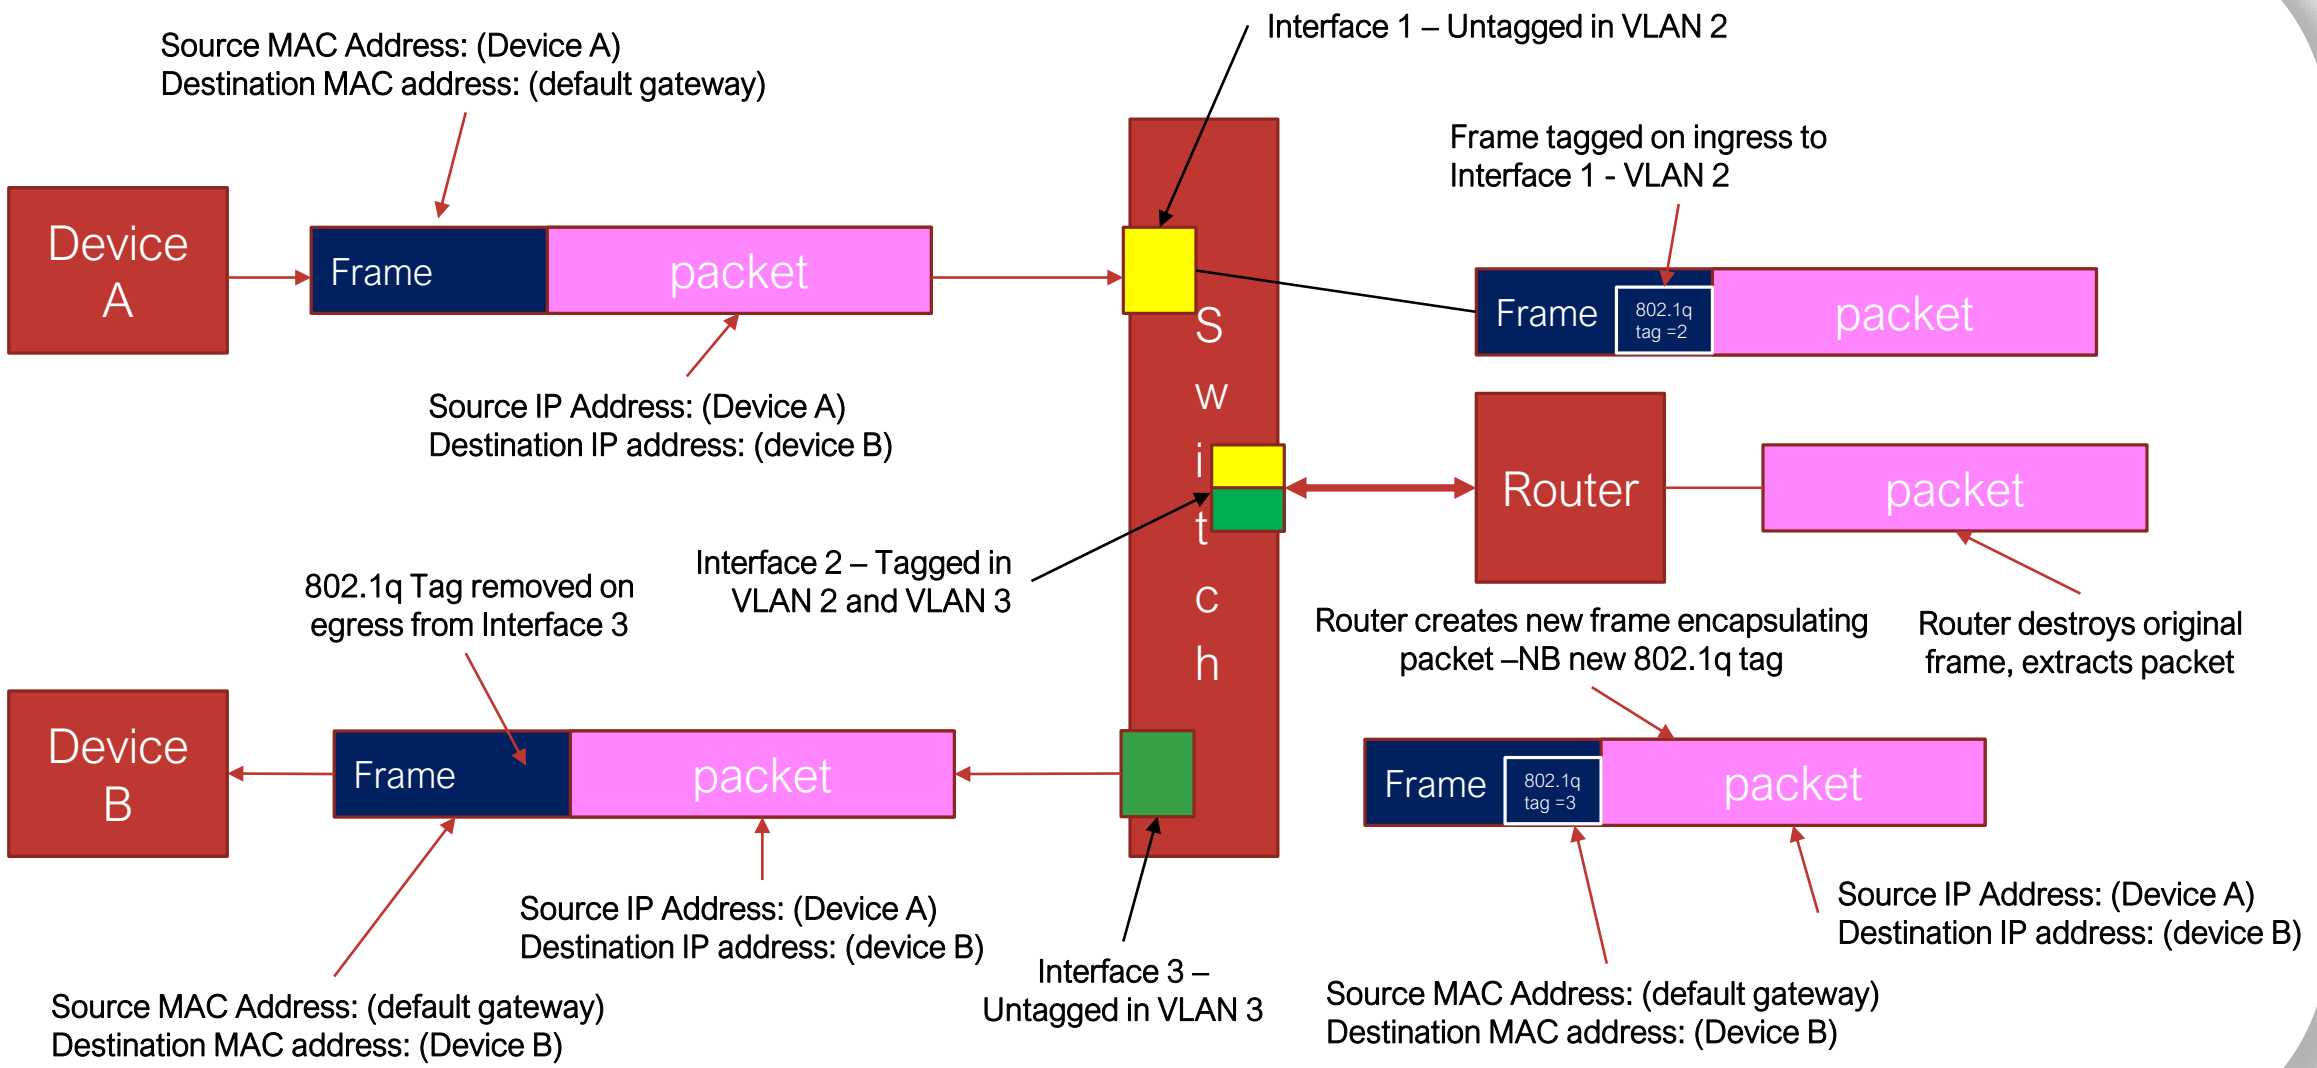

WHY HAVE DIFFERENT SIZED SUBNETS?

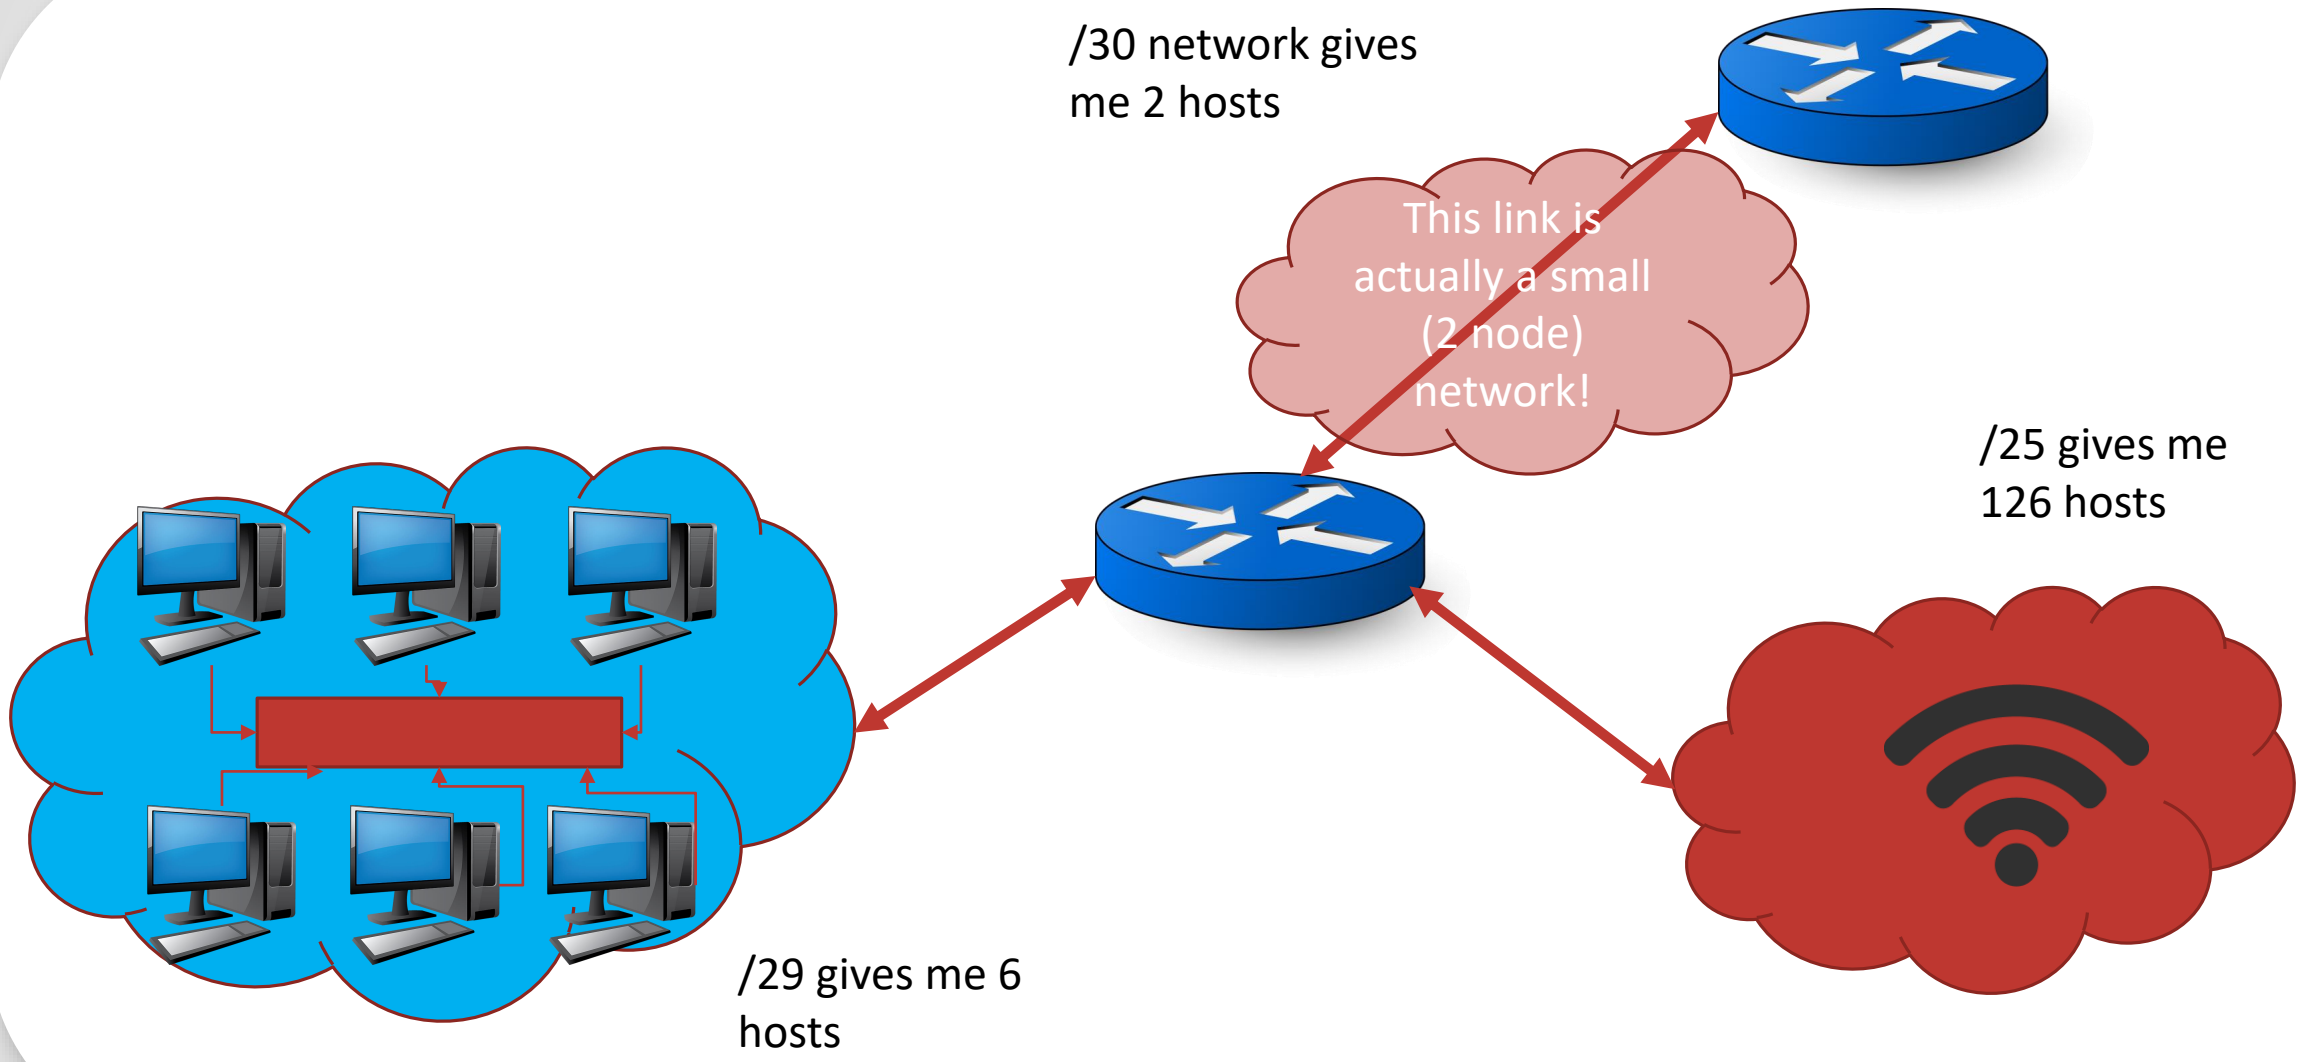

FOR MOST EFFICIENT IP SCHEME SEGMENT BIGGEST TO SMALLEST

Network IP	First Host	Last Host	Broadcast IP	Subnet Mask	Number of Hosts
192.168.1.0	192.168.1.1	192.168.1.126	192.168.1.127	255.255.255.128	126
192.168.1.128	192.168.1.129	192.168.1.134	192.168.1.135	255.255.255.252	6
192.168.1.136	192.168.1.137	192.168.1.138	192.168.1.139	255.255.255.254	2

Should I use Static or Dynamic IP addresses?

