WHITE PAPER

Redundancy in live IP media infrastructures

Is the concept of ST2022-7 flawed for infrastructure applications?

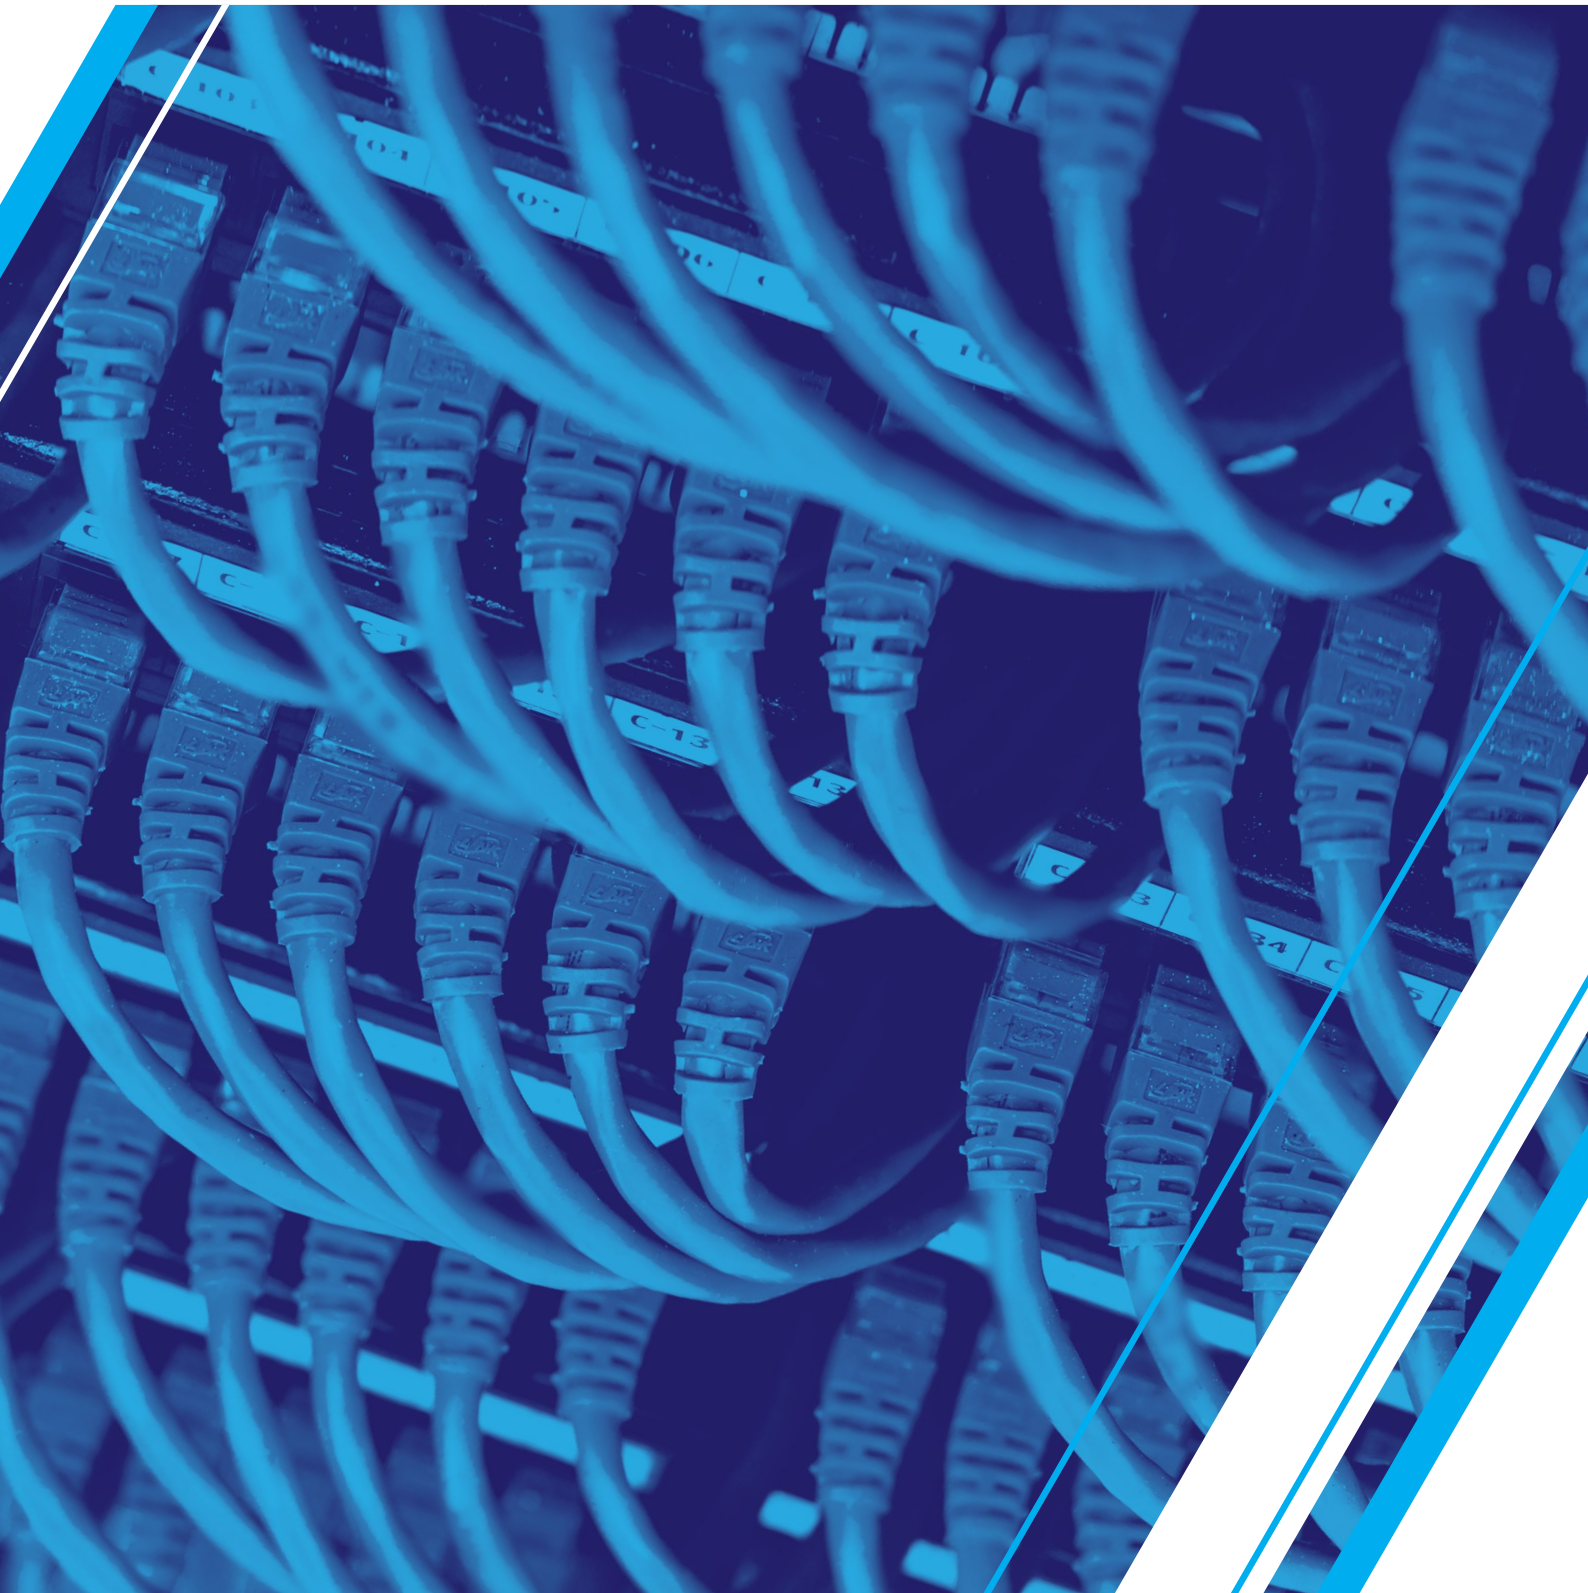

Introduction

This white paper covers some of the challenges faced by the broadcast industry about how to address redundancy and remove as many single points of failure as possible in an IP environment.

Focusing on live baseband media applications rather than generic IT infrastructures or compressed media distributing systems, this paper addresses the different challenges, and highlights the different methods to ensure redundancy in the IP domain.

Redundant “red” (main) and “blue” (backup) networks can be found in almost every live production facility. Although a redundant network is designed prevent system downtime caused by issues with switches, it does not prevent black screens or audio silences when faced with issues linked to media processing equipment. In this case, where is the redundant audio and video processing? What will happen with ST2022 and ST2110 streams when things go south? Is packet loss acceptable? Will a failover switch be visible? These are the kinds of questions we shall try to answer in this paper.

Contents

Moving from an SDI to an IP infrastructure has a massive impact on the backbone's resilience

