




ARCHITECTS OF VIRTUALIZED MEDIA PRODUCTION

Architecture & control

The two keys to IP infrastructure success in broadcast

A Nevion Whitepaper



Introduction

The adoption of IP technology across the whole broadcast workflow is now well underway. IP has already been used for many years for the purposes of broadcast contribution over wide area networks (WANs). The technology is now also beginning to be used in local area network (LAN) environments for transporting broadcast signals within studio- and campus facilities.

For many broadcasters, the move to IP is driven by the need for increased productivity, i.e. being able to do more with the same or less. Over time, IP technology is expected to be cheaper than broadcast specific baseband technology. This is primarily because it can easily handle any existing and new video and audio technology (HD, 4K/UHD, HDR, 8K, etc) since everything is transported as data (IP packets). IP also offers the prospect of media, control and data all being carried by the same network providing savings through multiple usage (economies of scope) that do not exist in baseband.

In many cases though, broadcasters are initially considering a like-for-like network replacement of baseband with IP, which, for the time being, still implies higher initial cost, not least because of the need to convert the SDI output of existing equipment to IP (until all broadcast equipment becomes IP capable).

There is much more to IP than simply mimicking existing baseband networks though. The seminal VRT/EBU LiveIP project (the first practical demonstration of using a multi-vendor all IP environment for live production, back in 2015-2016) made the point very eloquently: IP enables workflows to be “remote, shared and automated”.

IP brings the opportunity to harmonize local and long-distance media networks around a single technology – so-called IP LAN/WAN (local and wide area network) convergence. This means that it becomes much easier to share equipment, studios and control rooms, and even production staff, across locations – bringing further savings and much greater production flexibility than could be possible with existing technology. This is the opportunity broadcasters need to seize.

“IP brings the opportunity to harmonize local and long-distance media networks around a single technology”



To do this, broadcasters need to consider carefully how they design and architect their IP infrastructure. If they focus on the products, prioritizing features and cost before anything else, broadcasters may find themselves 'locked out' of the real potential that can be delivered through IP networks and find their IP network unable to cope with growth and changing needs.

If broadcasters are to get the most out of their IP infrastructure investment, they must prioritize two things: how the IP network is designed and architected from the ground up, and how the new network is controlled.

Keeping broadcast in mind

Implementing an IP-based infrastructure is a huge investment for any broadcaster, not just from a financial point of view but also logistically and educationally. Most broadcasters taking the leap towards IP have previously relied upon legacy baseband equipment for many years, and so naturally very few of them have any expertise in IP and how to best use it to their advantage.

To seek advice and consultation, many broadcasters will turn to IP switch vendors, who are (unsurprisingly) extremely knowledgeable about IP technology. But herein lies another problem: while broadcasters lack the knowledge of IP vendors, IP vendors lack the knowledge of broadcasting and its unique challenges. For example, broadcasting needs many relatively small audio signals and a few large video signals to be transmitted to multiple destinations in as close to real-time as possible. Few industries have such requirements in terms of number and size of data flows, combined with ultra-low latency, and zero tolerance for data loss.

Therefore, rather than taking advice directly from the IP switch vendor's guidebook, broadcasters need to insist on an approach that takes the special needs of broadcasting

"The initial focus should be on getting the network architecture and the control right for broadcast applications and workflows."

into account, while at the same time ensuring that standard IP equipment can still be used. For that reason, the initial focus should be on getting the network architecture and the control right for broadcast applications and workflows.

Architecture - building the foundations

As IP transport and switching technology is adopted more commonly across LAN environments like studios and campuses, there are many different options available to broadcasters when it comes to building the media networks. Many forces are at play here, and of course there is the strong temptation for broadcasters to stick to what they know and mimic existing baseband set-ups.

Getting the IP network architecture right from the outset is fundamental to a successful implementation. The ways in which equipment and components are connected to the main routers can make all the difference in ensuring everything can communicate and operate seamlessly without any issues.

