Data Networks in Audio-Visual and Lighting Part 2 – Todays reality

What you probably "know" a recap.

Traditionally, at this point I should launch off on an "explainer" of the OSI model. I've done this many times, and to be honest, it took me years to get a decent grasp of it. Its one of those things that just takes time. This is in no way to dismiss it (far from it, it is foundational). However, in today's AVL reality, it makes close to no sense! This is because it is (rarely) fully implemented, and for a user, even when it is, it loses practical application rapidly.

So, we have a piece of equipment with some kind of network connection on it. We have some cable (one assumes you have the "right" cable. And we have some other mysterious boxes that we normally use as breakin/breakouts that get called "switches" or Routers or some other differentiating name,

We plug devices into these boxes in the middle. We know which devices can be controlled by plugging a computer into said same box, we know which software to open, and we make it work.

Some folks talk about "VLANS" or "subnets" all we want to know is how to plug it in, and just get on with the day.

We found out that if we connect our computer to the "wrong" VLAN, we can't "see" the devices we want to see. Sometimes, we need to dive in and change this thing called an IP address, and this suddenly makes stuff work… there you go… that's networking isn't it?

Practically yes, there is a lot of truth in that perspective. For the vast majority of entertainment networks, this is sufficient. It only becomes an "issue" when things get larger (now the "boring philosophy of section 1 might make some more sense).

We might touch on a little bit of (practical) OSI – just the "layers" that make sense and are useful today. Consider that the OSI model has 7 layers (consider Dante's Inferno and the seven layers of hell!). But, I promise to keep it relevant. In "reality" we actually look at OSI as a journey from above Layer 7 down to layer 1, and back up again. This is the same as the layer above layer 7 being a singer singing into a microphone… its nice, but not our thing today. We want to just get on with the day and do something interesting. How can someone be so strange as to find this stuff interesting? As one who does find it a little "interesting", I wont bore you with why I think you should be interested. I also appreciate the misunderstanding that I may look frustrated when products don't necessarily conform to my personal view of "good enough" (some might say pointless perfection!).

At layer 1 we have the physical layer – that's cabling and/or wireless. Easy- plug it in!

Layer 2 is the "switches" these somewhat ensure that messages only use the shortest physical path to get from device A to device B.

Layer 3 is where we have IP addresses. IP addresses are considered a bit pointless – see AVB in its 1722 (most common) version. This has some merit. We rarely have so many devices that we need

consider “segmentation” for a particular service. It is rare in entertainment to ever connect to an IP router (also layer 3). So why do we care?

Simply put Layers 3 through 7 don’t normally need touching. They exist, and (to a varying degree) are implemented in equipment. Why? This is because if we want to use a “normal” operating system (Windows or Mac OS) we have to talk to it through all 7 layers. This may seem overkill (and yes it could be) but from a software engineers’ perspective, it is very much the easiest way to make a solid connection to a lot of “stuff”.

The main (and common) thing I see is the connection of a forest of network cables from a single laptop/desktop. This scares me (a lot) in effect, this is making the Operating system believe that it is an IP router! Please try to avoid this!!! I do appreciate that deploying systems in a temporary environment has its challenges (I’ve done a few shows as you might notice). My preferred mitigation is to have multiple Virtual machines running desktop operating systems in a server environment. This means that the “router” can be a designed for purpose device. We can connect a desktop OS to just one VLAN and use a remote desktop server to allow remote control in the safest way.

So why do we have VLANs Are they important?

VLANS or ieee802.1q to give it its full name, exist to segment “broadcast domains”. Whilst IP at layer 3 is also used to achieve this higher up the stack. 802.1q is a practical solution to not using many small switches. It cuts down cabling significantly, and reduces the number of power connectors. It is tidier, and as we all know, therefore much safer and reliable! VLANs split a “big” switch into a number of small individual switches by maintaining groupings of interfaces by “VLAN”. This is what prevents a box in vlan a from directly communicating with a box in vlan b (if just a switch is used). This (for example) allows us to stop the lighting or video from “interfering with the sound, and vise versa. The somewhat implausible notion that “the Internet” might break into your network and take your show out is sadly a bit of a myth. Simply, the IPv4 Public Internet is entirely Unicast – this means that a connection cannot be randomly made, and random traffic sent spewing into your amplifiers or mixing consoles! Anyone who has relied on this as a defence is either poorly informed, or trying to take advantage of others that they assume are poorly informed!

Did I miss something, they still don’t seem important?

No you didn’t miss anything. I have avoided (thus far) the reason why we consider collision domains and broadcast domains. Somewhat depressingly that time has come.

1. Collision Domains – in a modern network, we can somewhat assume that these have ceased to exist. Likely prior to about 2003 these horrid devices called “hubs” were in abundance. “Hubs” are layer 1 devices (they do look like switches). However the difference between an hub and a switch is that a switch contains a “MAC address table” whereas an hub doesn’t. the MAC address table reconciles a MAC address (layer 2) to a physical connection (layer 1). This is the thing that stops unicast traffic “flooding”. As long as the device communicates its MAC address in the standard way, the switch will “learn” it, and populate its MAC address table. From then, any message at layer 2 (called a frame) that is destined for a particular MAC address will only be sent via the corresponding physical connection! By contrast an hub simply floods the traffic. If 2 messages coincide, they will be mixed together, and their crc (on the end of the frame by convention) will be “wrong” and the receiver will drop it. In a busy

network, collisions can cause the whole thing to grind to a halt! Switches get around this, by storing the message if the interface is busy, and forwarding it when there is space.

2. Broadcast domains. Unfortunately, at layer 2 (and in ipv4) broadcast traffic exists. This is characterised by a destination mac of ff:ff:ff:ff:ff:ff this means that the switches will flood the message everywhere. The receiving devices may “filter” it, but the damage (in terms of capacity on the cable) is already done. Ipv4 does use a fair bit of broadcast traffic. This is why we have vlans and to an extent IP subnets. Constraining the number of devices that can generate broadcast traffic means we keep the network functioning. Typically for “random” devices, it is considered a broadcast domain of around 8 bits is “safe” or 254 functional IP addresses. Mainly due to history, the “other” kind of traffic (multicast) took a very long time to mature. To all intents and purposes, multicast in IPv4 is broken. Whilst ipv4 multicast can be passed through a Router, the complexity of doing this is astounding. Multicast feels quite obvious in entertainment – anything that its “one to many” like we do all the time, is obviously beneficial. “one to one” is called unicast (and is the default, some might say only option in ipv4) “broadcast” is “one to all”, in the past it may have been sensible, however multicast in “source specific” mode rather than multicast in “any-source” mode is less useful. It goes without saying that whilst “source specific" mode has come into ipv4 lately, it was there from the beginning in ipv6.

In short – VLANS break up broadcast domains – this is useful. If you want to connect two vlans – you use a router at layer 3 and IP addresses.

For day to day networks, routing is not a thing for users or technicians. Beyond this we get into architectural concepts, which I shall cover in other articles.