- Normally Dynamic is easier, and less prone to human error
- Static is sometimes preferable for servers (but not essential)
- Static IP addresses are only essential when using Organization SSL certificates (the only “excuse” acceptable to obtain a public static IPV4 address in Europe)

ADVANCED DANTE NETWORKING: SECTION 4

IN THIS SECTION…

Quality Of Service

- What QoS is, and what it is not
- When should you use it?

The “Third Kind” of communication

Multicast

- Why is Multicast useful
- How do we manage Multicast?
- IGMP – Deep Dive

QUALITY OF SERVICE (QOS)

WHAT IS QUALITY OF SERVICE (QOS)?

- Quality of Service (QoS) is the business of traffic prioritization and management
- The aim is to provide a specific level of service to an application using the network
- How do I specify a level of service?
  - Throughput
  - Delay
  - Delay variation
  - Packet loss

WHAT IS QUALITY OF SERVICE (QOS)?

- QoS comes in two basic forms:
  - *Class based*, using techniques like queueing and prioritization (e.g. Diffserv)
    - Relative
    - Specify what is important, what is less important
    - Simpler to implement
  - *Reservation based*, using techniques like rate control, traffic shaping, and admission control (e.g. QoS NSLP, or Intserv)
    - Absolute
    - Specify how much, how often, and on what terms - then make a decision if it is possible
    - Complex to implement

WHAT IS QUALITY OF SERVICE (QOS)?

Audio over IP only really cares about some aspects:

- *Delay variation* – helpful to minimize latency and achieve good PTP performance under a wider range of conditions
- *Aggregate throughput* – we need to get a lot of audio packets through the network
- *Packet loss* – we need all the audio to arrive

WHY DO IT?

- QoS techniques can help extend performance further into a region where bandwidth limitations might otherwise cause problems
  - Help to squeeze more out of the infrastructure
  - Running a link at over 60-70% utilization
  - Mixing 100Mbps and 1Gbps in the same network
  - Running a shared services network (e.g. audio and video, enterprise IT etc)
- QoS is not magic, it cannot manufacture better performance

WHY DO IT?

- Simple fact: The best QoS available is bandwidth
  - Cheap, plentiful, easy to use
  - Dante will present a predictable peak load to the system – this should enable you to ensure that the network is appropriately dimensioned

- **Key things to remember:**

- Dante devices transmit relatively small, fixed size packets that are “nicely” distributed by the transmitter

- This helps to minimise queueing in the switch, and greatly reduces the need for QoS to achieve good performance

CLASS BASED QOS

- Class based QoS is the most widely deployed technique in COTS Ethernet switches
  - We can thank VoIP for this!
- Uses a tag written into the IP or Ethernet header of each packet
  - Just a number, no special meaning
- Switch needs to inspect each packet and map to a specified queue at the output port
- The queues are then emptied following a set of rules that govern the process
  - These rules define the level of service experienced by the traffic in each class

CLASS BASED QOS

- Has many advantages
  - Simple
  - Low per-packet overhead
  - Implement in switch silicon
  - Distributed
  - Flexible – network can apply tags, change them, interpret them using a system specific policy etc.
  - Very good at protecting high priority traffic from lower priority
- Disadvantage
  - Can overload within a traffic class – e.g. everything high priority…
  - When everything is important… nothing is
  - Not a limitation in vast majority of cases

DIFFSERV

- DSCP values are mapped to a queue at each output port in the switch
  - In a Dante centric network the highest priority queue is used for PTP traffic (very small payload)
  - The next level of priority is used for Audio data – slightly larger
  - The level of priority below that is used for time critical control data
  - And finally the last queue is available for everything else
- Queues in switches send frames onto the wire after applying this “sorting” criteria
  - Don’t forget this is per-switch, so all switches should have the same configuration if you want consistent behaviour
- Extremely well suited to ensure that high priority traffic is not affected by lower priorities
  - PTP amounts to only a few packets per second per device
  - Audio is several thousand packets per second per flow

EXAMPLE

- Airline check in counter or boarding lanes
  - Customers are assigned to a class based on importance to the airline (Platinum, Gold, the rest…)
  - Equivalent of a DSCP Tag
  - Important customers have a priority check-in counter
  - Equivalent of a dedicated queue
  - When it is not busy, all customers served quickly
  - When it gets busy, important customers are served quickly

- Why does this work?
  - There are relatively fewer priority customers than other categorie
  - Minimal wait for priority customers at check-in, at boarding…
- When does this fail?
  - At Chicago O'Hare United terminal, where everyone is Platinum
  - If everyone is important, no one is…

SIMPLIFIED QUEUE MODEL

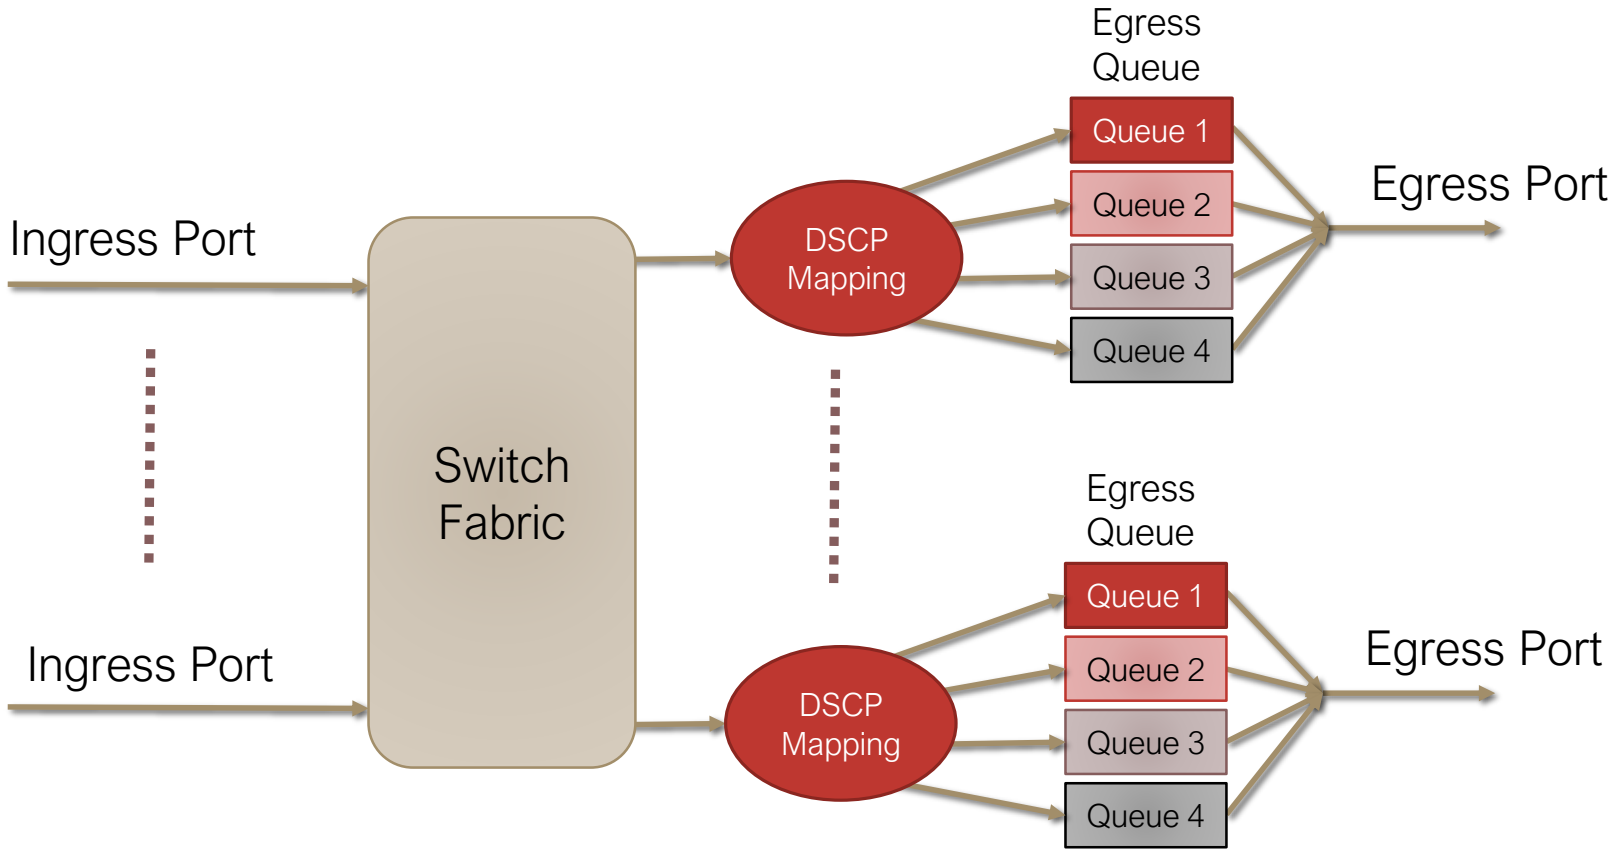

QUEUEING MODES

Switches can be configured to manage queues using different algorithms

- **Strict Priority** – Always serve the most important class
- **Weighted Round Robin** – Share service time amongst the queues following per-queue weights
- **Shaped Round Robin** – Share service time amongst queues in a statistically shaped manner (Cisco use this term)

QUEUEING MODES

- Dante needs a *strict priority* queue for PTP traffic
  - This will ensure that if a sync packet tagged DSCP 56 arrives it will always be passed through the as quickly as possible
  - Using Strict Priority minimizes delay variance experienced by PTP traffic through the system, particularly if under heavy load
  - WRR or SRR will increase delay variance and reduce PTP performance

- Many data-centre switches default to SRR or WRR modes – you need to check for this if using QoS on higher spec switch products

DANTE DSCP VALUES

Priority	Usage	DSCP Label	Hex	Decimal	Binary
High	Time critical PTP events	CS7	0x38	56	111000
Medium	Audio, PTP	EF	0x2E	46	101110
Low	(reserved)	CS1	0x08	8	001000
None	Other traffic	BestEffort	0x00	0	000000

These values chosen to align with default mapping in many switches

CISCO SG300

- Switches like the Cisco SG300 can have QoS queueing set up very simply
- All egress queues can be set to strict priority
  - Highest priority queue is emptied first, followed by next in order etc.
- It is simple to set up – set switch to Trust DSCP, and then assign DSCP tags according to previous table… job done

SOME OTHER SPECIFICS – TRUST DSCP?

- Tells the switch to trust that the endpoints will tag traffic correctly, not abuse high priority
- Useful with services like Dante that do not require user configuration – the endpoint can only generate correct DSCP tags
- What if I turn this off? - the switch will “strip” DSCP tags from all packets and ignore them
- Most campus, enterprise, and carrier networks will NOT trust endpoint DSCP tags, and instead apply a policy as defined by the system admin

SOME OTHER SPECIFICS – DSCP REMAPPING

- Many switches allow you to strip or remap DSCP tags on ingress
- Define ACLs on ingress and handle matching packets
- Very common at borders between networks
- Useful for protecting PTP in non-trusted environments, or making PTP fit within the scheme in operation on that network
- **Example:** Match PTP traffic on UDP port 319, 320, 321, and apply DSCP 56

SOME OTHER SPECIFICS – CLASS OF SERVICE (COS)

- What about COS, or Class of Service?
  - COS relies on a Layer 2 method for tagging frames
  - Part of the 802.1q header (remember VLANs?)
  - Only tagged VLAN traffic will have an 802.1q header
  - Most endpoints can't generate 802.1q tags
  - Not useful end-to-end and will most likely be stripped once the traffic hits a gateway
  - Note: Even though they might change, DSCP values can be retained end-to-end, no matter what link type
- What about IPP or IP Precedence?
  - Supports less resolution than DSCP, so used less frequently in larger systems

ENTERPRISE SWITCH FEATURES

- On larger (more powerful, and more expensive) enterprise switches QoS can be more “flexible” (read complicated)
- Switches in this category allow the network operator to cater for many different services
  - Leverage the flexibility and power inherent in IP
- It is important to understand exactly how the QoS settings on the particular platform are implemented
  - Manufacturer training courses are best for this
  - There can be a lot of detail here

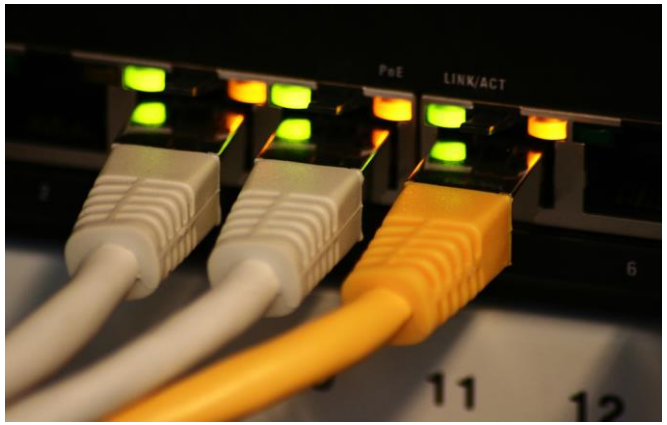

IS IT LIKELY TO MAKE A DIFFERENCE?