Network architecture cannot be totally dissociated from control though, and to satisfy broadcast requirements, it is important to know how signals flow through the network. If the architecture is relatively simple, cabling may make that knowledge explicit, in the same way as it does in baseband. But for more advanced architectures, control is needed, and indeed some methods of control may not be adequate enough to meet the expectations.

Three main architecture layouts are typically used in broadcasting and other industries. While each has its own unique set of pros and cons, there is one layout that is significantly better suited to operations within a broadcast environment.

“It is important to know how signals flow through the network”



Centralized star network

When it comes to architecture, the tendency for most broadcasters is to adopt what is known as a centralized star network (or “monolithic switch”, as it’s sometimes known).

This star network is very consistent with the traditional baseband architecture - all connections transit through a large IP router (usually with another large router that performs as a backup) that can be located in the master control room (MCR). This approach makes it very easy to route signals across the infrastructure, as all endpoints are connected to a single device, and of course it’s a layout that many broadcasters are already familiar with.

The main disadvantage with the centralized star approach, however, is that there is no signal aggregation at the edge: everything needs to travel to the central router. This has several consequences.

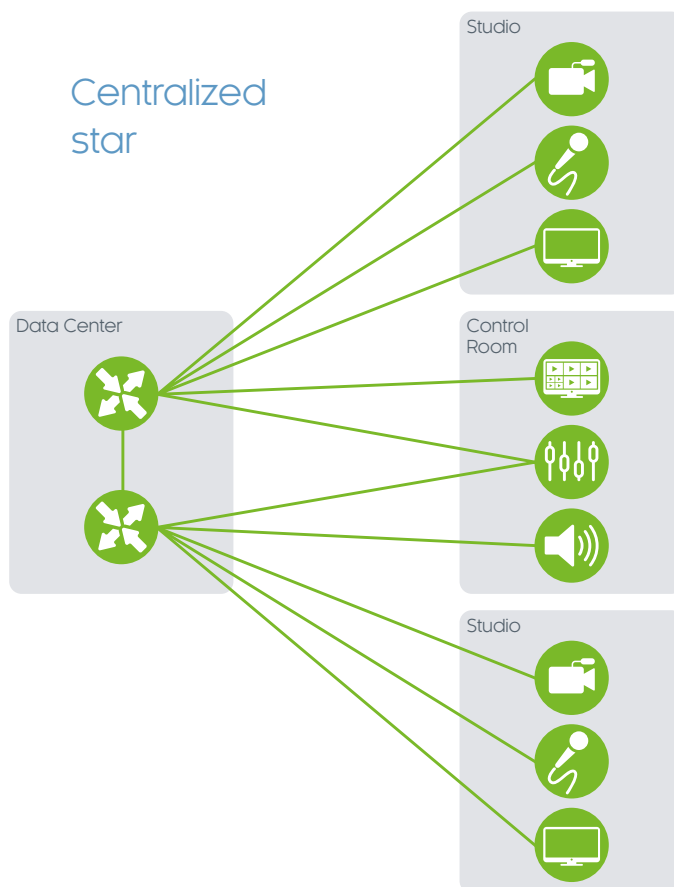
The first issue is that fibers need to be laid to connect every single device to the central router. This is expensive, because of complex fiber management and poor utilization of the fiber infrastructure.

Another issue is scalability. Like traditional baseband networks, the size of the central router is a potential bottleneck. To cater for the highest demand and anticipate future needs, broadcasters have to purchase an oversized router from the very beginning of the project. Despite the best plans, capacity is often reached sooner than anticipated (not least because of the adoption of new standards like 4K, 8K and HDR with ever higher bandwidth demands). At that point the central router needs to be replaced with a bigger and better version.

Scalability is not just a theoretical issue: some broadcasters who initially adopted the IP-based star-architecture are already finding that their network cannot cope with expanding needs, and are having to spend heavily to upgrade.

