CHANGING THE WAY OPERATORS DELIVER TV

Appear TV X Platform

Dual Backplane and Redundancy Features

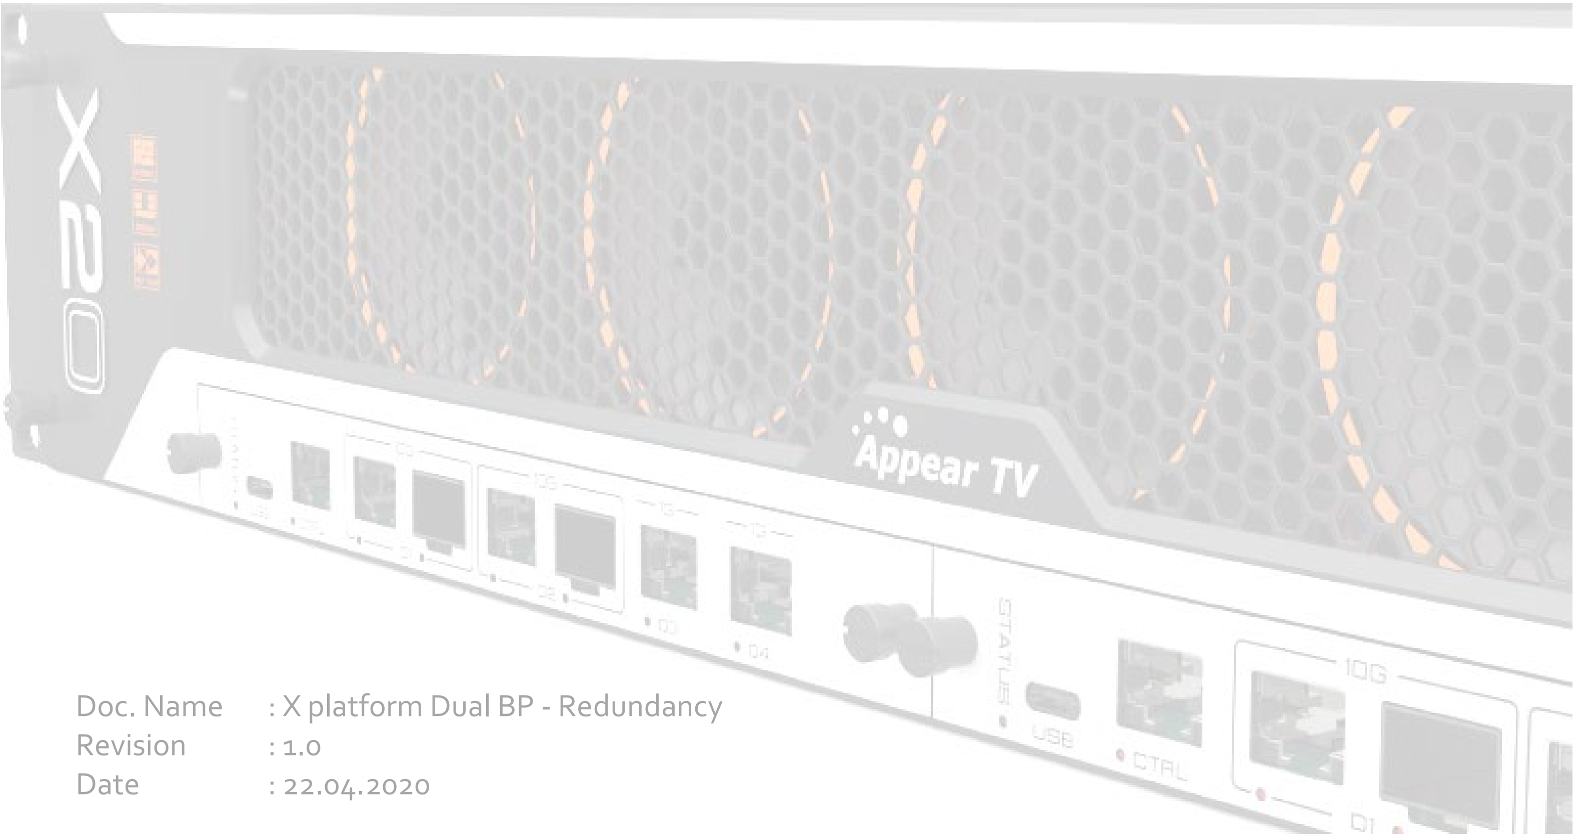

Doc. Name : X platform Dual BP - Redundancy
Revision : 1.0
Date : 22.04.2020

Table of Contents

1. Introduction

The following document provides an overview of some of the key features available in Appear TV's X platform. The X Platform has been developed to be a highly flexible, customisable Contribution and Distribution solution used by Network and Broadcast operators for compressed and uncompressed video transport. Redundancy has always been one of the key points that had to be taken into account when development started on this product.