- Using QoS will “catch” momentary “spikes” in demand
  - These are generally caused by untagged DSCP traffic
- If you are moving 512 channels of Dante over the network link in question, you still have “headroom” to accommodate other services (more than 100mbps)
  - How many people have 100mbps Internet connection?
  - How many people see 100mbps of Internet traffic on the same network segment as AV?
- You will generally need to have a pinch point somewhere in the network in order to see a positive benefit from QoS configuration
- You are also likely to see negative impacts from over configuration in environments where QoS configuration is not necessary

MULTICAST

MULTICAST IP ADDRESSING

- A multicast group is defined by a “multicast destination IP address”
  - Any host within the address scope can “listen” to the group address
  - Really really useful when there are many more receivers than transmitters
- Special reserved IP and MAC addresses indicate multicast destinations
  - MAC addresses start with 01:00:5E
  - IPv4 range: 224.0.0.0 – 239.255.255.255
  - IPv6 Range: FF00::/8
- Multicast IPv4 addresses do not map 1:1 to Multicast MAC addresses
  - Up to 32 Multicast IP addresses for each Multicast MAC address
- This creates challenges for a L2 switch uniquely identifying multicast streams
  - Hence the use of “IGMP Snooping”

INTERNET GROUP MANAGEMENT PROTOCOL (IGMP)

- Internet Group Management Protocol (IGMP) is a protocol used to manage multicast traffic in IPv4 routed networks
  - IPv6 replaces IGMP with “Multicast Listener Discovery (MLD)
- Used to help manage multicast flooding
  - Without it, multicast traffic needs to go everywhere - not ideal
  - Multicast management is generally a good thing for high rate, media traffic
- Protocol is responsible for working out which networks or hosts need a specific multicast group
  - The router sends query messages to all hosts on each network
  - Hosts respond with the groups they want to join
  - What about inside the VLAN?
- All Dante devices implement IGMP (as a host)

IGMP QUERIER

- Most switches today are able to function as an IGMP Querier
  - Querier periodically sends query message to all hosts on 224.0.0.1
  - What group addresses do you want?
  - Hosts respond with “membership report” to 224.0.0.2 (or 224.0.0.22)
  - I want A.B.C.D
  - You should only have one active Querier per VLAN
  - Multiple queriers *can* be configured amongst some manufacturers equipment

- Intermediary L2 switches “snoop” on IGMP traffic
  - Membership report identifies which interfaces need each multicast group
  - Switch then maps groups to ports and builds a multicast forwarding table
  - Multicast traffic only appears on the ports it is needed

IGMP AT WORK

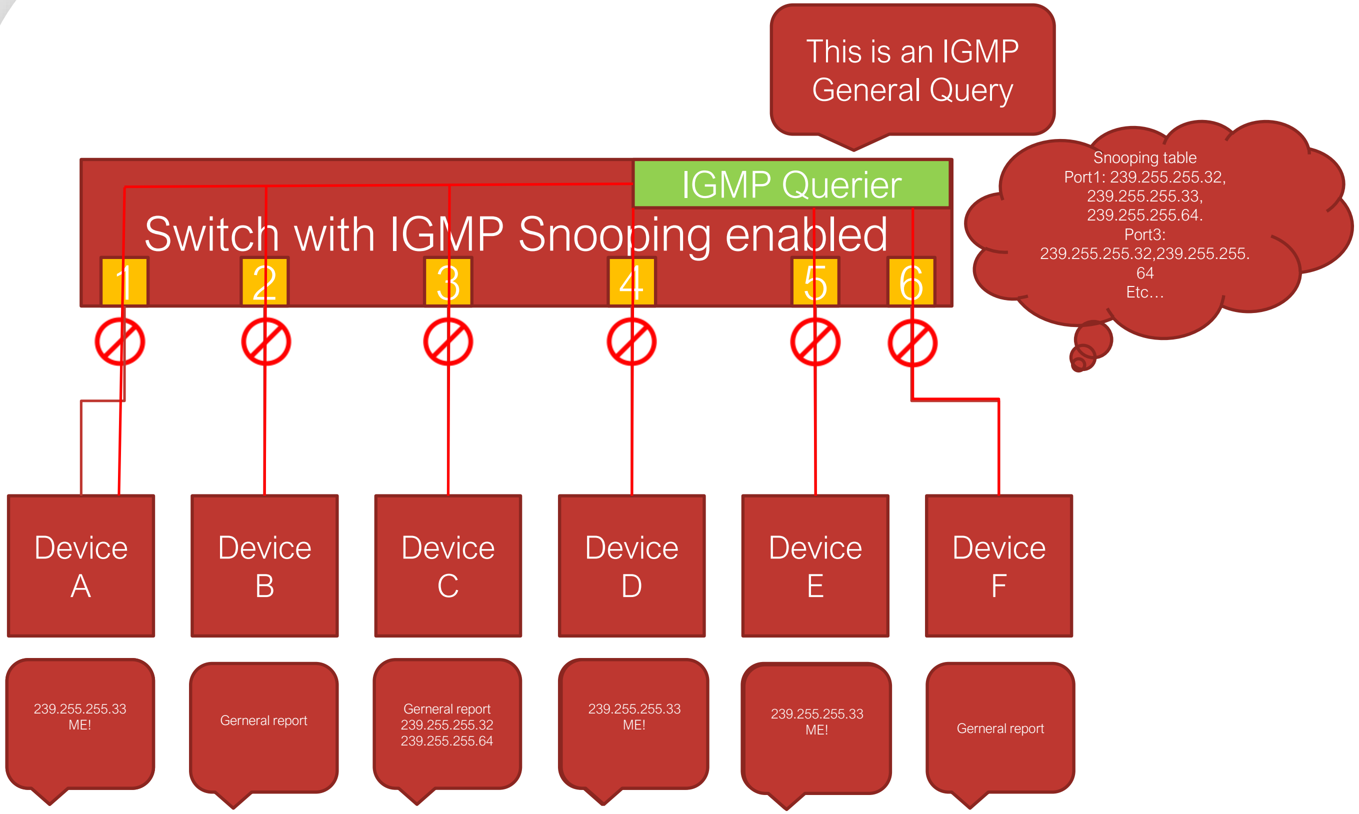

1. Turning on IGMP snooping will cause the switch to block all multicast
2. IGMP query messages must be present to give IGMP snooping something to “snoop” on
3. On many modern switches turning on IGMP snooping also turns on a querier inside the switch (as on SG300)
4. Queriers send queries, devices (receivers) send “reports”
5. The “snooper” builds a forwarding table based upon queries and reports
6. Queriers send two main types of query
   1. General query – asks which multicast groups the receiver is interested in
   2. Group specific query – asks who is interested in a specific multicast group
7. Devices send reports in response to queries
8. When a device is connected it will send an “unsolicited” report
9. This means that the IGMP snooper will forward the multicast group as soon as the device is connected (no need to wait for the querier)

Dante by Audinate

IGMP SNOOPING OPTIMISATION

- Running IGMP to manage IPv4 multicast involves:
  - Specific roles in the network Querier
  - Timers: Query Intervals, Timeouts etc.
  - Specific Multicast Addresses: 224.0.0.1, 224.0.0.2, 224.0.0.22 etc.
  - Multiple versions: IGMPv2, IGMPv3 etc.
  - Protocol implementations from different switch vendors – vendor specific optimisations: e.g. fast leave, querier proxies etc.
- It can be complex…
  - Especially if you start randomly changing timers or switch configurations
  - Random changes will have random effect

SIMPLIFIED RECIPE

- Set the Querier interval quite short – 15 or 30 seconds is OK
  - Minimizes potential for gaps in audio, increases responsiveness
- Leave timeout values at default settings
  - Longer is better in most cases
- Ensure that there is only active querier per VLAN.
  - Ideally this is a core network switch, or the router port
  - Symptom: Intermittent audio
  - Diagnosis: Wireshark shows IGMP Query messages from multiple sources

SIMPLIFIED RECIPE

- Avoid using fast leave
  - Doesn't offer anything
- Avoid IGMP proxies
  - Unless you are CERTAIN you know how it behaves
- Ensure that "block unregistered multicast" is not set on Netgear switches
  - Blocks traffic it should not (mDNS, PTP)
- Dante devices support IGMPv3

IGMP SNOOPING PERFORMANCE CONSIDERATIONS

- Remember – switches operate using MAC address table
- There are possibly multiple IP addresses per MAC address
  - Work needs to be done by the switch to resolve this
- Some switches also use a CPU to forward multicast traffic
  - This can lead to variability in forwarding delay
  - Potential for overload as the amount of multicast increases
- This means some switches can be slow dealing with multicast
  - Potentially can cause unacceptable latency for multicast audio or PTP
- Common Symptoms Include:
  - Devices losing sync periodically
  - Choppy audio when receiving a multicast stream

Dante by Audinate

MULTICAST PROPAGATION

- It is a common requirement that the network provides L2 VLANs over a L3 core network
  - Multicast is sometimes a challenge in those environments
  - Many system admins don't like multicast traffic in these environments
- PTP likes the forwarding delay through the network to be consistent and symmetric
- In some scenarios this may not be the case
  - Asymmetric non-multicast technologies such as MPLS, GPON etc.
  - ACLs incorrectly configured etc.
- Common symptoms include:
  - Multiple PTP masters
  - Devices failing to maintain sync
- Turning ON unicast delay requests for all devices can help diagnose, and possibly resolve the problem

TESTING MULTICAST SWITCH PERFORMANCE

Audio Testing

Create some Dante Multicast flows → Subscribe them to a receiver → Check the latency view in Dante controller → Turn on IGMP snooping in the switch → See how far the latency view has moved → Increase number of multicast flows through switch → Check latency performance again

PTP Testing

Create some Dante Multicast flows → Subscribe them to a receiver → Check the clock histogram in Dante controller → Turn on IGMP snooping in the switch → See if clock histogram has spread → Increase number of multicast flows through switch → Check sync performance again

ADVANCED DANTE NETWORKING: SECTION 5

IN THIS SECTION…

Understanding network resources for Dante

- Bandwidth use by Dante
- Effect of topology on bandwidth use

DEFINING RESOURCE REQUIREMENTS

- Any project in anything requires this phase of planning

- Required resources for a Dante network:
  - ✓ Enough Transmit flows to serve all receivers
  - ✓ Enough Bandwidth to carry flows
  - ✓ Unblocked logical connections (“wire” is not cut, and is “plugged in”)
  - ✓ Enough Receive flows available on devices connecting to transmitters

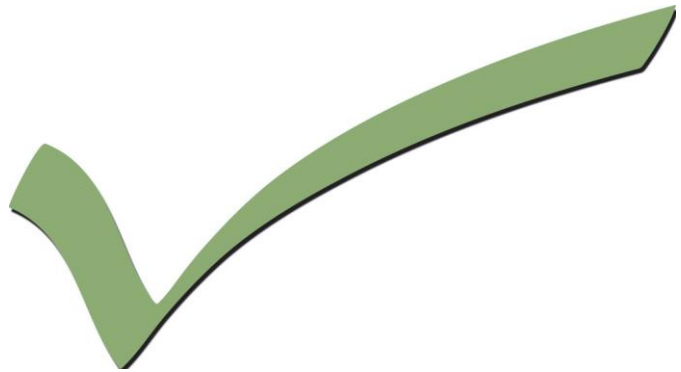

WHAT DANTE PUTS ONTO THE WIRE

- In order to make a complete audio network solution we need:
  - Audio transport
  - Clocking
  - Device discovery
  - Device and routing control

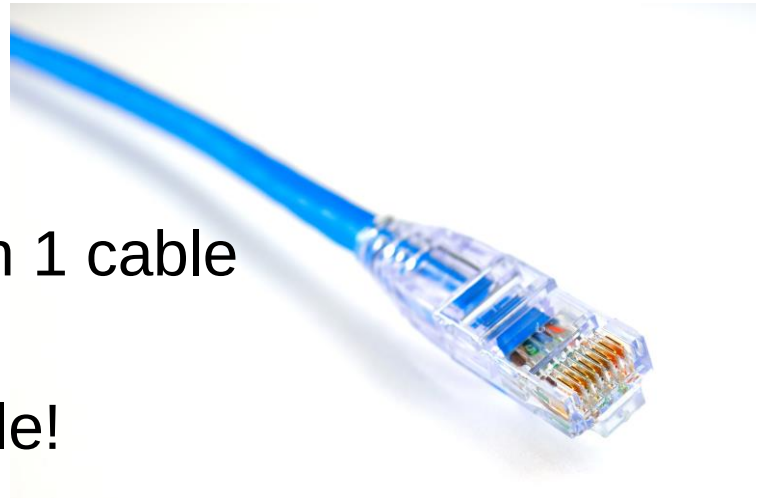

- Using the IP over Ethernet model allows us to do this on 1 cable
- Some devices also take their power from this same cable!
- The OSI model gives a framework to do this very tidily

ARCHITECTURAL PRINCIPLES

- What are the main considerations when architecting a network?
  - Identifying required services
  - Bandwidth utilization
  - Specifying an optimal infrastructure for predicted bandwidth demands
  - Identifying potential “bottlenecks” – designing them out or managing them
  - Ensuring a scalable architecture (the network will grow over time)
  - Understanding possible future service requirements
- A long and daunting list – luckily there are established “rules”