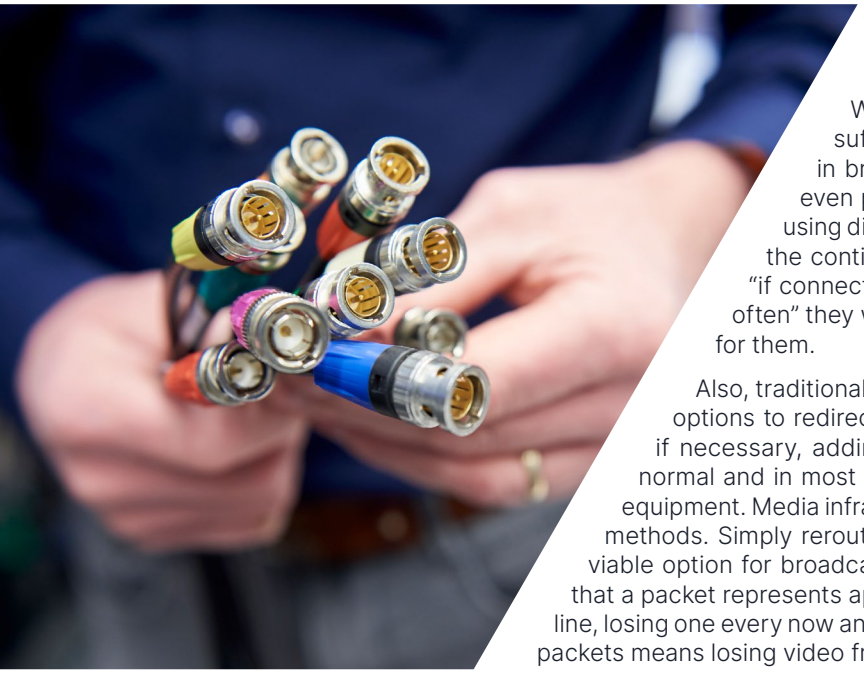

IT challenges

Moving to IP involves a huge increase in bandwidth usage. While a 10 GbE IT backbone is considered modern and sufficient to support an entire office, a 10Gb/s connection in broadcast terms is quite low bandwidth, since it cannot even provide a single baseband UHD stream. This requires using different switches, all needing to work harder to handle the continuously high bandwidth load. The question is not "if connection problems will happen", but "when" and "how often" they will occur, so broadcasters need to be prepared for them.

Also, traditional IT backbones rely on a mesh network, with options to redirect packages through different routes and, if necessary, adding a delay. Resending packets is quite normal and in most cases does not cause any issues in IT equipment. Media infrastructures however cannot use these methods. Simply rerouting and resending packets is not a viable option for broadcast operations. When we consider that a packet represents approximately a quarter of a video line, losing one every now and then is not acceptable. Losing packets means losing video frames.

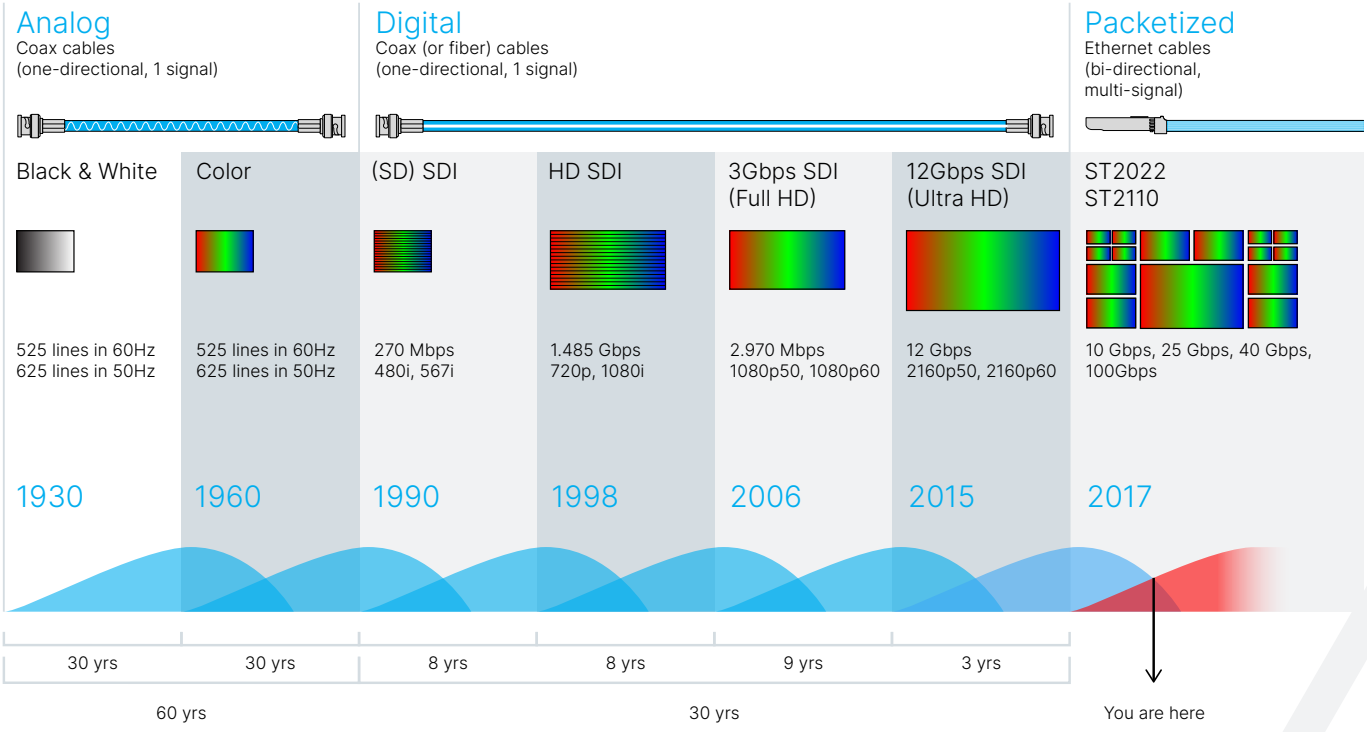

Characteristics of IP media infrastructures

Besides packet loss or resending infeasibility, another major necessity in the broadcast industry is sustaining full capacity of the main Ethernet switches, which are continuously loaded with very high bitrates. Internal switch buffers are quite small for baseband requirements, and traffic shaping has a limited effect on the switch output, as packets remain small. If a few, identical, traffic-shaped streams enter a switch and are combined on an output port, the shaping might not even be visible.