Redundancy is one of the features frequently requested by operators that are running distribution or contribution network. There are many ways that redundancy can be added to a solution from something as simple as dual PSUs going through to 1+1 chassis redundancy, 1+1 redundancy within the same chassis, Input redundancy, seamless redundancy, OSPF network redundancy…. the list goes on.

This document covers several of these scenarios and should provide a good overview of what is possible to deploy when using the X platform.

The following solutions are included:

- Dual Backplane
- Input Redundancy
- Hot Standby
- Identical Sources
- Further Redundancy and Features
    - IP Input policing
    - OSPF Output Redundancy
    - Cloned IP Output
    - Seamless IP Input
    - Muto on Error
    - Monitor-In-Out

2. Dual Backplane

When operating Appear TV's X chassis in dual MMI / dual backplane mode to increase fault tolerance or even provide full 1+1 redundancy within the confines of a single chassis, it is important to understand how the system works.

The two key items to understand are as follows:

- Operation in dual backplane mode and what this means.

- The INPUT REDUNDANCY concept since this will determine whether the output services originate from the main or backup inputs. INPUT REDUNDANCY is a key tool that is often required to make dual MMI operation work.

The X platform has been designed to allow customers to build fully integrated solutions with great ease. An important part of this is the internal backplane, a fully integrated low latency network that connects all the modules together. The diagram below provides a general overview of how the switch module connects to the installed modules over the two available backplanes:

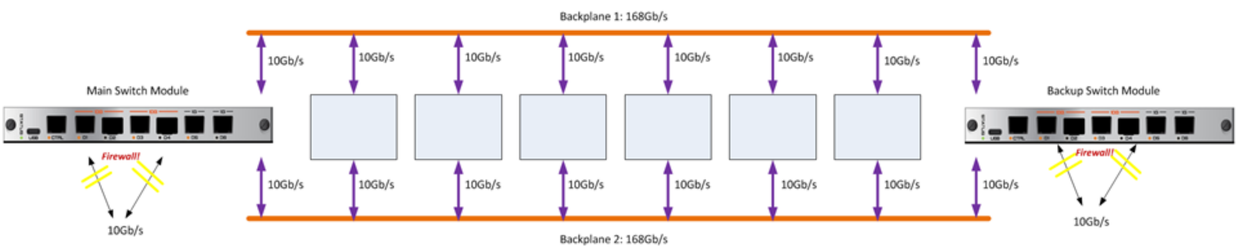

This network is provided in star form, with the switch/management module forming the central node of the network. All connections via the backplane to the modules operate at duplex speeds of up to 10Gb/s.

Many customers equip their chassis with one switch/management module which implements a single star network linking the modules together; one internal network is usually enough if a chassis is backed up using a typical 1+1 arrangement. If the chassis is running with one internal network and the worst type of fault possible were to occur (for example the switch/management module completely fails) then it is possible the entire chassis could stop functioning. For 1+1 chassis redundancy where a separate main and backup chassis is deployed, a switchover to the backup unit would recover the failure but an alternative approach is to fit a main and backup management module within the same chassis.

It is possible to install dual switch/management modules in either the X10 or X20 chassis. Providing dual switch modules may be sufficient to meet a customer's needs, for example by reducing the 'granularity of failure' to the point where the worst possible fault that may occur can be limited to one module and therefore a limited subset of services. However, it could be a step towards implementing a fully 1+1 redundant solution within a single chassis to ensure no single point of failure can be service affecting. In every case where adding redundancy within a single chassis is the goal, adding 1+1 MMI module redundancy is always the first and most important step. It protects the internal network ('chassis backplane') against failure by implementing these in duplicate, and also backs up the chassis management function seamlessly.

With dual switch modules installed, two control access points are provided to manage the chassis. X platform has a largely distributed control architecture, so the net result of a management failure is zero when two switch/management modules are installed; all chassis functions will continue as normal.