A further disadvantage is that every connected device occupies one expensive high-bandwidth port on the central router, regardless of its actual bandwidth requirement. This makes the cost-per-port for low bandwidth devices very high.

And of course, the lack of aggregation means redundancy needs to be handled by the edge devices. This requires devices to



have two connections to the central router, or one connection to each of the main and backup switches, which, aside from the extra cost, may not actually be possible in some cases, such as microphones or speakers that have just one connector.

Finally, a star network architecture is not inherently suitable for treating remote locations as extensions of the main location, as it assumes that all traffic will transit through the central router. This means the true potential of LAN/WAN convergence, sharing resources across locations, is virtually impossible to realize.

To summarize, star networks may appear at first to be simple and convenient, but in practice are highly limiting, potentially vulnerable and likely to be very expensive over time.

“A star network is very consistent with the traditional baseband architecture - all connections transit through a large IP router”



Spine-leaf

While the centralized star layout above might be advantageous for those looking for something that is simple and familiar, the negatives clearly outweigh the positives. This is why, in most industries, IP networks are typically distributed. In other words, the network relies upon a number of small to medium-sized routers and aggregation, rather than a single large centralized router.

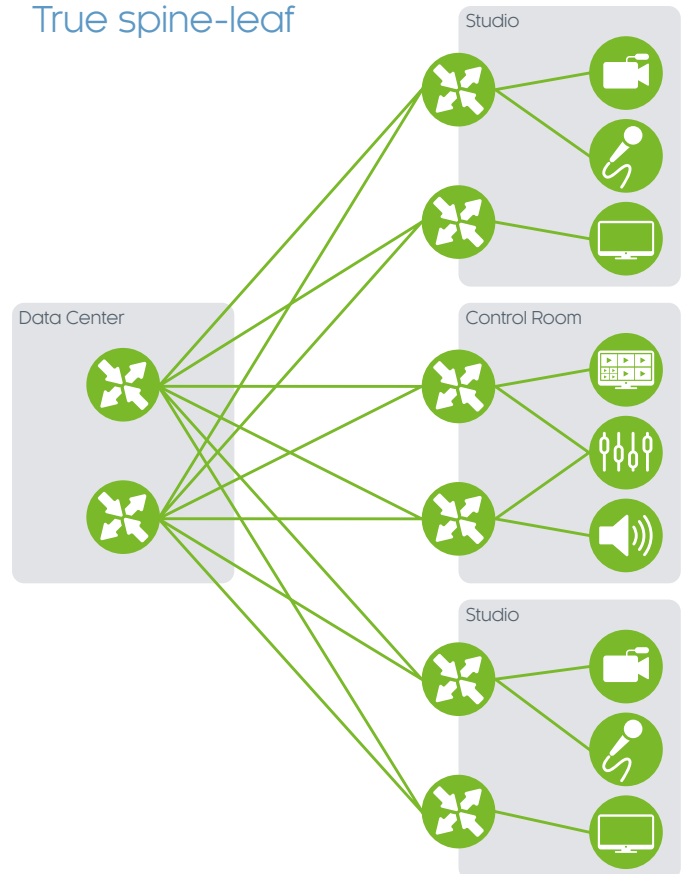
An established distributed IP network approach is one used for large data centers, and is commonly known as a spine-leaf architecture. This involves two or more routers at the core (spine) and other smaller routers at the edge (leaf). The leaves take care of aggregation and connectivity, while the spines take care of the inter-area connectivity.

A spine-leaf architecture typically includes at least two spines (a single switch could be a potential point of failure in the network). Each leaf then has at least one connection to each spine, which accounts for any potential issues regarding load distribution and redundancy over the core routers.

“In most industries, IP networks are typically distributed.”



True spine-leaf



As anyone can see by comparing the spine-leaf diagram with the centralized star diagram, the layout looks a lot more complicated to implement, with many more connections running from one point to another. But this layout offers plenty of advantages.