BANDWIDTH UTILIZATION – DANTE FLOW

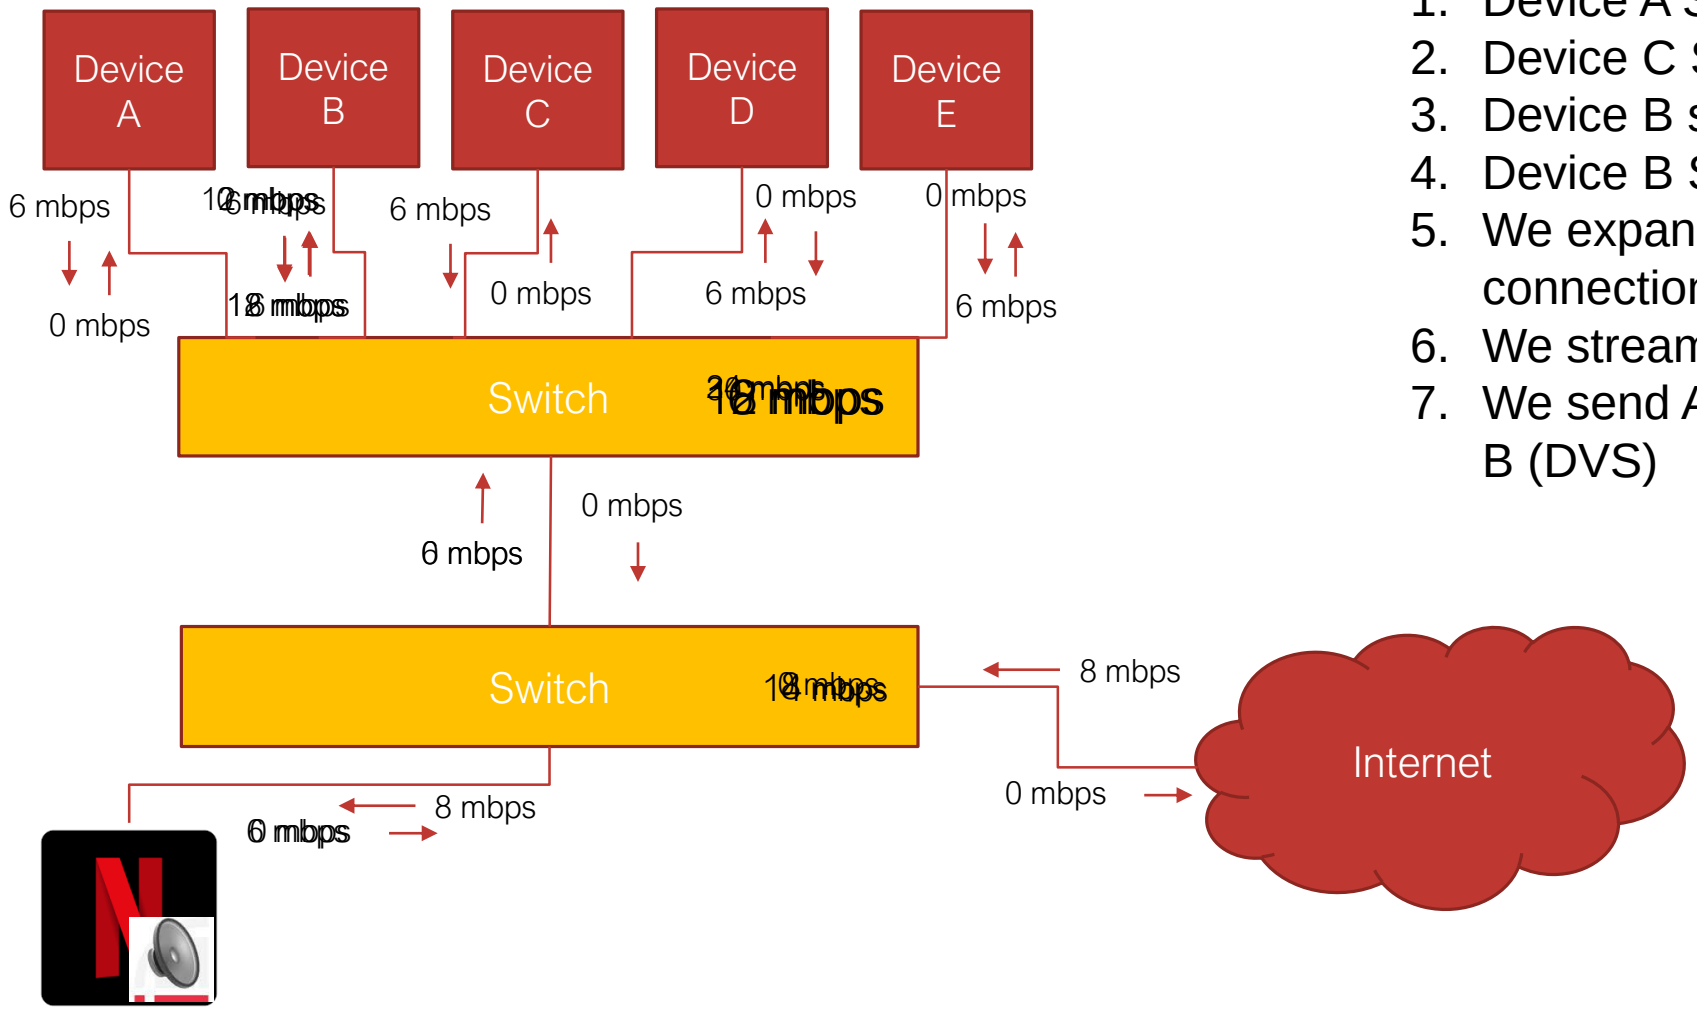

1. Device A Sends a signal to Device B
2. Device C Sends a Signal to Device B
3. Device B sends a signal to Device D
4. Device B Sends a signal to Device E
5. We expand the network (add Internet connection)
6. We stream Netflix from Internet
7. We send Audio from Netflix to Device B (DVS)

UNDERSTANDING RESOURCES - BANDWIDTH

- A flow is a sequence of packets sent to an IP address / port
- A Unicast flow contains 4 channels
  - 2ch on Ultimo2 devices
  - At 48KHz sample rate a unicast flow uses 6mbps bandwidth
  - If only one channel is “subscribed” the other 3 channels are “sent” as silence
  - As other channels are added between the 2 devices audio content from the successive channels “dumped” into the flow
  - New flows not created if there is “space” in existing flows
- A Multicast flow contains up to 8 channels
- At 48KHz sample rate a multicast flow uses up to 12mbps bandwidth
  - Only selected channels are included in the flow
  - Cannot add new channels to an existing multicast flow
  - Avoid creating multiple “single channel” Multicast flows that follow similar paths

UNDERSTANDING RESOURCES - FLOWS

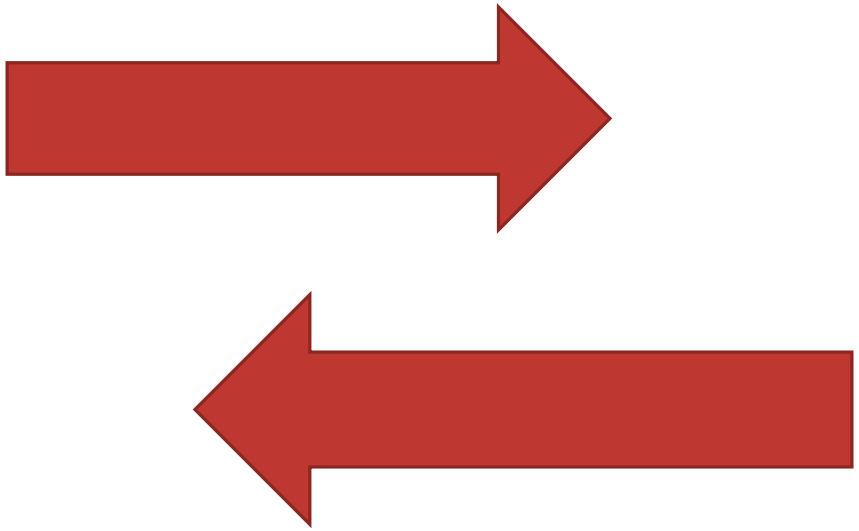

- A flow limit is an internal device resource limit
- Different device types have different limits
- It does not matter to the device whether the flow is unicast or multicast (as the destination IP address simply changes)

UNDERSTANDING RESOURCES - FLOWS

- Each Dante device has a number of "**Transmit Flows**" and **"Receive Flows"** available to it
  - An Ultimo device has 2 transmit and 2 receive flows
  - A Dante Virtual Soundcard has 16 transmit and 16 receive flows
  - A Brooklyn II device has 32 transmit and 32 receive flows
  - A Dante PCIe card has 32 transmit and 32 receive flows
  - A Dante HC device has 128 transmit and 128 receive flows

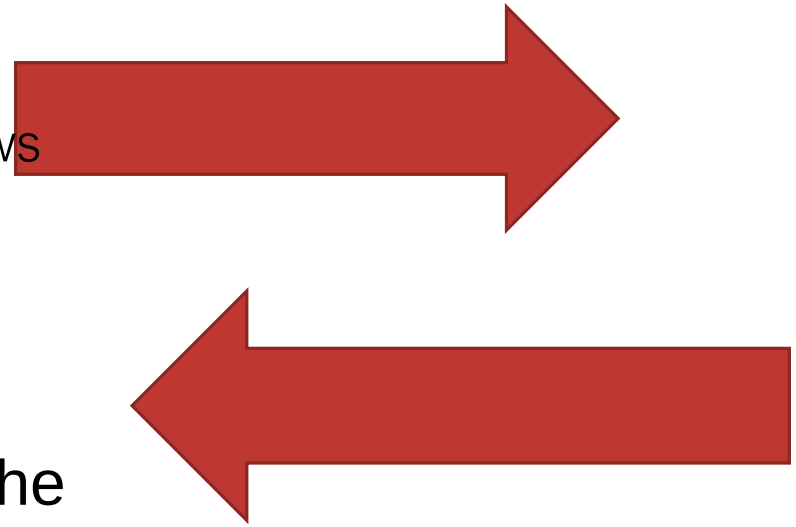

- Flows are used to increase the overall efficiency of the system

UNDERSTANDING RESOURCES - TX FLOWS

- A Tx flow is created in 2 different ways

  In Unicast mode –

  - The Tx flow is created when the receiving device requests one (or more) channels
  - If room exists channel is added to existing TX flow
  - A single unicast flow can contain up to four channels

  In Multicast mode –

  - The Tx flow is created immediately (it doesn't need a receiver to request it)
  - A single Multicast flow can contain up to eight channels
  - Cannot add channels to a multicast flow

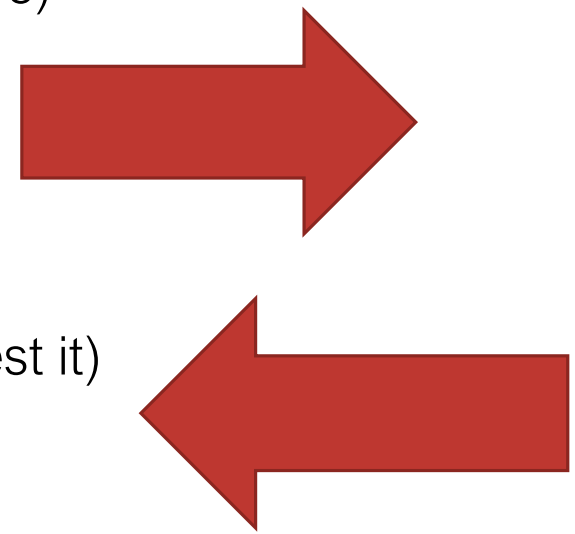

- All Dante flows are UNICAST by default
- Only intentional programming can create a multicast flow – in Dante controller – this is done by a human

UNDERSTANDING RESOURCES – TX FLOWS

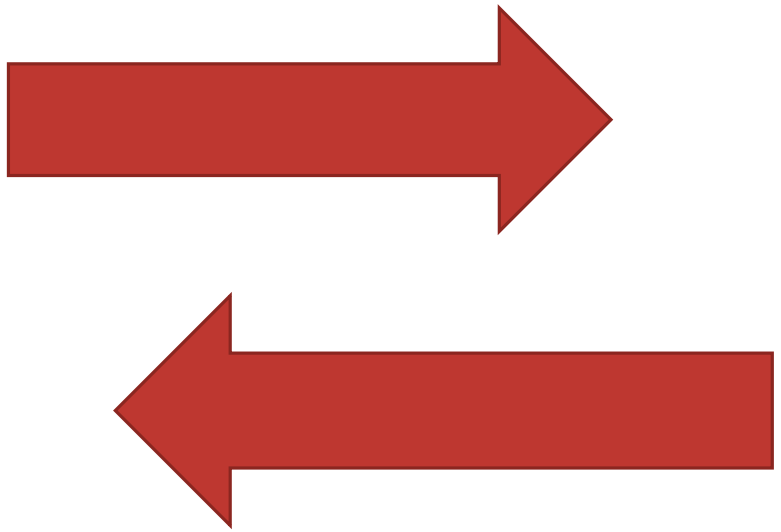

- A Dante flow is Unicast by default
- Optimisation can be achieved using Multicast when distributing one to many
- When channels are made Multicast – those channels will be transmitted in a new multicast flow
- If a device is using all flows – no more subscriptions can be made by receivers
- Tx Flow usage is visible on the Transmit tab of Dante Controller

UNDERSTANDING RESOURCES – RX FLOWS

- A receiver will utilize a new RX flow resource each time a subscription is created that either:
  - Requires a new unicast TX flow on the transmitter, or
  - Is the first subscription to a channel already in a multicast flow

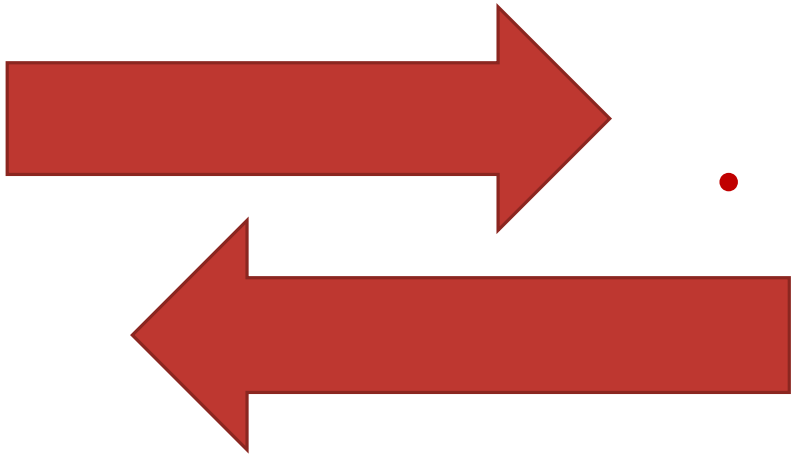

- Rx Flow usage is visible on the Receive tab of Dante Controller

SIMPLE NETWORK RESOURCE UTILIZATION

Recap

An audio network consists of 3 things

- Devices that create flows (microphones, stage boxes, DSPs, playback sources etc.)
- Infrastructure that carries the data
- Devices that receive flows (DSPs, Amplifiers, Powered Loudspeakers, Recorders etc.)

Dante by Audinate

ADVANCED DANTE NETWORKING: SECTION 6

IN THIS SECTION…

Subscriptions in Dante

- Subscription process
- Multicast subscriptions
- Topology and QoS

SUBSCRIPTION PROCESS (UNICAST)