Another factor is that Ethernet connections contain many streams and channels nowadays. 32 video channels on 100Gb/s or 5000 audio channels on 10Gb/s is not exceptional. This means that media processing equipment suddenly has to process multiple channels at the same time. Even the best broadcast equipment can (and most probably will) fail at some point, causing a single point of failure for numerous channels at once.

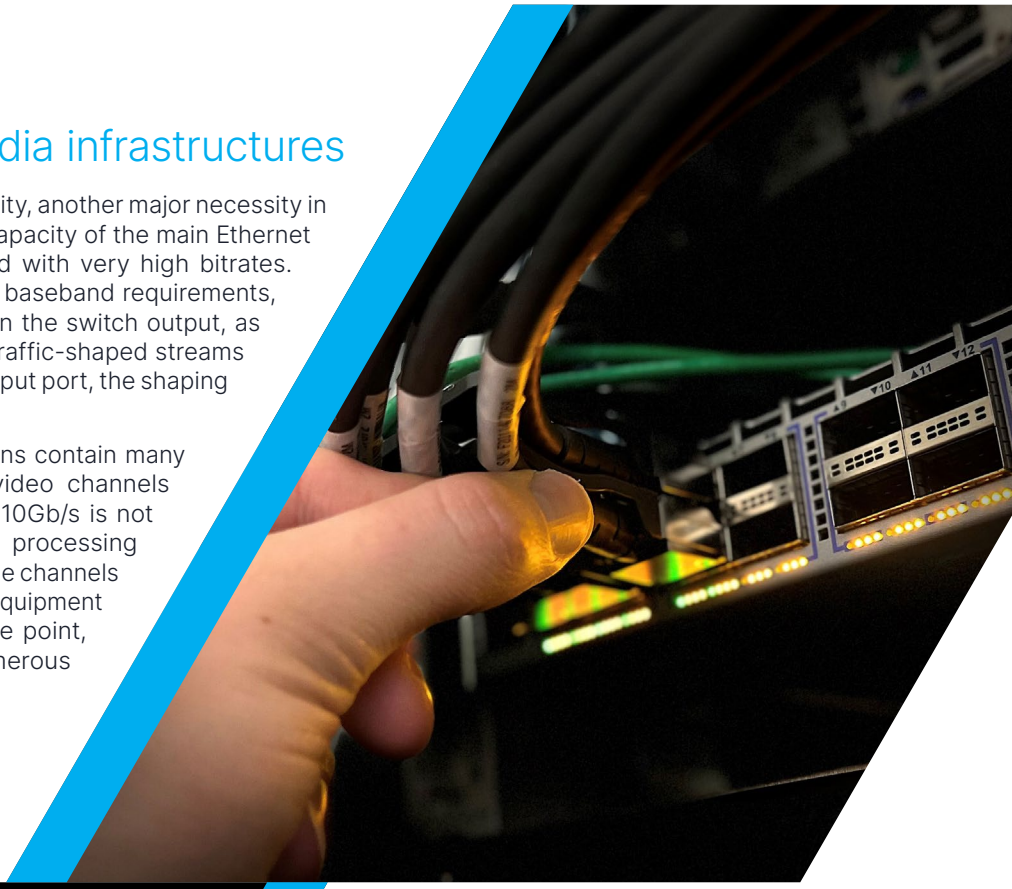

A typical broadcast infrastructure has a range of different shapes and flavors of incoming media streams. In the IP domain, for the most part, these streams currently comply with ST2022 and ST2110-based standards, with all the different dashes and suffixes that this implies. For some broadcasters, ST2110 is seen as the most modern infrastructure standard. While this might be largely true, it is not always the best way forward. In the SDI domain, there were good reasons to keep the audio embedded in the video blanking, and to lock the video and audio together throughout the production workflow, as it prevents lip-sync errors and mixing incorrect audio channels.

The ST2110 standard dispenses with this concept, by splitting video and audio into separate streams again. If only minimal manipulation of audio is required, such as for a playout facility or master control room, locking the streams together using ST2022, makes a lot of sense. The redundancy requirements nevertheless remain the same for both ST2022 and ST2110. With ST2110 however, there are a lot more streams, each of which requires a back-up stream. As result, thousands of IP streams are created in a single network, which all have to be managed correctly and redundantly.

ST2022-7 redundancy and packet loss

ST2022-7 (or simply -7) is a mechanism designed for long-haul applications and situations where major packet loss and corruption is expected. It can be used in an infrastructure to select between main and backup signals. The concern is that -7 can give a false impression of a connection's resilience and might only work if both streams are present.

There should be a possibility, for instance, to switch off one source or one switch for maintenance purposes. A -7 structure is perfectly capable of handling such an action, but the catch could be that the infrastructure is only resilient for -7 and not single -6.

In other words, even though a -7 may look like a resilient connection on paper, the actual requirement is for two perfect -6 streams to be able to rely on one half not being present. A system based on -7 resilience – cherry picking between two streams to make a perfect single stream – will not allow one half to be totally unavailable. A resilient solution for a real long-haul application with natural packet corruption therefore requires a dual -7 setup, because packet loss, again, is not an option.

The SDI backup switch

With a typical active backup switch in the SDI domain, when an error is detected on the main input, the active backup switch changes to the backup input. To help prevent video loss in a situation where there are also errors on the backup switch, a passive relay bypass ensures a continuous feed, by circumventing the processor entirely (see diagram below).