By connecting all the equipment in each area to aggregating leaf routers, and then connecting these to the main routers, broadcasters can dramatically reduce the number of connections going directly to the main routers, leading to simplified fiber management. For example, if a broadcaster had 15 cameras in studio one, instead of having 15 different connections going to the main router, it could connect all the cameras to the aggregator router and use a single fiber connection to feed all this back to the main routers. This reduction in the number of fibers needed between the leaves and the spine also takes advantage of the exponential increase in Ethernet bandwidth becoming available now – 100G today, 200/400G soon.

Fewer fibers also mean fewer ports are needed on the central router(s), and a much more effective cost-per-port, especially for low bandwidth devices such as microphones and speakers.

The spine-leaf architecture makes it easy to build redundancy into the network for all devices, and at a much lower cost: all of the main connections to the data center are also duplicated, which provides a high level of resilience. It also ensures that no endpoints can be blocked by the loss of a single spine.

This approach also provides optimal flexibility and scalability - if an additional ten cameras need to be added to the network, for example, broadcasters can simply add another aggregator router to accommodate - which therefore makes growth far more organic.

If big infrastructure changes result in a lack of ports on the spine, then additional spine switches can be added to the solution. Conversely, if the initial needs for the network are relatively small, there is no need to build huge over-capacity “just-in-case”. Capacity can be added over time, when needed. The spine-leaf architecture is really about “thinking big, starting small”.

This true spine-leaf architecture is the model that Nevision would recommend to any broadcaster looking to move towards an IP-based infrastructure. While the layout is more complex, it is an extremely scalable, resilient and high-performance network structure that is perfectly suited to the needs of broadcasters.

“The spine-leaf approach provides optimal flexibility and scalability”



Dual star (pseudo spine-leaf)

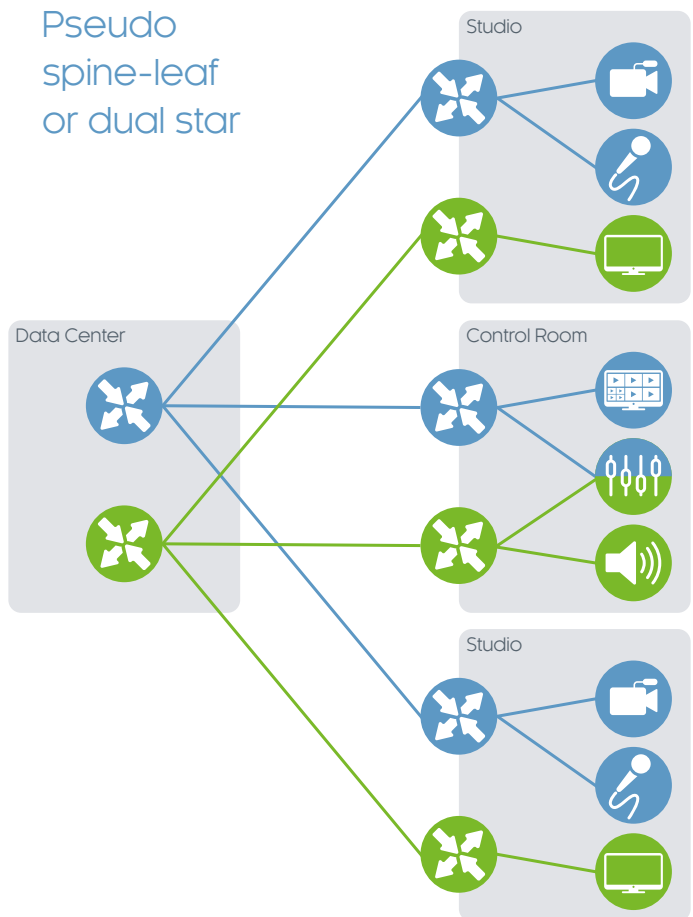
This third architecture model is what some might call a spine-leaf, but in reality, it is a ‘dual star’ architecture model that fails to deliver all the benefits that come with a true spine-leaf structure. This architecture still involves the use of two spine routers, but each leaf in the network is only connected to one of the spines.