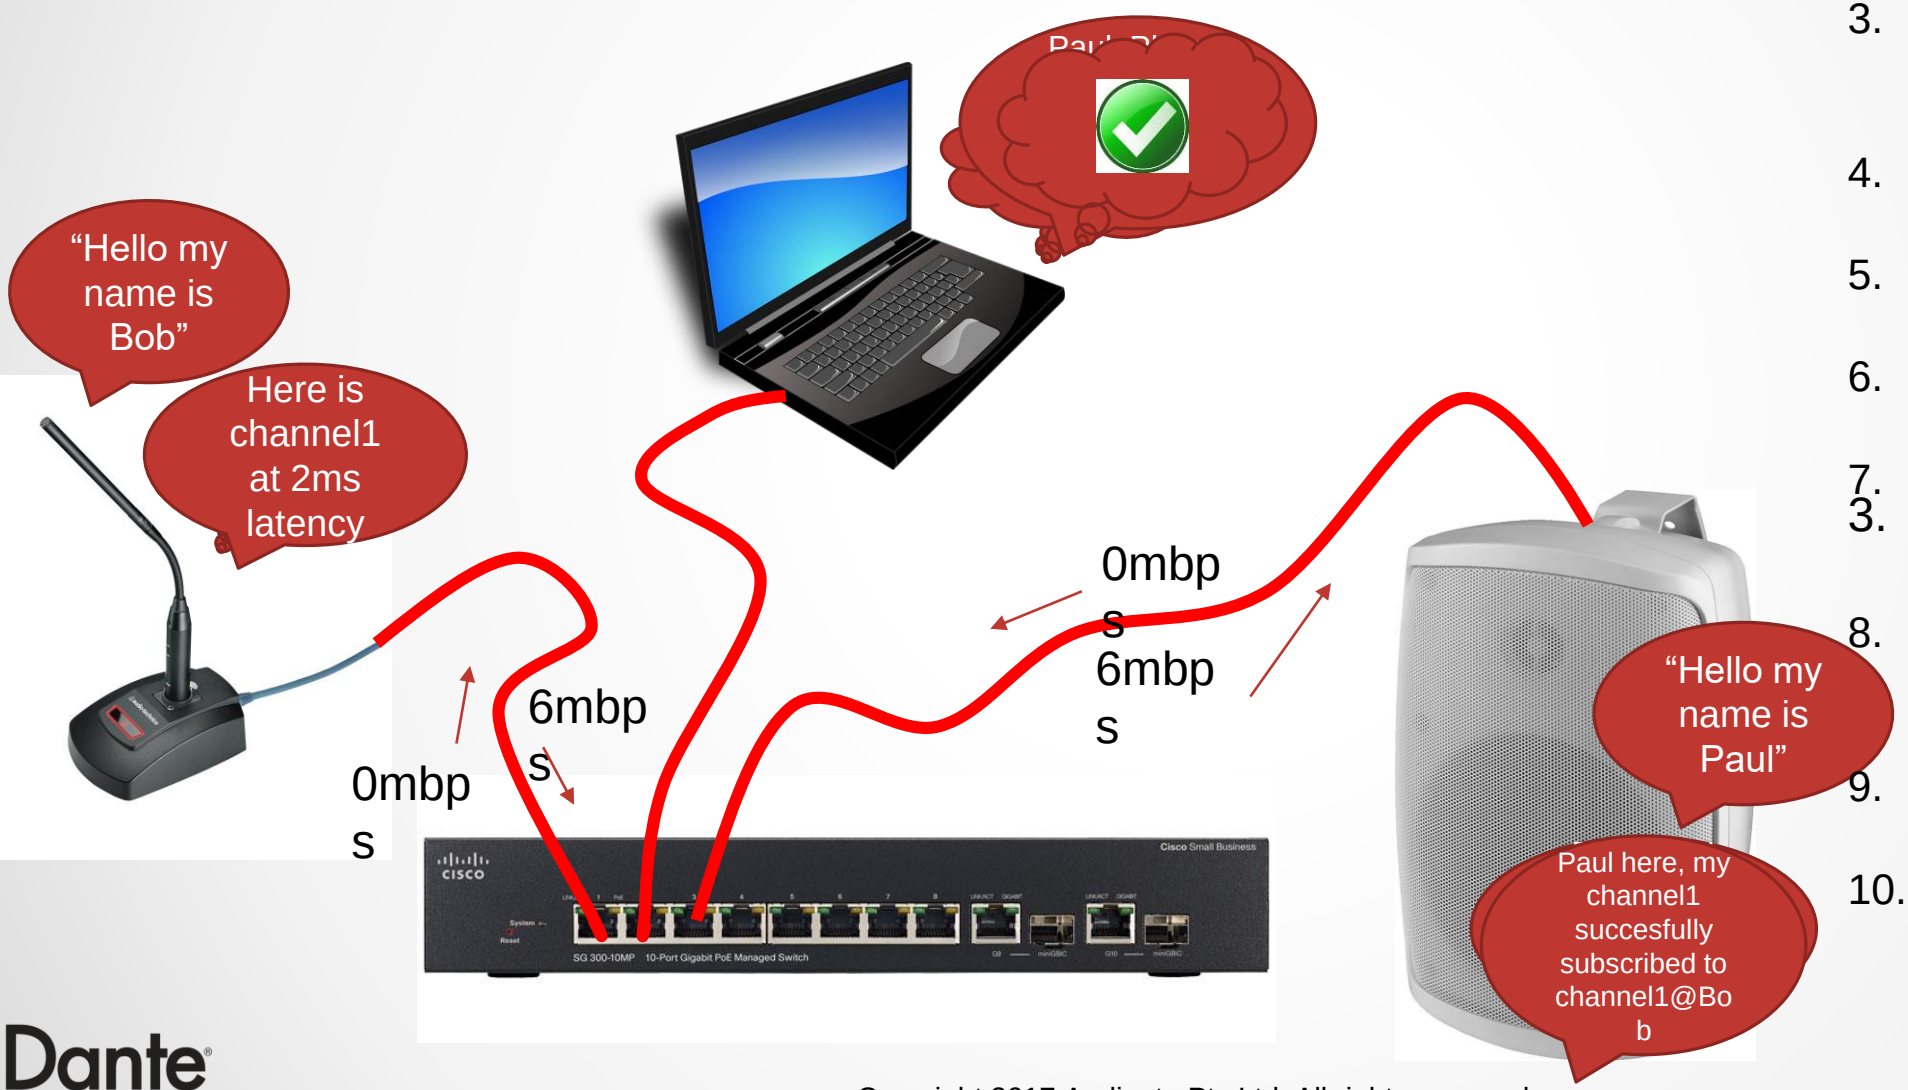

1. Plug laptop (running Dante Controller)into switch
2. Plug Microphone into switch
   1. Microphone “announces” itself (multicast)
   2. Dante controller “hears” the announcement and populates routing page
3. Plug Speaker into switch
   1. Speaker “announces” itself (multicast)
   2. Dante Controller “hears” the announcement and populates routing page

1. Paul to Bob (from Dante Controller)
2. Dante controller “tells” Paul to ask Bob for channel1
3. Paul “asks” Bob for channel1 at the latency setting
4. Bob checks that latency setting is correct, if not
5. If this is true, flow is created with Paul’s latency
6. If not flow is created with Bob’s latency
7. 6mbps bandwidth is used between Bob and Paul (in one direction)
8. 0mbps bandwidth is used on cable between switch and computer
9. Paul reports “good” subscription (multicast)
10. Controller updates record (green check)

MULTICAST SUBSCRIPTIONS

- Dante is Unicast by default (for good reason)
- In Unicast – bandwidth is only used between a receiver and a transmitter:
  - When an active subscription is made
  - And on the shortest path between receiver and transmitter
- When a multicast flow is created
  - It doesn't care about a receiver
  - If there is no subscriber – it still uses bandwidth - everywhere (in an unmanaged network)
- Dante CANNOT dynamically switch Unicast to Multicast
  - BUT if a multicast flow is created AFTER unicast subscriptions have already been made, they will switch to use the new multicast flow

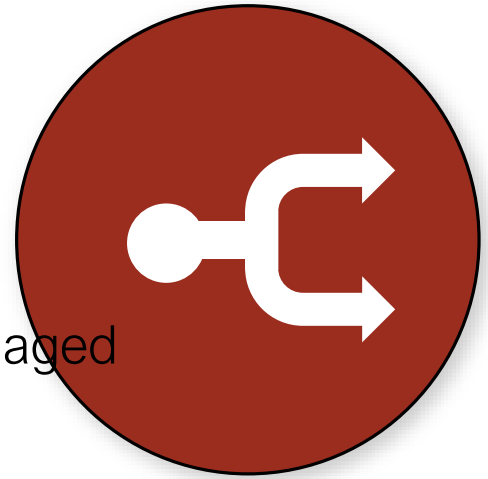

EXAMPLE - IN A DANTE NETWORK

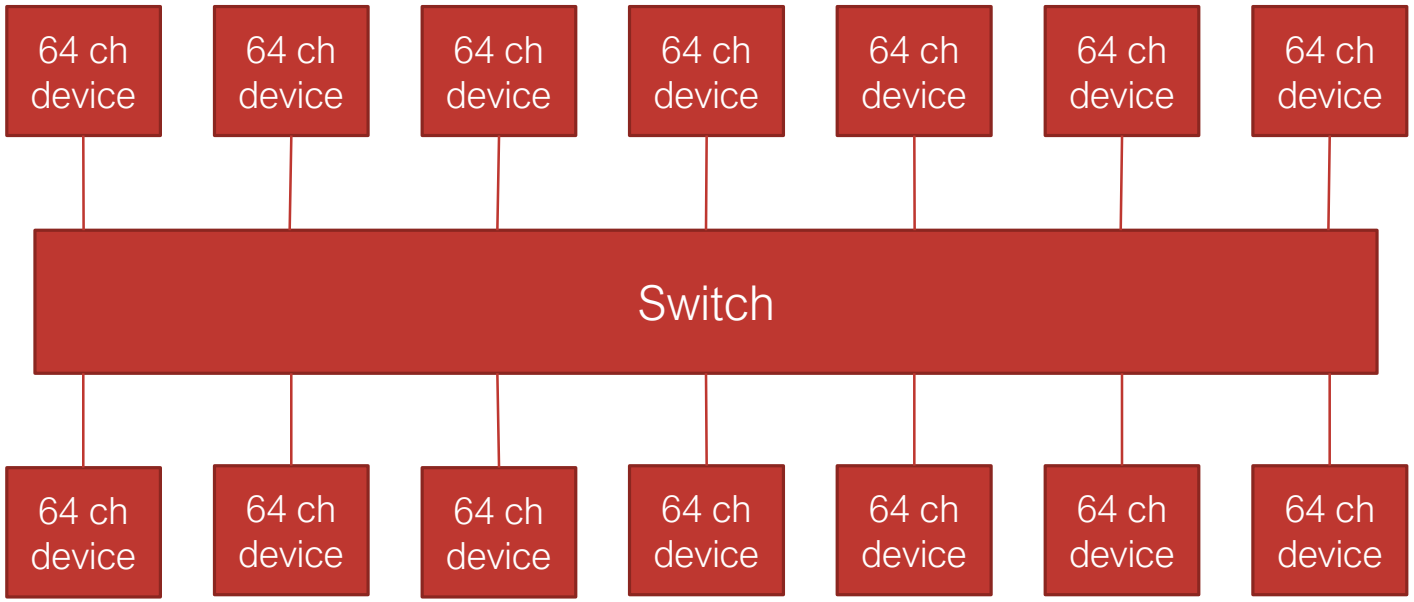

1. 14 Devices, 64 channels each, 32 flows each
2. 32x6 = maximum unicast b/w per node =192mbps per node
3. 192x14 = 2688mbps = more than a gigabit
4. Maximum Rx bandwidth per node is still 192mbps unicast (run out of flows)
5. Minimum backplane bandwidth of switch 28gbps (still more than 10x actual bandwidth)
6. “worst” flow structure possible = 32x 2 channel flows on a device (to still send 64 channels)
7. Only a ”2 way” split achievable unicast
8. IF all devices set to full Multicast – switch would still be fine… Devices would not!
9. 8x12mbps multicast flows = 96mbps x 14 nodes = 1.344gbps = more than each cable can handle! BUT only still 4.8% of the backplane bandwidth
10. Sending 832 channels into a network where they “flood” and the maximum receive possible is 64 channels seems like a huge problem…

LET'S REVISIT OUR ORIGINAL EXAMPLE

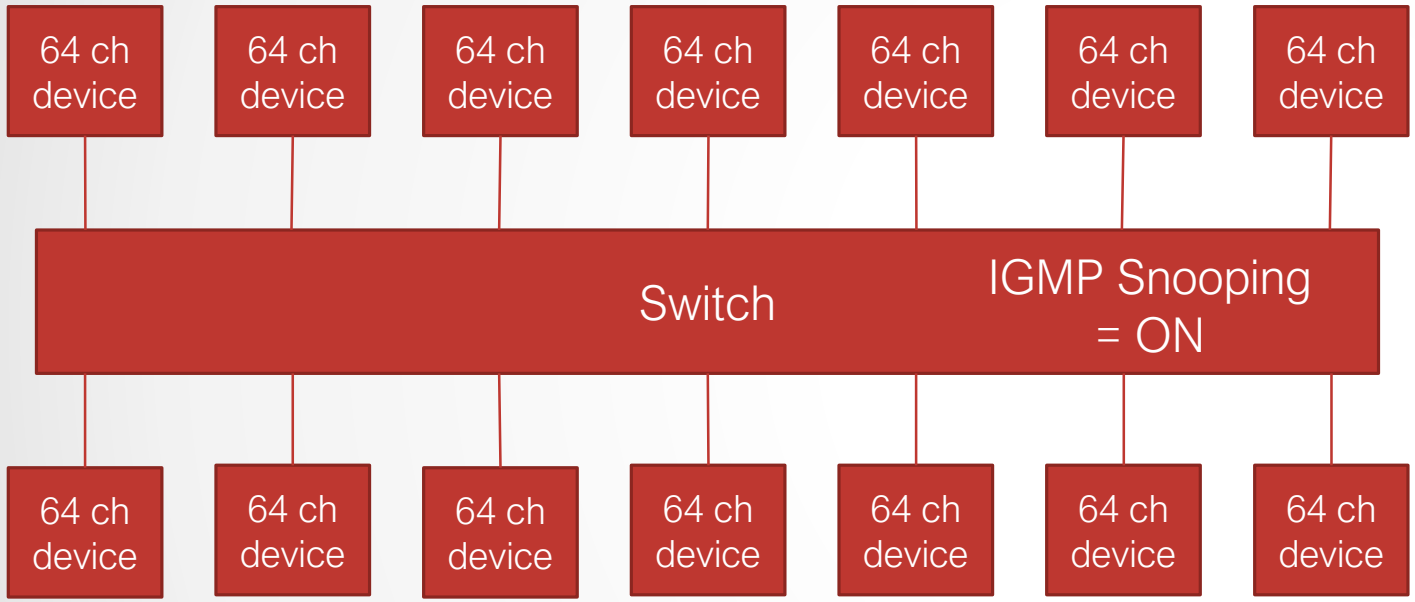

- Remember – with IGMP snooping off and everything multicasting – we were trying to push 1.344gbps down each gigabit cable
- Obviously never going to work
- Turn on IGMP snooping (assume querier is working)
- Each device can only subscribe to 32 flows
- This is a maximum bandwidth of 384mbps per port (and in reality it is never that)
- The backplane bandwidth is unchanged at about 4.8% capacity (no contention)
- This is why multicast management is always a priority over QoS (by orders of magnitude)