The SDI backup switch is a straightforward concept to achieve a passive bypass with a relay when there is a power failure, or for instance if the wrong processing module is accidentally removed (as the relay switch sits on the connector panel, which is powered by the active processing module).

Of course, since almost everything is different in the IP domain, this SDI backup switching technology cannot be used in an IP environment. A completely new solution is required.

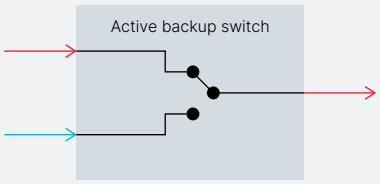

SDI, normal situation

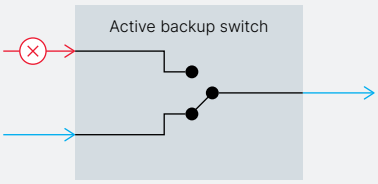

SDI, backup situation

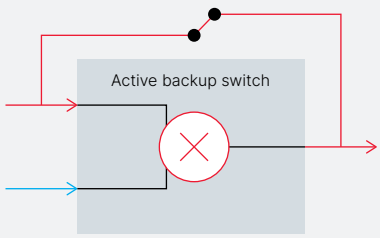

SDI, relay bypass situation

The -7 backup switch

The industry's most common solution for redundancy in the IP domain is the -7 backup switch method. There is a main (red) IP switch and a backup (blue) IP switch. The switches are exact copies of each other, feeding all streams to an active backup selector. It is the task of the backup selector to create a blue and a red layer for all devices and destinations behind it. In the "All OK" situation, both the blue and the red layers are sourced with the signals from the red IP Switch.

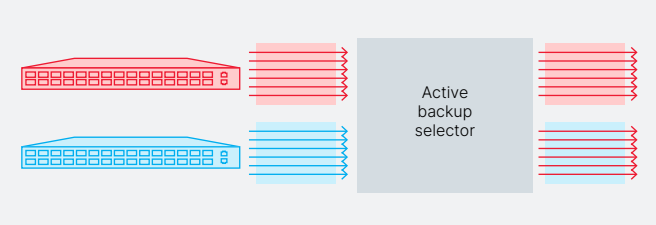

To illustrate how the -7 backup selector functions, let's take the example of the red switch failing entirely, because the power cables were inavertendly disconnected. In this situation, the backup selector immediately switches to the blue sources for both the red and the blue layer.

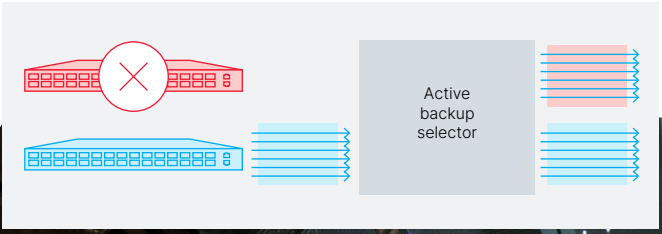

Of course, an entire switch failure will not happen often. In most cases only a couple of sources will fail, for reasons varying from bandwidth overloads to connection losses or subsequent devices failing. In these cases, the backup selector "fills in the blanks" with blue sources in both the blue and the red layers of the facility.

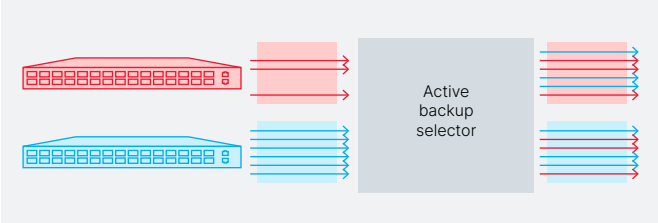

Protection against failure of the source is possible in -7 if a switch fails, or if one of the up-stream sources fails, or even when carrying out maintenance on upstream equipment. But what if the backup IP selector fails? There would be no active sources for any of the destinations. In other words, it would be the ultimate meltdown situation.

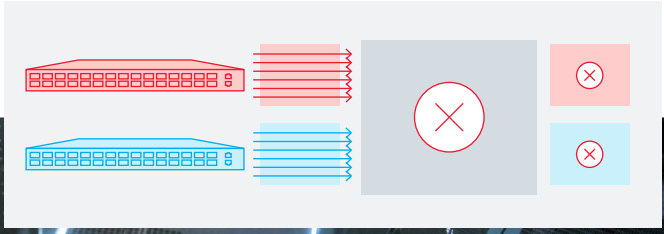

Legend

Red IP switch

Blue IP switch

Red signals on red layer

Blue signals on blue layer

Blue signals on red layer

Red signals on blue layer

The need for redundant backup selectors

Protecting against such single points of failure requires a dual backup selector: one to provide a red layer, and one to provide a blue layer. This will allow for redundant red and blue outputs and, in the worst-case scenario, when a backup selector fails, provide at least a red or a blue output. However, this would imply doubling the number of ports on the core switch fabric. Nevertheless, it provides far greater protection in the case of a single point of failure for a large number of streams as described above.

Here is what this solution would look like in the diagram below. The top, active backup selector provides the red layer, and the bottom selector feeds the blue layer. In the illustrated "All OK" situation, both backup selectors are using the red sources.