For example, the diagram to the right indicates that it is not possible to establish a connection from the camera in the top studio to the monitor in the bottom studio - there is simply no path available.

This solution is not flexible when it comes to load distribution and optimization of total network capacity. As the network evolves, this pseudo spine-leaf approach puts special requirements on end devices that need redundant connections. It also suffers from the same redundancy issues that are common with the centralized star approach: if one of the central routers goes down, a large portion of the overall connections is lost.

There is, however, a reason why some push the dual-star approach, despite its clear limitations, and that is control.

The proponents of this architecture usually also favor automatic protocol-based routing (see below) rather than Software Defined Networking (SDN), and from that perspective the dual-star has one big advantage over true spine-leaf: it does not have any network loops – something that automated control finds difficult to cope with.



But this issue with automated control does not make the dual-star the better option overall.

In short, broadcasters must be conscious of the significant differences between a true spine-leaf model and a pseudo spine-leaf. While a pseudo spine-leaf might be initially more appealing due to the simplicity in setting it up, and may be appropriate for some remote production connectivity, only a true spine-leaf architecture will enable broadcasters to get the most out of their IP infrastructure investment in the facilities.

“The dual-star solution is not flexible when it comes to load distribution and optimization of total network capacity”



Control - connecting things

In tandem with the network architecture, broadcasters need to make another important choice: how to orchestrate and control the IP media network, i.e. to connect sources (e.g. cameras) and destinations (e.g. switchers or monitors), and switch between sources and destinations (e.g. to switch between cameras) - at a fast pace. This control decision is as important as the overall architecture of the network, because it can enable or restrict the agility and flexibility of an IP network.

Control in this context is not the same as management. This is not about which control panel or management system to use, but refers to how the commands from such systems get carried out in the network.

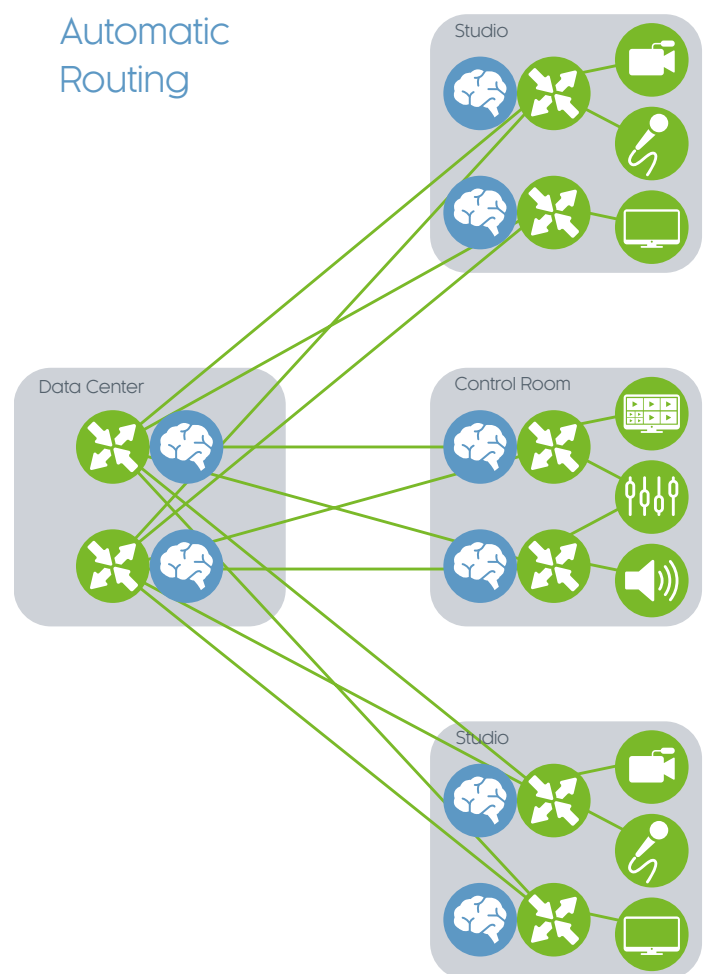
In simple terms, there are two main ways of controlling how signals and connections are routed in an IP network: automatic routing and SDN.