The same generally applies to the backplane function since with dual management modules there are two independent networks functioning within the chassis.

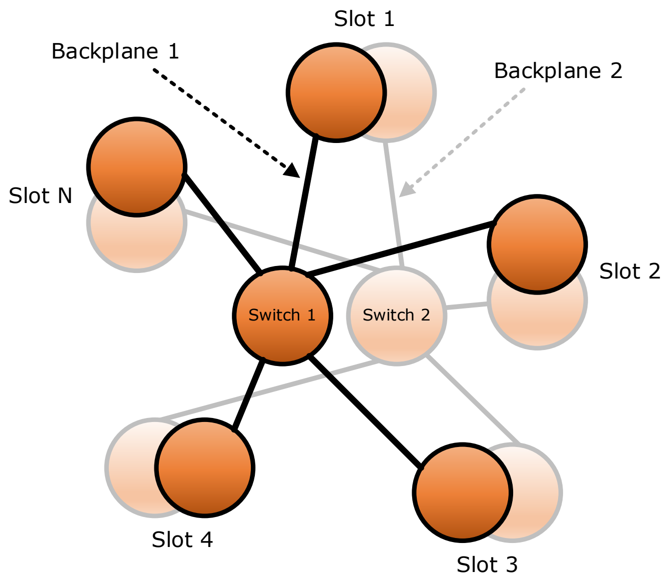

Management module 1 forms the core network of BP1 and management module 2 forms the core network of BP2 as shown in the diagram above.

Note: To implement dual backplanes, two switch/management modules need to be installed into the chassis.

A simple example of this is where a chassis is equipped with dual management modules and so has dual backplanes in operation. The traffic flow between two rear mounted modules (for example an HEVC encoder module sending packets to a DIP module (IP output)) are shown below.

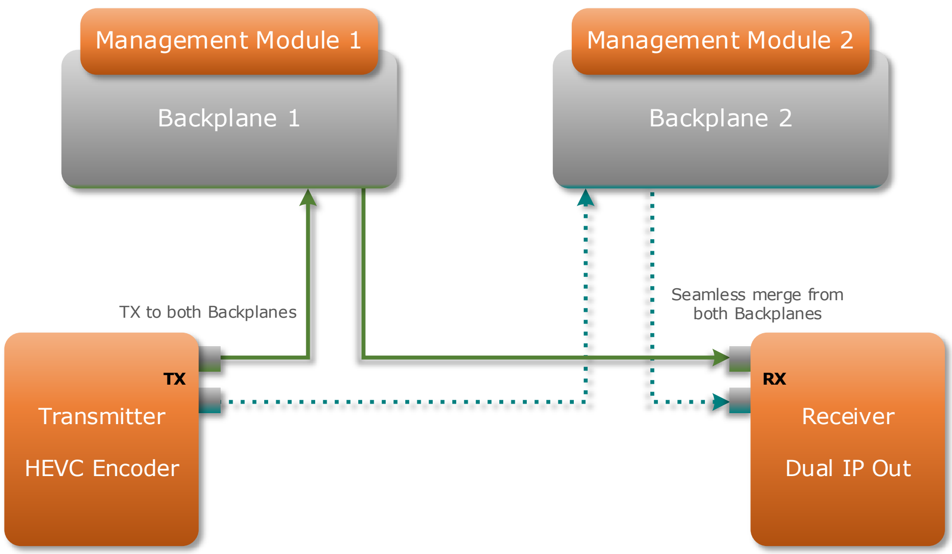

This is a simple case and highlights the basic transmit and receive principles:

- The transmitter (in this case the HEVC encoder) always duplicates and sends its traffic over both backplanes.
- The receiver (in this case the DIP module) always performs a seamless merge using both backplanes.

Operationally, this provides complete and seamless backplane diversity. For example, if either of the management modules will cause one of the backplanes to fail then this will have no impact on transmission between the encoder and DIP module.

If the HEVC encoder and DIP functions were duplicated, by performing these functions in parallel on separate modules a fully redundant 1+1 solution could be configured within a single chassis. However, since these modules are not duplicated in this case, providing 1+1 management and backplane redundancy serves to improve the granularity of failure of the chassis by ensuring no single management or backplane issue can bring the entire chassis down.