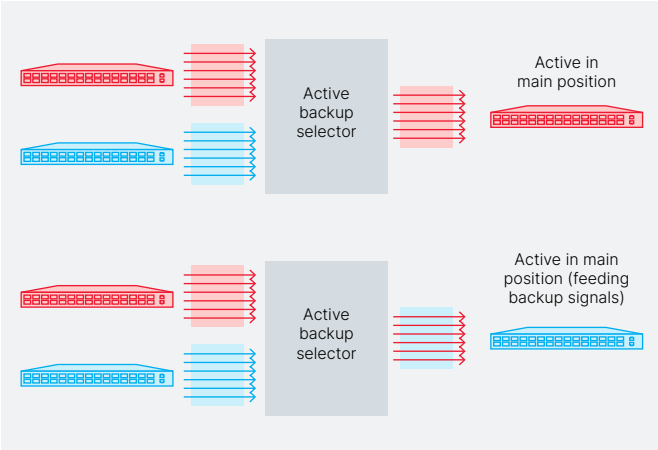

To compare with the usual -7 method described earlier, if a core switch fails, both backup selectors will switch to blue sources for the layers they are responsible for.

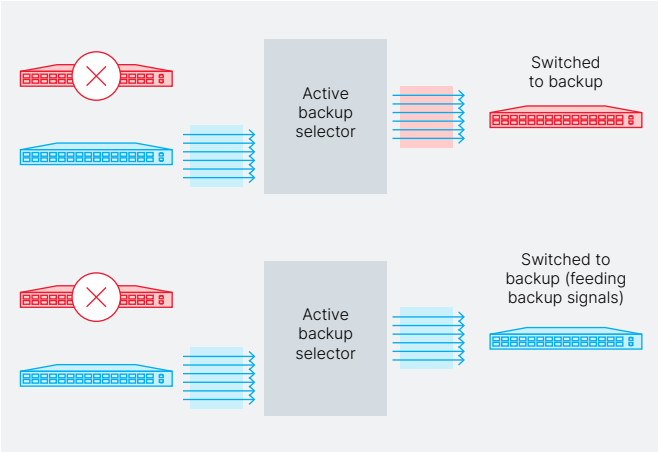

If one of the backup selectors failed, all the streams would still be available on the other layer of your facility. This is illustrated below, where the red layer of the top backup selector has failed, but the blue layer is still active, because the blue layer backup selector is a physically separate device.

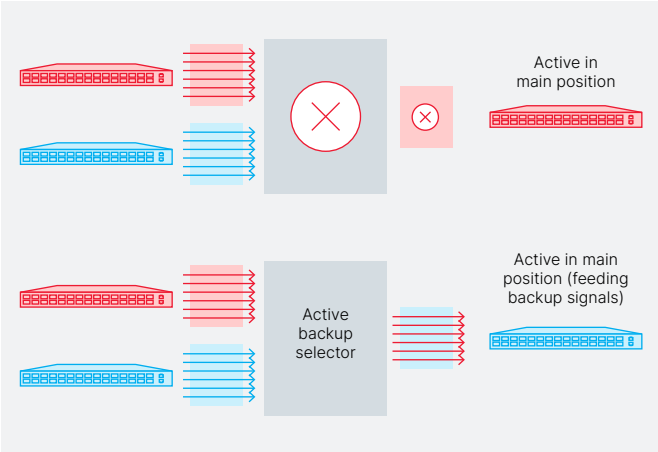

Adding intelligence and security to your backup selectors

The situations described above assume that an input loss is detected upstream from the backup selector switches. Of course, in real situations, the backup switch needs to be a little more intelligent, as losing an input is not the only reason for switching away from a stream. In the SDI domain there are usually integrity checkers, continuously probing for video freeze, video black, audio silences, loss of subtitle data and other incorrect ancillary data, etc. Whenever an issue like this is detected, it is possible to switch to a backup SDI input. The same should be considered for the IP domain. This is especially true when the incoming IP streams are from an external, less trustworthy source, to be able to ensure that feeds to IP devices are absolutely clean and reliable.

For security reasons, some level of isolation is required, to remove any potential hostile "code" from the incoming IP stream. This can easily be achieved by decapsulating the AV streams, bringing them back to baseband media, and then encapsulating them again in fresh new ST2110 or ST2022 packets. In this way, no hacker code would be passed through.

Intelligent backup selectors are also called for, to communicate and let the other selector know the state of its signals; not only monitoring for input loss, but also for video freeze and black, audio silence, ancillary data presence, etc. On top of this, if they could provide each other with their signals, i.e. the main selector sharing its red streams with the blue backup selector, and the blue selector sharing its blue signals with the red selector, this would avoid having to duplicate the number of cables to the core switches.