Dante by Audinate

DANTE WITH QOS

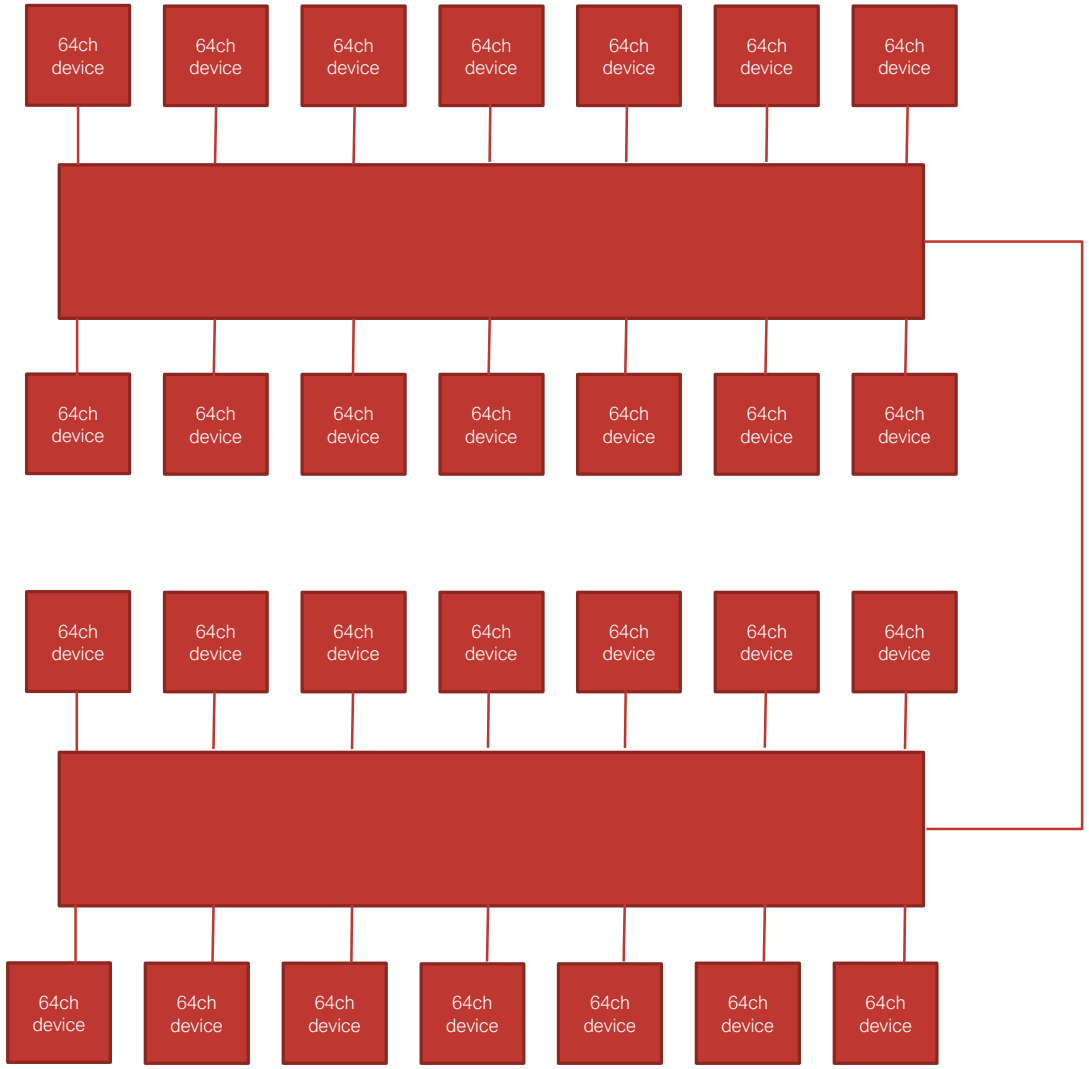

- Assume all 14 devices on top switch are transmitting on a 1:1 mapping to the 14 devices on the bottom switch
- 14x64 = 896 channels
- 16 flows per device @ 6mbps x 14 = 1344mbps
- Simply put – on a gigabit link…
- The Bandwidth has ran out!
- Solution? 10 gigabit link (not ridiculous)
- Problem solved

- All of these devices are transmitting a constant data rate
- QoS will do Nothing here (its not the solution)

TOPOLOGY AND QOS

The previous example shows a bottleneck being created by a sub optimal topology

Luckily there are established methods for designing network topologies

The “three layer hierarchical model” describes the concept

As Switch and router speeds have increased “collapsed core” has become more normal

“Classic” model

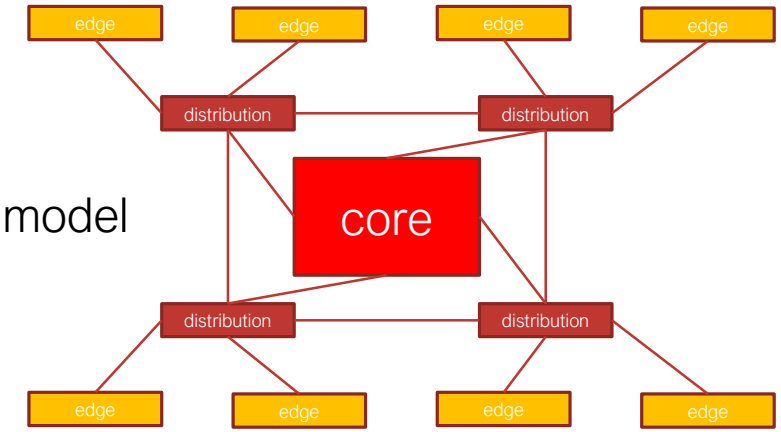

“Collapsed core” model

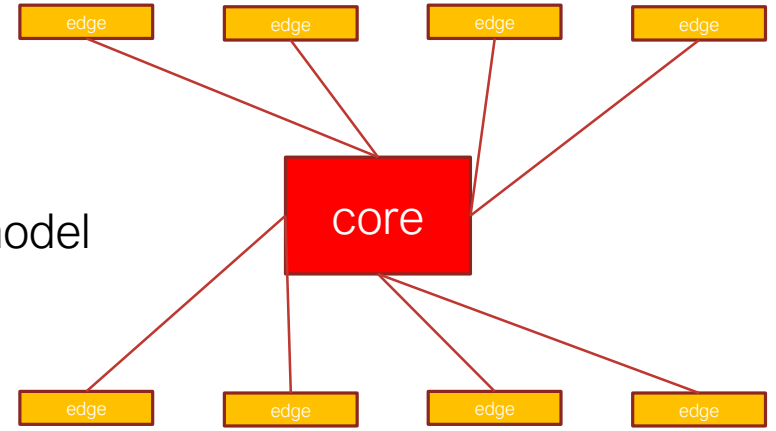

ADVANCED DANTE NETWORKING: SECTION 7

IN THIS SECTION…

Advanced clocking

- How does a network synchronize media clocks?
- What does an external clock do?
- Network synchronization
- Synchronizing time
- Time in PTP

HOW DOES A NETWORK SYNCHRONIZE MEDIA CLOCKS?

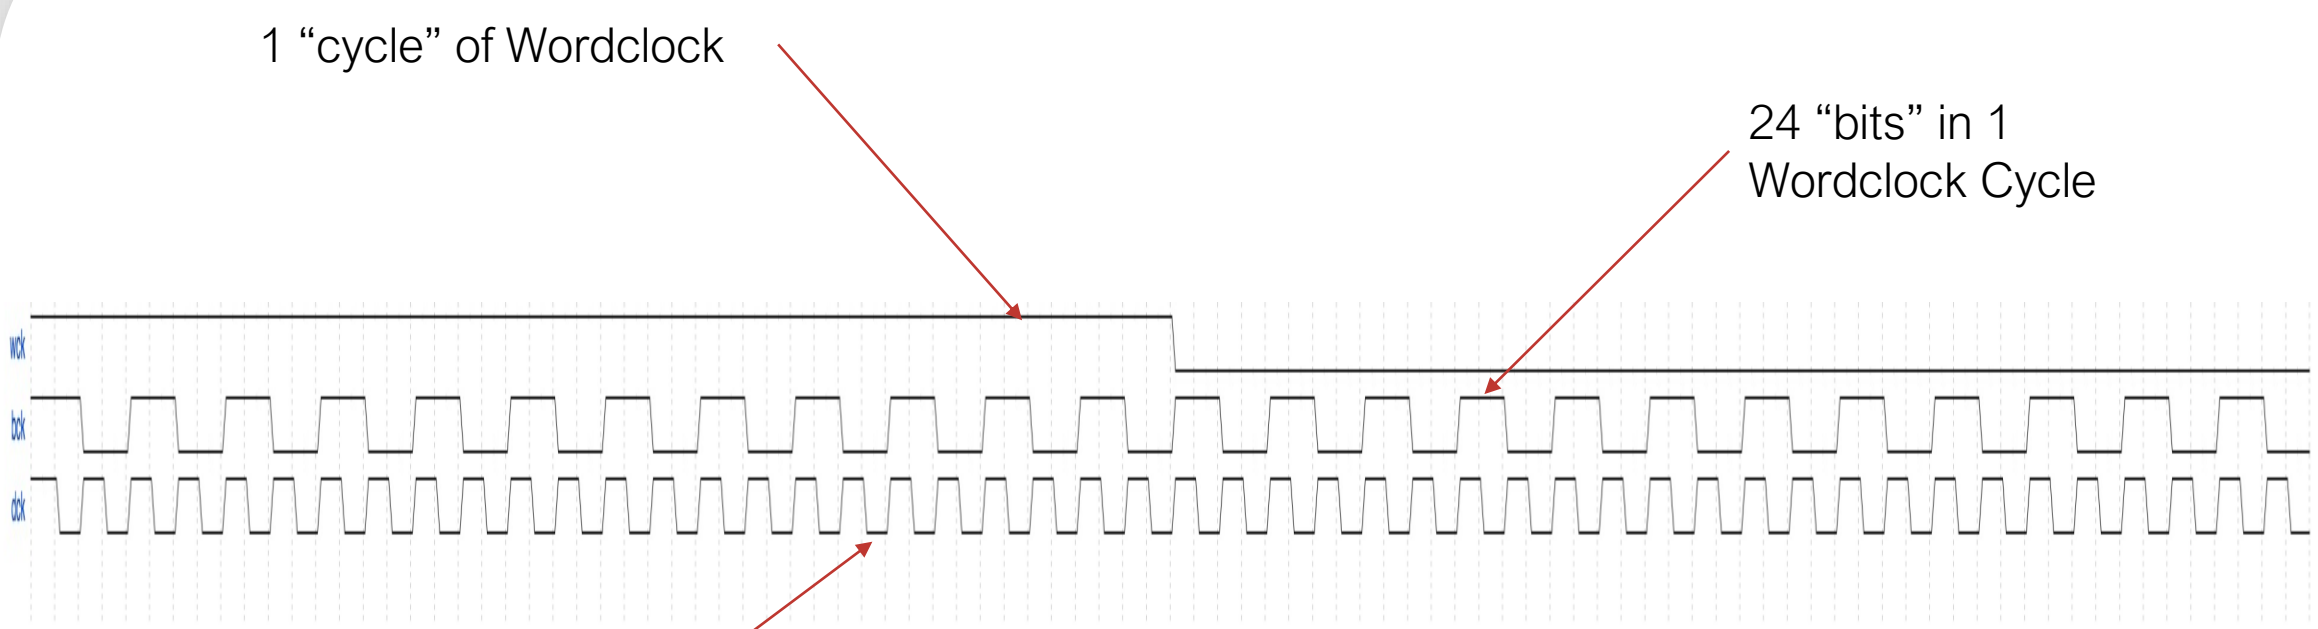

- Simplified example
- “Real” AES3 contains other data frames too
- This is more like a format called LJ I2S

SO WHAT DOES AN "EXTERNAL" CLOCK DO?