Automatic routing (aka in-band or protocol-based routing)

The protocols that typical IP switches run (e.g. IGMP/PIM) enable the network elements to make decisions about routing based on the IP traffic. This means that the decision of how to transport individual media flows across the network can be left to the network itself, rather than to the operator. This is what is more commonly known as automatic routing.

To execute the request for a media destination (e.g. a studio monitor) to connect to a source (e.g. a camera), two things need to happen. Firstly, the source needs to put its output flow (e.g. video signal) onto the network as a multicast. And secondly, the destination needs to be told where it can find that source, i.e. what multicast to subscribe to.

While automatic routing, and in particular the IGMP and PIM protocols, are used widely in IP networks across the world, they have some disadvantages when it comes to professional real-time media production networks.



IGMP AND PIM

The IP switch fabric typically uses IGMP (Internet Group Management Protocol) and PIM (Protocol Independent Multicast) to control routing. IGMP is used by destination devices to request a connection to a multicast stream from a particular source. PIM is used to enable forwarding of multicast streams within the network.



Performance is a major issue: depending on implementation and requirements, automatic routing may not be fast enough to deliver the significant number of simultaneous switching events required in live production, especially at key times in a large distributed network.


Automatic routing can also get into trouble with networks where loops exist (such as true spine-leaf networks). This can be fixed by applying sophisticated automatic routing configurations in each network element, but it will lead to significantly higher operational complexity. Therefore, loops are often avoided by blocking links around the network instead. However, closing links means that some network capacity is never used, which is wasteful and expensive.

In addition, the switch fabric typically does not handle bandwidth management. This means that, unless great care is taken in designing and controlling the network (for example significantly over-engineering the bandwidth), there is a risk of oversubscribing it. This could cause instability and drop-outs of signals (especially given the huge volumes of data involved in the transport of uncompressed video in real-time).

This problem is not helped by the fact that automatically routed networks typically only make routing decisions based on the initially configured network topology. They do not take overall traffic load into account. They also do not have the concept of planning and reserving capacity - for example for a live production known to take place at a specific date and time. In other words, automatically routed networks deal with bandwidth demand as it happens, which is neither efficient nor effective when handling high volume traffic like uncompressed video.

On top of this, there are also concerns around protection and security, as streams to destinations are not explicitly controlled, which means potentially any destination could subscribe to a multicast - including 'rogue' equipment. This is typically fixed by configuring complex access control lists on each IP router, resulting in significant operational complexity.

In short, despite being widely deployed in other industries, automatic routing does have limitations that can affect media networks adversely.



If automated routing was used for car navigation

One way to picture automated routing is to imagine a totally hypothetical car navigation technology, that would not rely on a centralized, overarching view of the road network and the traffic (as is the case for satnavs). Instead, this navigation technology would use special routing equipment located at every road junction.

This routing equipment would be in constant communication with the routing equipment located at the previous and next junctions, but not beyond. When the car being navigated would approach a junction, the routing equipment located there would suggest which road to take, based on the less busy next junction. The routing equipment would not take account of any congestion further down the road in that direction, as it would not be aware of it.

This would probably be appropriate in the countryside where traffic is limited or for drivers who are in no hurry to get to their destination, but much less so for critical deliveries!