This ultimate IP redundancy scenario would look as follows:

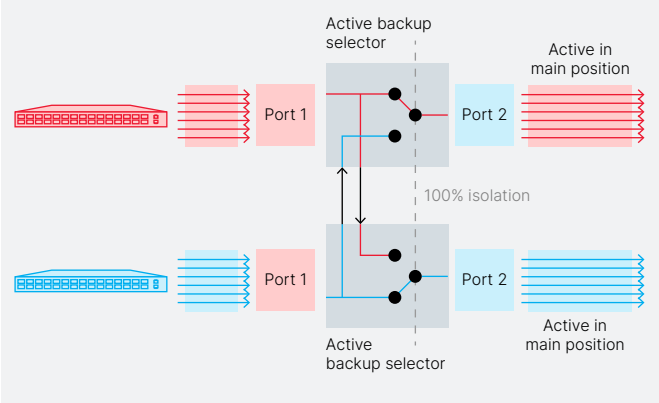

If the red backup selector detects an issue in one or all of the red streams, it checks with the blue backup selector if its streams are all OK. If they are, the red selector, which also receives all the streams from the blue selector, switches the failing red sources to the blue sources from the blue backup selector. At the same time, it also takes care of decapsulating and encapsulating, and thus isolating the IP streams.

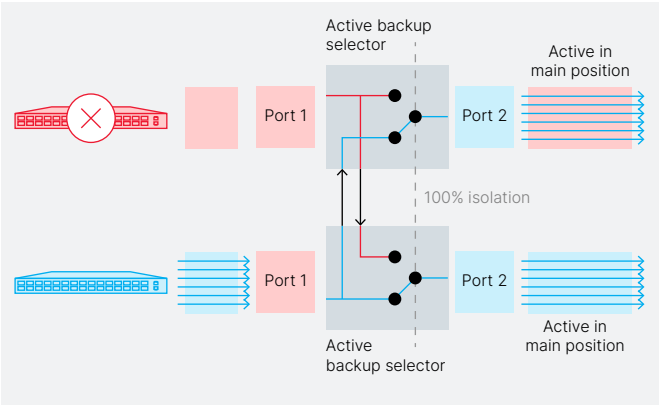

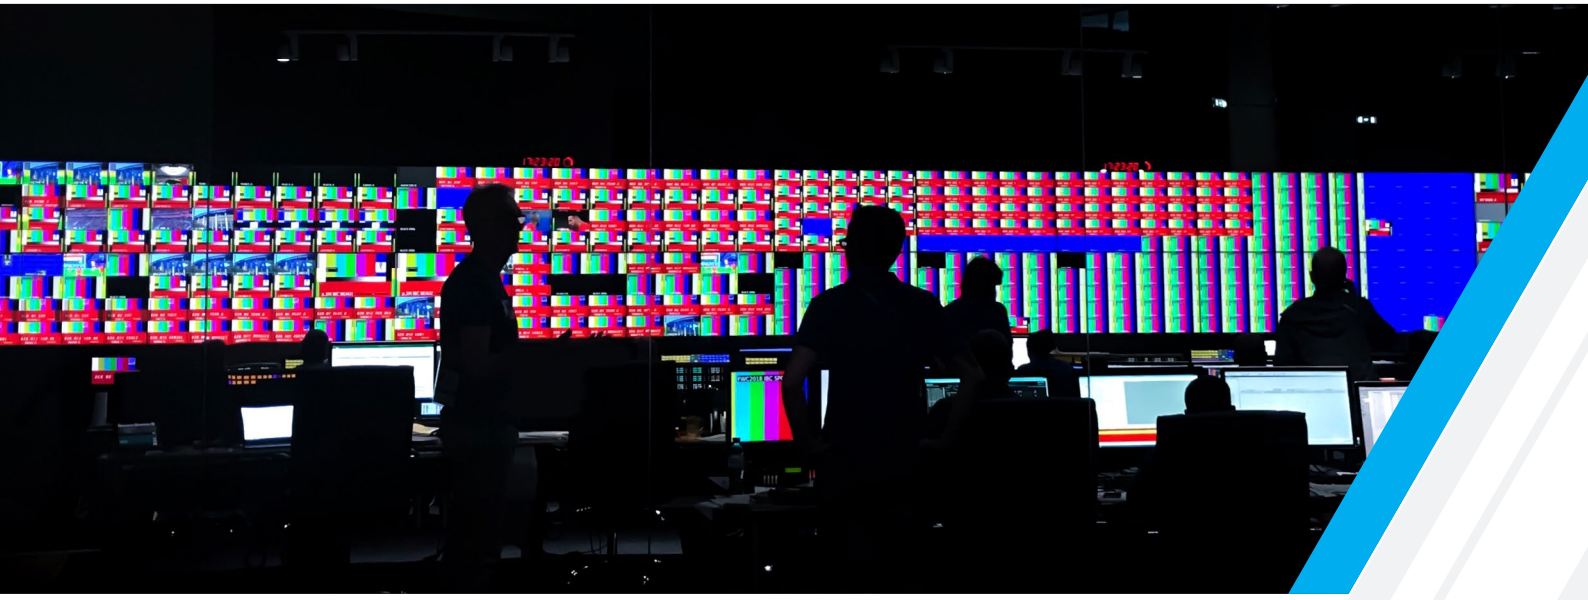

Conclusion

In an IP world, the problems and solutions work differently compared with the more familiar SDI world. IP has more, bi-directional channels over fewer cables, which is the whole reason for moving to IP. As multiple signals are handled through a single connection, broadcast media processing equipment can become a single point of failure for multiple channels at once. Having a red and a blue network with one media processing device in the middle can be risky, but there are ways to overcome this.

Intelligent, redundant, isolated backup selectors avoid having to worry about “what if?” scenarios. They provide full peace of mind and protection from unreliable or even hostile sources.