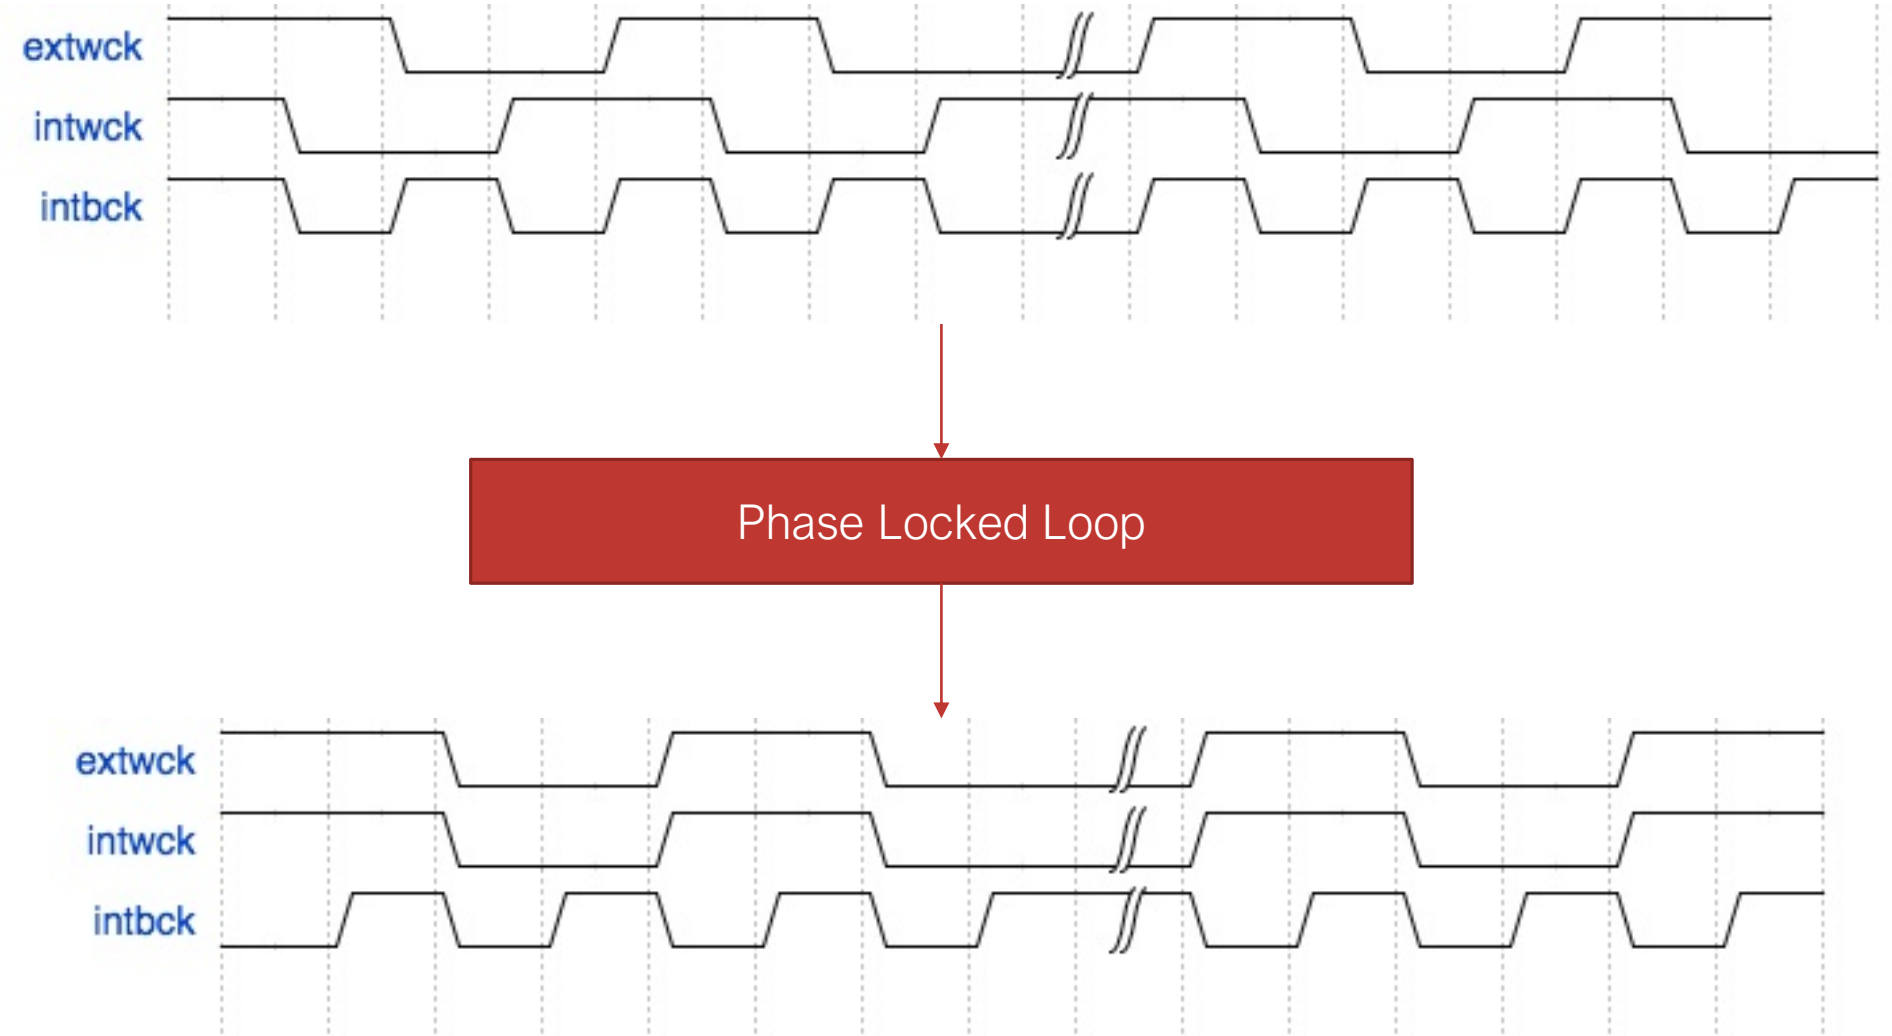

HOW A DISTRIBUTED WORD-CLOCK WORKS

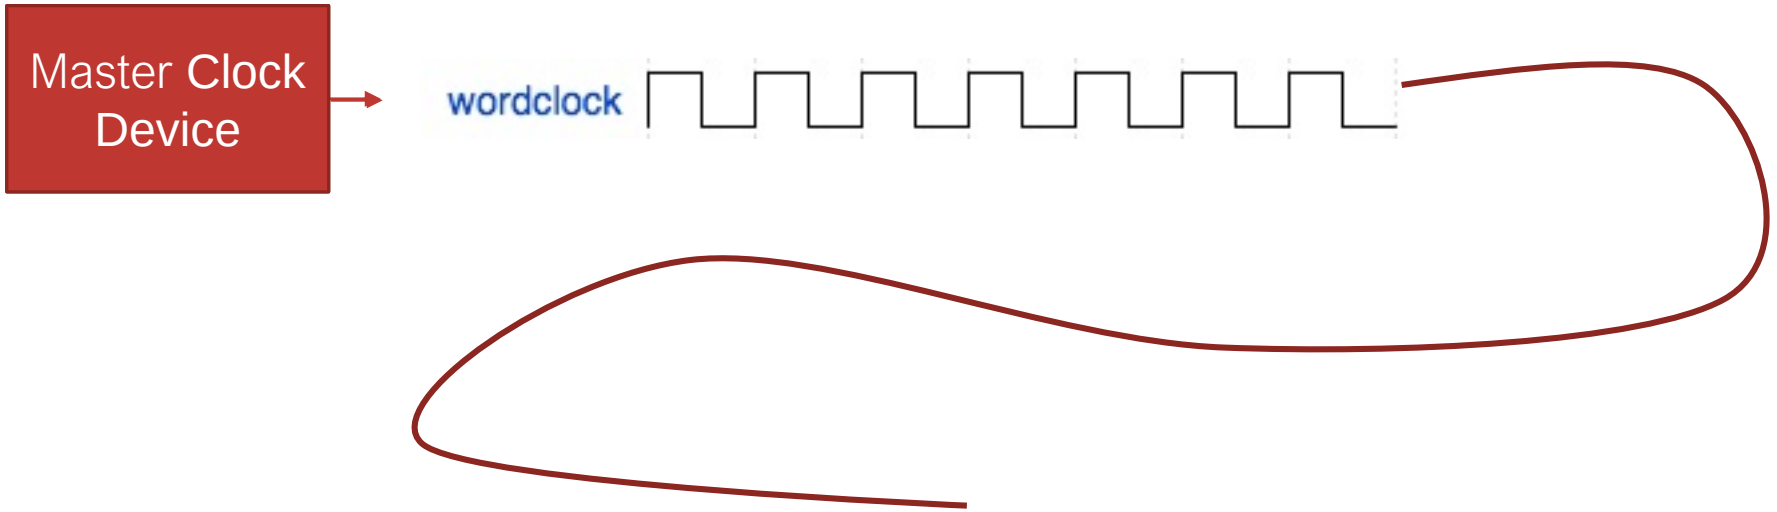

- Speed = Distance/time How far does the signal travel in 1uS?
  - $2.8x10^{8}$ x $1x10^{-6}$ = 280m
- How often do we need to synchronize over this distance?
- In Live sound and Broadcast – a lot of the time (think Golf, Football, Stadium Concerts, Motorsport… etc)
- Dante guarantees microsecond sync at any point in the network…. How?

NETWORK SYNCHRONIZATION

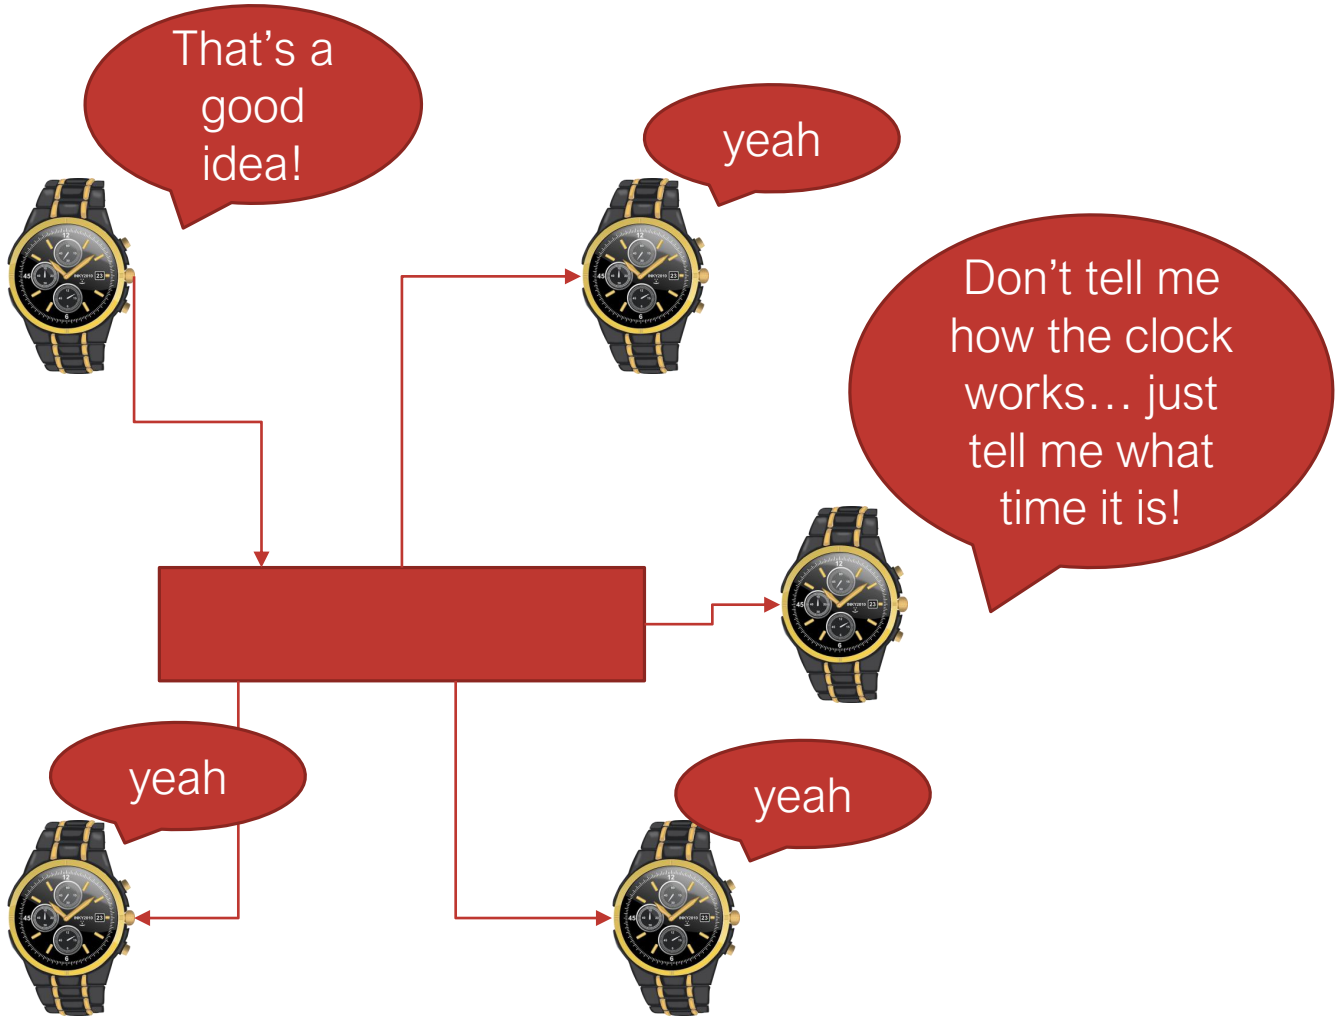

SYNCHRONIZING TIME

- The idea of distributing time over a network started with British Railways
- As rail networks grew – it became necessary to have a timetable
- It was not acceptable for a train to just show up, and leave at arbitrary times – think of how annoying that was!
- How did this problem get solved? (remember – no telephones!)

SYNCHRONIZING TIME

London

Set watch to Big Ben

Put watch on train

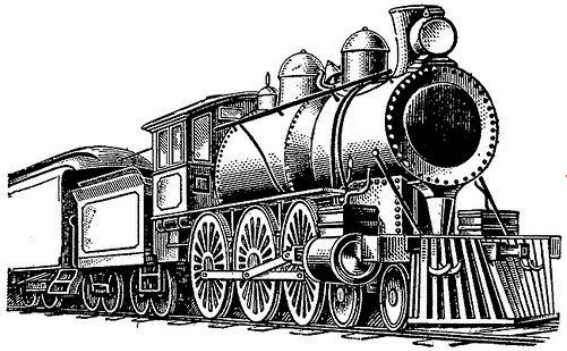

Go to

Set Edinburgh station time to watch

Edinburgh

TIME IN PTP

- Of course – synchronizing in seconds is nowhere near good enough for audio!
- Luckily IEEE1588 aka PTP Precision Time Protocol is a lot better!
- IEEE1588 has a resolution of nanoseconds
- The PTP master device has the reference time
- It transmits this reference time onto the network (think train)

TIME IN PTP

- The problem with a railway system is that the physical clock would have to be loaded and unloaded from a physical train – only one station could synchronize at a time
- Wouldn't it be better if we could simultaneously achieve this?
- Luckily on a network – that's exactly what we do
- Just like the railway – it takes time for the message to propagate across the network – PTP also nicely takes care of that!