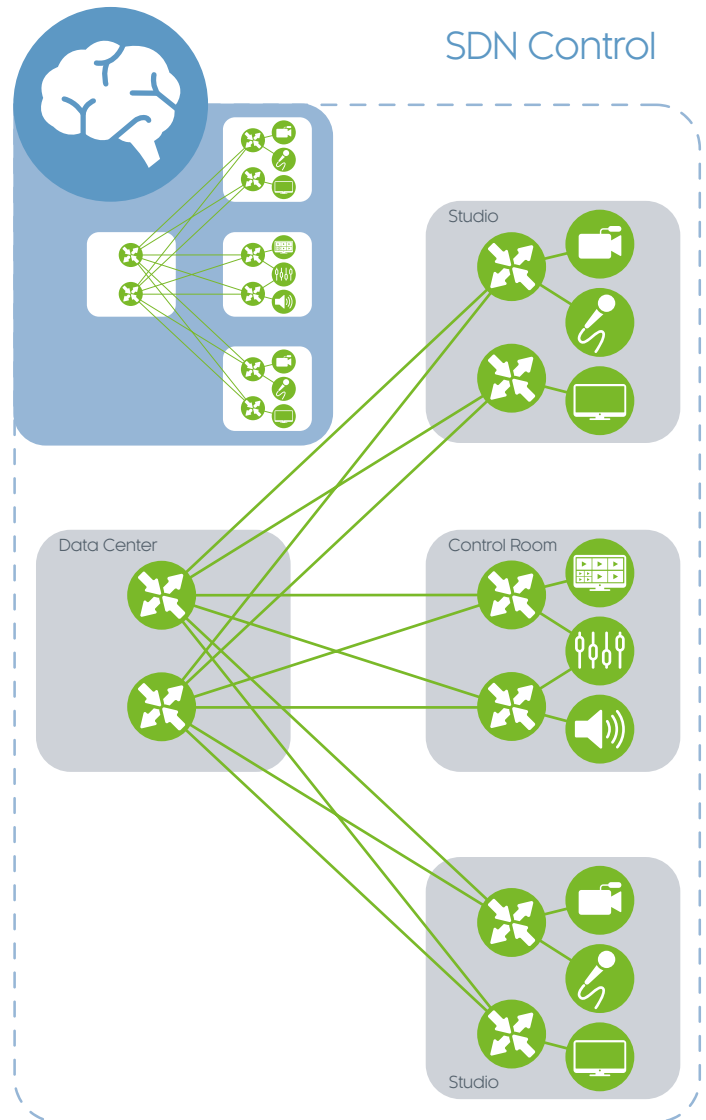


Software Defined Network (SDN) routing

The concept of Software Defined Networking (SDN) is about taking the routing control away from the individual network elements, and putting it in the hands of a centralized control layer - for example Nevision's VideoPath orchestration and SDN control system.

SDN has been around in the IT industry for several years, and its adoption is increasing – particularly in data-centers. SDN as a concept is also both vendor neutral and supported by industry standards, like Netconf/Yang, OpenConfig or Openflow.

With SDN, the management and orchestration software holds a complete view of the available equipment, the network infrastructure and the services (media flows) - not just those currently in place, but also those that are planned. Thanks to this, it can make intelligent decisions on routing and controlling flows far more efficiently than is often possible with automatic routing. It can also provide the explicit routing capability that broadcasters expect and need.



“Software Defined Networking (SDN) is about taking the routing control away from the individual network elements, and putting it in the hands of a centralized control layer”



When a control element (e.g. a control panel, switcher or mixer) needs to connect a source and destination, the orchestration and SDN control software layer determines the best routing options for the signal based on the situation (current and planned network load) and the requirements (e.g. required bandwidth and redundant paths), and pushes routing

information into the network by updating the routing tables in each network element along the identified path. This can be done using the same IP switches and routers as in an automatically routed solution - it does not require special network components.

There are many advantages to opting for an SDN routing approach.

First off, it guarantees a much higher level of performance when compared to automatic routing. As the management and orchestration software has all the required information about how sources and destinations are interconnected and the processing power to make the switching decisions fast, it can meet the clean switching speed requirements of studios. The software is also in control of every media flow, which means it is much more aware of, and much better at dealing with, existing and even planned (scheduled productions) bandwidth requirements.