EVS has the perfect solution for this: the PROTECT line in the Neuron real-time IP media processor. The PROTECT cards use 100G external connections to the red and the blue IP switches, while sharing red and blue signals via an internal TDM connection. They include 100% isolation and can even continuously probe for audio, video, and ancillary data errors.

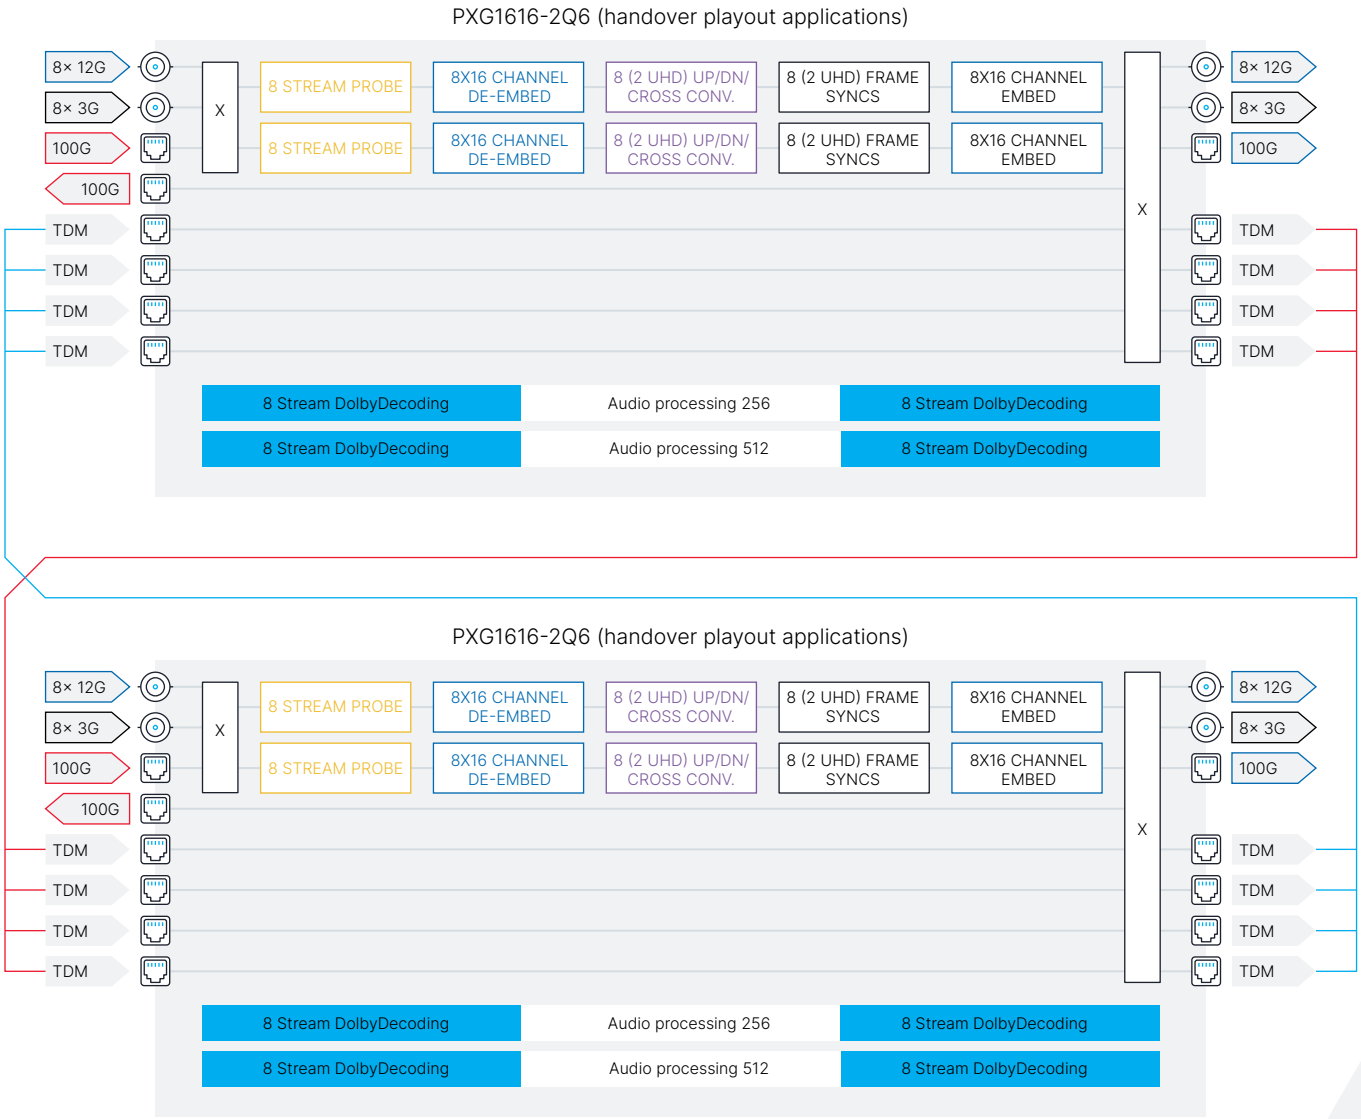

Designed around IP technology, Neuron sets new benchmarks in performance and flexibility

→ learn more → contact us