PTP MECHANISM

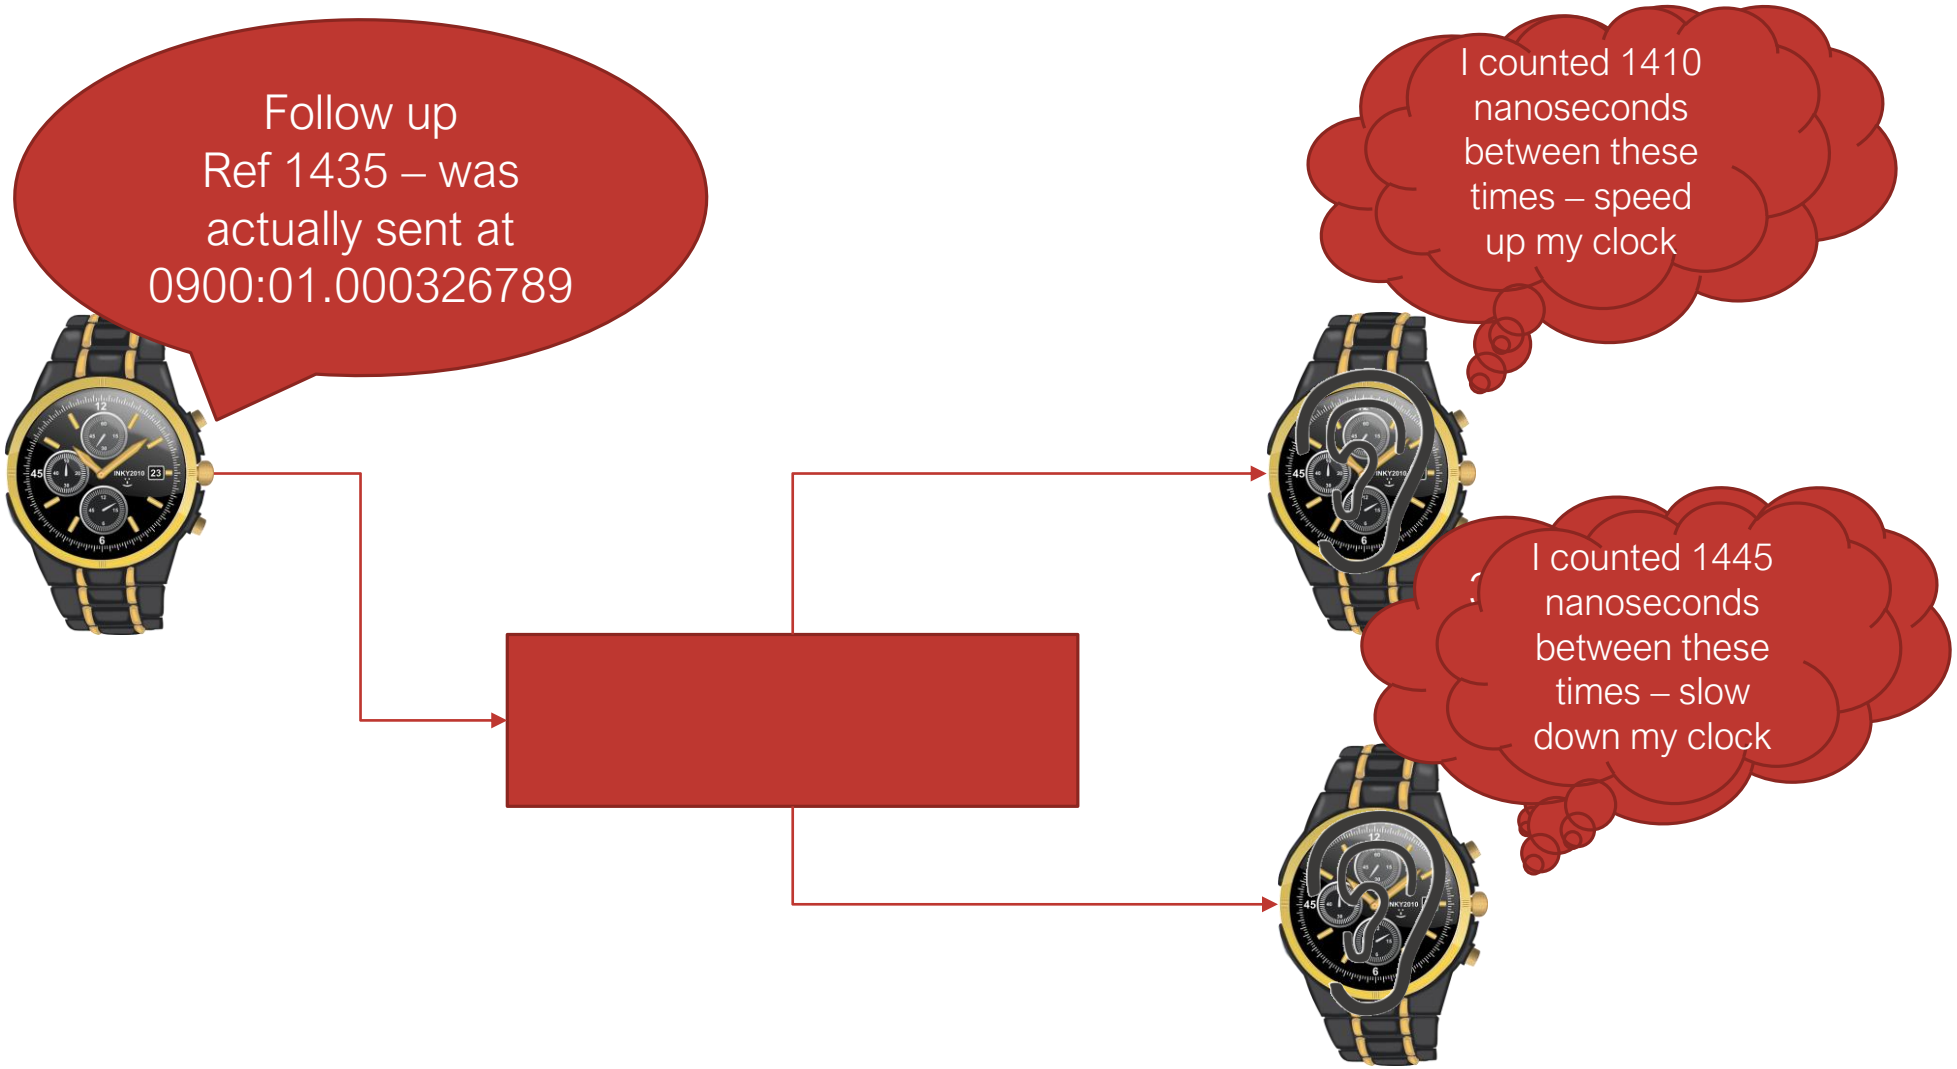

PTP HAS TO WORK ACROSS A NETWORK

- We saw how frames can even slightly potentially “get in the way” of each other
- This is also true for PTP frames (they aren’t magic)
- In reality the slave device needs to “get an average” of the messages it receives from the master to not “jump to conclusions”
- PTP also considers how to be robust in a network

PTP HAS TO WORK ACROSS A NETWORK

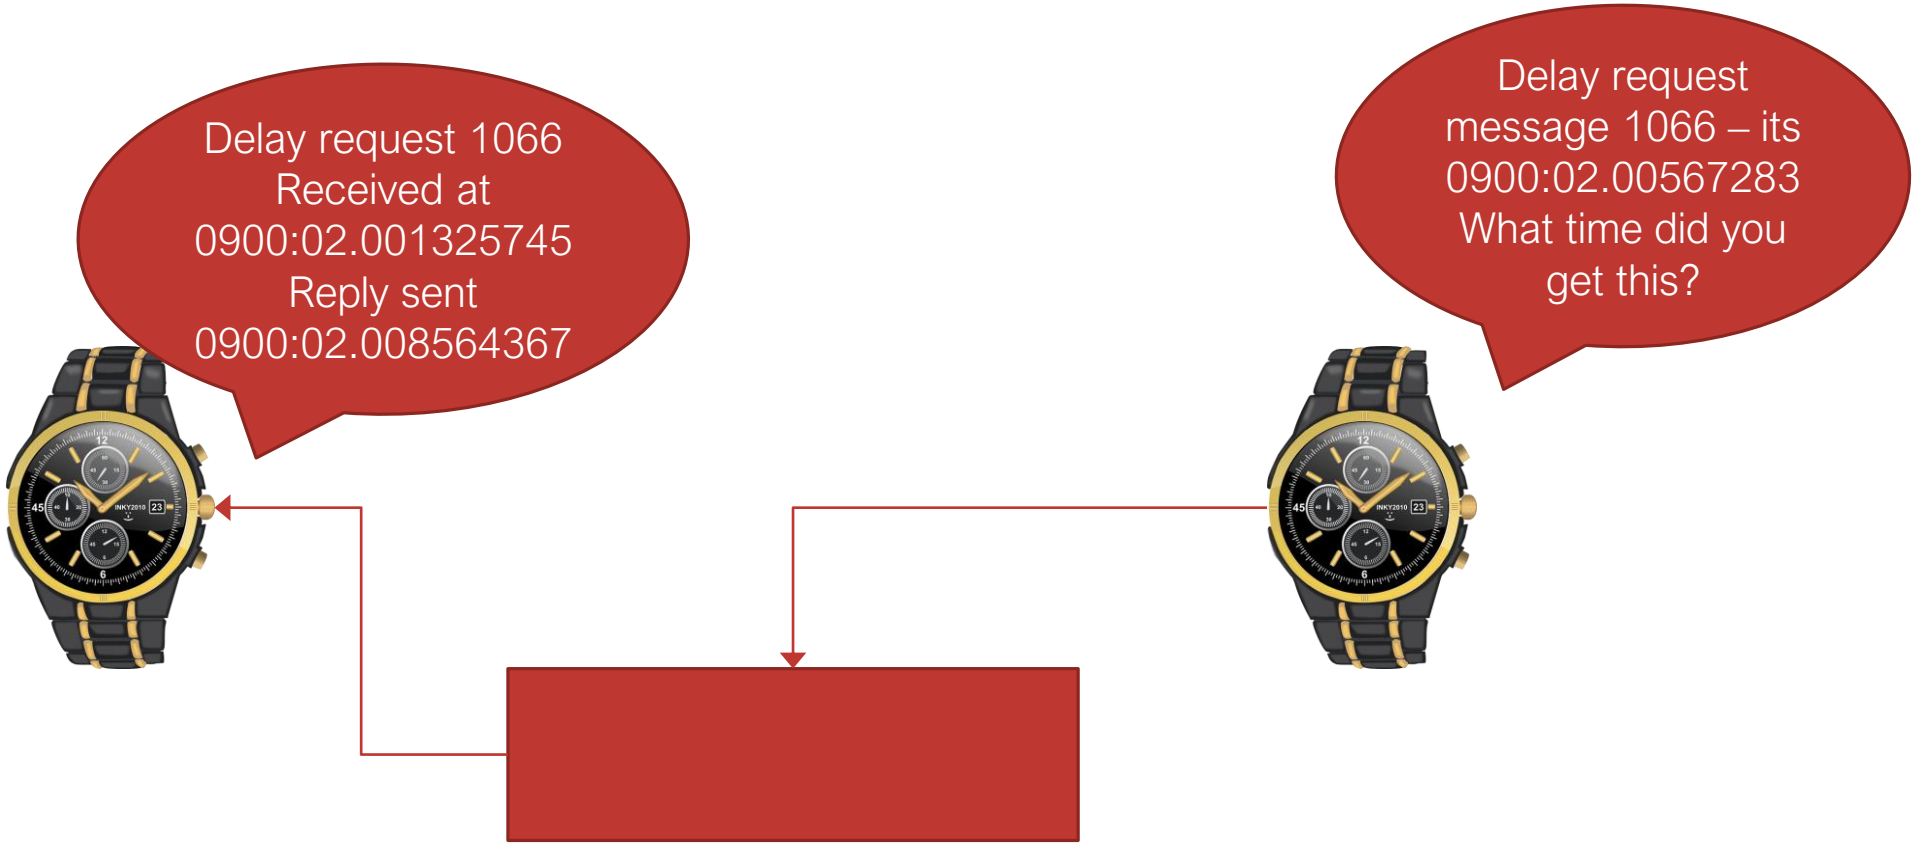

CONTINUING ON…

- The master continues to send sync and follow up messages
- The slaves continues to send delay request messages
- Dante has a “sync interval” of 0.25 seconds – 4 packets per second
- The slaves obtain tight synchronization very quickly

CONTINUING ON…

- Now… if we remove the Master clock some time later
  - The slaves have a very good idea of what the “relative speed” of the master clock was (and can maintain this on their own) … think a speed limit sign on the highway (and assume you follow the rules)
  - The slaves also have a very accurate idea of what the time is at the master device – and will have set their own timepieces accordingly
- This means
  - Any Slave can now easily (and quickly) become the master
  - The subsequent slaves are synchronous enough with the “new master” that their audio clock is unaffected
  - The synchronization of the system is not dependent on a single permanent fixed master
  - Time is Time – I can “chop” it into however many parts (Hz) I like
  - Audio and Video clocks are simply the inverse of time – 1Hz = 1S^-1 after all

ADVANCED DANTE NETWORKING: SECTION 8

IN THIS SECTION…

Troubleshooting

- Fundamentals of troubleshooting
- Dante port addresses
- Network ports
- IP addresses
- MAC addresses
- Common symptoms and causes
- Talking to an IT department

TROUBLESHOOTING

- Networks are just about connecting devices together
  - Understanding how connections can be “broken” is the key to troubleshooting
- **Recap** – Dante puts the following “traffic” into the network
  - Timing – PTP IEEE1588
  - Audio Data
  - Commands for creating routes and controlling Dante functions
  - Discovery services (so that commands can be sent to devices)

- Just because you cannot “see it” doesn’t mean it cannot be seen!
  - Think about applications attaching to the network stack using “port addresses”
  - Many Operating systems use software “firewalls” to prevent certain applications accessing the network stack
  - If you can “see” a master clock (identified at top of Dante controller routing screen) but no devices, or device names with no information, 99% of the time – the firewall on the machine running Dante controller is blocking something

TROUBLESHOOTING - DANTE “PORT” ADDRESSES

Traffic used by Dante is as follows:

mDNS	224.0.0.251:5353
Control and Monitoring	224.0.0.230 – 232:8007-8706
PTP	224.0.1.129 – 132:319-320
Multicast Audio	239.255.0.0/16:4321
Unicast Audio	RX Unicast IP:14336 – 14600
AES67 Multicast Audio	239.XX.0.0/16:5004

- All Dante traffic is UDP/IP
- This means that if any traffic to/from these ports and IP addresses is blocked, then the ”wire” carrying that particular service can be considered “cut!”

NETWORK PORTS – 2 WAYS TO CUT THE WIRE

- A computer's operating system may use a "firewall"
  - Look at which applications within the Operating System are trying to access the network
  - Apply rules to these connections (permit or deny)
  - Look at which destination IP addresses (and ports) are being used – can permit/deny on this basis
  - Can determine whether a network is "private" or "public" and decide to permit/deny on this basis
  - Can look at how much data has been transferred to/from a specific port and/or - IP address and decide to trigger rules on a "threshold"

- Firewall Appliances are simply computers in the network at gateway points – they work in a similar way to an OS Firewall

NETWORK PORTS – 2 WAYS TO CUT THE WIRE

- The Network Switch Fabric can have ACLs applied (Access Control Lists)
  - ACLs are very powerful tools for filtering traffic in the network
  - Many advanced IT Networks will be applying ACLs
  - Normally IT departments will not wish to reveal or discuss these (potential security concern)

- It is reasonable to expect that the required service be provisioned with the required resources – at a Port level this is defined in previous slides

IP ADDRESSES

- 2 kinds of IP addresses to be concerned about:
  - In Dante, Multicast IP addresses only work in a single IP subnet
  - All communication on a Dante network spanning multiple subnets is unicast (device IP to device IP)
- Multicast IP addresses
  - Some switches block IP multicast by default (this will prevent the following functioning properly):
    - Dante Discovery in the local area network
    - Dante Clock (PTP is sent multicast as standard)
    - Multicast Dante Audio

IP ADDRESSES

- 2 kinds of IP addresses to be concerned about:
  - In Dante, Multicast IP addresses only work in a single IP subnet
  - All communication on a Dante network spanning multiple subnets is unicast (device IP to device IP)
- Unicast IP addresses
  - For a unicast communication to be successful:
    - In a LAN – both devices must be within the same IP subnet
    - In a Routed network – the LANs containing the devices to be connected must be joined together by an IP Router
    - The IP subnet configuration of the Router must be correct (Interfaces correctly addressed)
    - Routes must exist between the IP subnets in question
    - More advanced routers can filter traffic using ACLs just like a switch or a firewall

MAC ADDRESSES

- MAC addresses rarely cause any issues (unless an ACL or security appliance is configured with a deny rule)
- At Layer 2 – ensuring all interfaces on a switch are in the expected VLAN is the most common configuration issue
- And Finally, Layer 1
  - Is the wire physically cut?
  - Is the interface open? (managed switches can have their interfaces shutdown using a software command)
  - If using fiber – is the signal integrity good (fiber ends cleaned? Interface dust free?)
  - Use respectable test equipment to certify cables (cat 5/6/7 etc) and check signal integrity and loss of a fiber assembly

COMMON SYMPTOMS AND CAUSES

Device Name shows in Dante Controller – no + visible to expand channels & status view missing Data

- Check Firewall settings on your computer's Operating System
- Check that the interface being used by Dante controller on your computer is in the same IP subnet as the Dante devices that you wish to control

COMMON SYMPTOMS AND CAUSES

Clock gives unlock/lock warning

- Possible Multicast block – check the clock histogram in Dante controller to confirm
- Unicast Delay Requests can be a quick tool to test if this is the case
- Dante clocks can run for a surprising amount of time before falling out of sync badly enough to affect audio… hours

COMMON SYMPTOMS AND CAUSES

Multiple Master Clocks shown

- Devices “cannot hear” multicast sync messages from other devices – assume they are master
- Normally caused by blocked multicast traffic in network

TALKING TO AN IT DEPARTMENT

- This course covers the requirements of a Dante Audio over IP network in some detail
- Configuration details for a particular IT infrastructure vendor's equipment are specific and unique
- HOWEVER – the concepts covered in this course are universal in IP networking (the terminology may be subtly different on different infrastructure vendors)
  - Understanding how IP connections are made (ports, IP addresses, Routing, Switching, Multicast, Broadcast etc) helps explain requirements
  - Understanding bandwidth requirements is useful
  - Understanding that QoS rules can be applied to optimize performance
  - Understanding that IGMP can be used to mitigate against the negative perception of multicast

TALKING TO AN IT DEPARTMENT

- Networking is about making a lot of parts of an unique “jigsaw” work together
- This course explains the requirements and performance of the “Dante piece” of the jigsaw
- Remember – an IT department have to make many pieces fit together – remain patient, it can and will be made to fit together nicely.

NEXT STEPS

TAKE THE LEVEL 3 TEST

http://www.audinate.com/certify

- Create an Audinate account if you don't have one
- Login to your account
- Take Level 3 test
- Certificate is automatically generated

THANK YOU