It is even beneficial from a protection and security perspective. The orchestration and control software knows all about the network topology and how to control the routing, which means it can easily create path diversity to protect failures, and can also fully control which destination is allowed to receive which multicast, thereby reducing the security risk substantially. The default behavior of an automatically routed network is to allow all traffic, unless it is specifically instructed to block it. In contrast, an SDN-controlled network works in the exact opposite way: it blocks all traffic, unless specifically instructed to allow it.

“SDN can, with the right orchestration and control software, easily handle any network architecture”



And of course, unlike automated routing, SDN can, with the right orchestration and control software, easily handle any network architecture, whether star, dual-star (pseudo spine-leaf) or true spine-leaf - without compromise.

Despite the clear advantages of a SDN routing approach over automatic routing, there continues to be passionate debate about which option is the best. Many leading names in this industry will actively champion automatic routing as the approach of choice, but there is a good reason behind this: many of their products have automatic routing built into them, and so there is a vested interest to promote a technology they have spent lots of money on, and which they believe provides a unique selling point.

Nevion, on the other hand, has no such vested interest. The VideoPath software can handle both automated routing (e.g. IGMP/PIM) and SDN, but the benefits of SDN make it the control of choice for the creation of truly flexible, scalable and high-performance IP media networks.

Conclusion

There is no doubting the influence that IP technology is having on the broadcast industry at this very moment, and the benefits it can provide are plentiful. It's an agile, flexible and scalable infrastructure solution that allows broadcasters to future-proof themselves while - most importantly - maintaining and even increasing reliability and maximum uptime.


“With a true spine-leaf architecture and an intelligent SDN routing system in place, broadcasters will be able to flourish in this exciting IP-based era.”



To achieve this, however, broadcasters have to remember that they must learn to walk before they can run. A successful IP infrastructure is not built around the products and components that comprise individual elements, but the way in which the infrastructure is designed and built from the ground up. A large factor of success is also determined by the way in which individual elements within the network are controlled.

As explained above, there are numerous options to choose from when it comes to both the architecture and control of an IP network. Each of these options has its own unique list of advantages and disadvantages, and depending upon the industry it's being used for, each option is credible. However, when the time comes for broadcasters to invest in an IP-based infrastructure, they should ideally be architecting the network using a true spine-leaf model, and controlling the elements within it using a SDN routing approach. This combination is ideally suited to an environment where agility, scalability and resiliency take priority over everything else.

This is incredibly important for broadcasters to keep in mind when building their own IP network - it must be designed with broadcast as the main priority. While it is perfectly reasonable for companies to reach out to IP switch vendors for their advice on how to go about building the infrastructure, they must remember that broadcast requirements are wildly different to those of other industries, and just because a certain method worked well for them doesn't necessarily guarantee



success within broadcast. Yes, sophisticated IP infrastructures can be scaled up or down to meet needs, but this isn't very useful if the infrastructure struggles to meet the overall business goals in the first place.

With a true spine-leaf architecture and an intelligent SDN routing system in place, broadcasters will be able to flourish in this exciting IP-based era. It is the most effective way of squeezing as much potential out of IP technology as possible, which in turn helps to deliver optimal return on investment and much higher chances of operational success.

This whitepaper is an extract of Nevion's "From Baseband to IP to Virtualization – Architecting the media production infrastructure for the future", which is available in print at most tradeshowes attended by Nevion and its partners.



Nevion near you!

Nevion has a presence in all the major regions, and an extensive network of partners to reach customers anywhere in the world.

Visit our website for your nearest sales contact

nevision.com

©Nevion 2018

No part of this document can be copied, reproduced, distributed or otherwise transmitted without the express written permission of Nevision.

Disclaimer:

The information and data contained in this brochure has been prepared by Nevision solely as a guide, and does not constitute guidance or advice, and should not be treated as such. By using this brochure, you accept this disclaimer in full. Nevision will not be liable to you in respect of any business losses, and without limitation loss of or damage to profits, income, revenue, use, production, business, contracts, commercial opportunities or goodwill.