When sending traffic from a management module (acting as IP input) to a DIP (IP output). The situation becomes different due to the multiple functions that this module is performing. For example (looking at the diagram below), if the management module 1 fails, all activities within the highlighted green section could stop which at a system level has the potential to cause a double fault condition where the streaming and backplane transmission may stop.

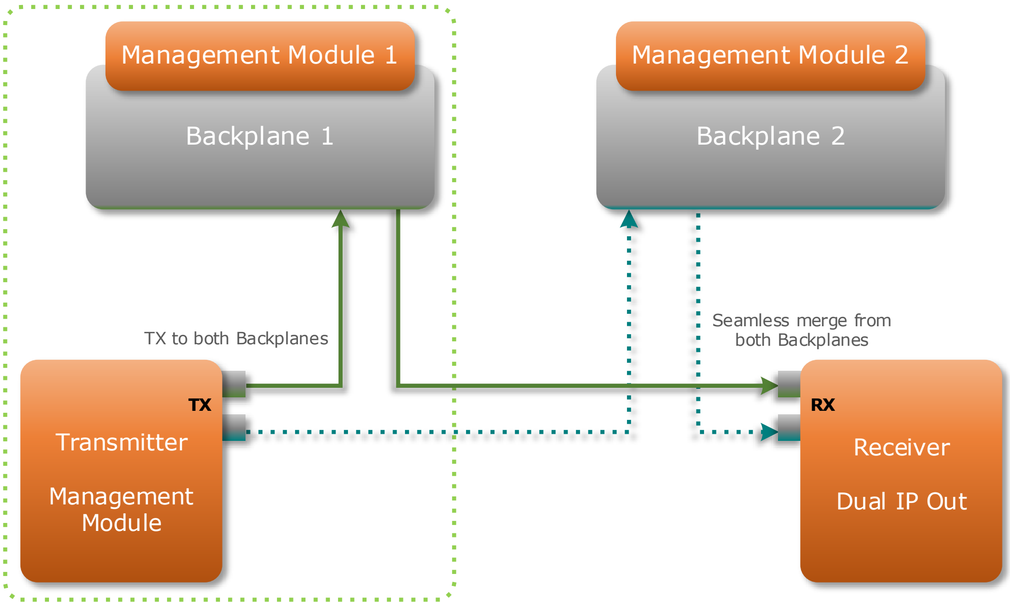

This means when a management module is used for IP streaming, its functions must be backed up 1+1 by the spare management module. This has an important implication since implementing a redundancy solution where the receiver (in this case the DIP module) can select EITHER the service being transmitted by management module 1 OR the service being transmitted by management module 2 requires the use of the INPUT REDUNDANCY feature. This is why INPUT REDUNDANCY can play such a central role with dual MMI architectures.

Furthermore, it is likely that both MMIs will be receiving the same traffic feeds (same multicasts). If these are MPEG_TS the SEAMLESS switching capability of the INPUT REDUNDANCY feature can be used and any loss of MMI function will be seamlessly concealed. For example;

- The loss of MMI management function will seamlessly transition to the backup MMI
- The system can still be managed externally, via the backup MMI control port
- The loss of backplane connectivity provided by the failed MMI will seamlessly transition to the remaining backplane.
- The loss of any feeds the failed MMI was providing will be seamlessly concealed by the INPUT REDUNDANCY feature which will run on the module accepting the feeds. Note: The flow must meet seamless requirements for seamless MPEG_TS switching, otherwise the impact will be approximately 100ms for switching.

Implementing input redundancy enables the output module to check if the service is being received using the primary route (for example, from management module 1) and if not, select the backup route (which will define management module 2 as the source).

3. Input Redundancy

The INPUT REDUNDANCY feature available on the X platform has been designed to be a flexible and powerful mechanism that enables any output module to choose which source to use, depending on whether healthy traffic is being received or not. In this way, the system protects against failure of an input source without caring what the source is. The solution is a generic redundancy scheme, so the user can protect any source with any other source (for example, an IP primary input with a satellite backup from the satellite demodulator module). However, it can equally well be used to protect 1+1 redundant feeds originating from a main and backup encoder module, or a main and backup MMI module, within the same chassis. This is why it is important when building fully resilient single chassis solutions.

The following diagram illustrates the end to end implementation of a chassis running in dual backplane mode (In this case the solution shows IP input to ASI output).

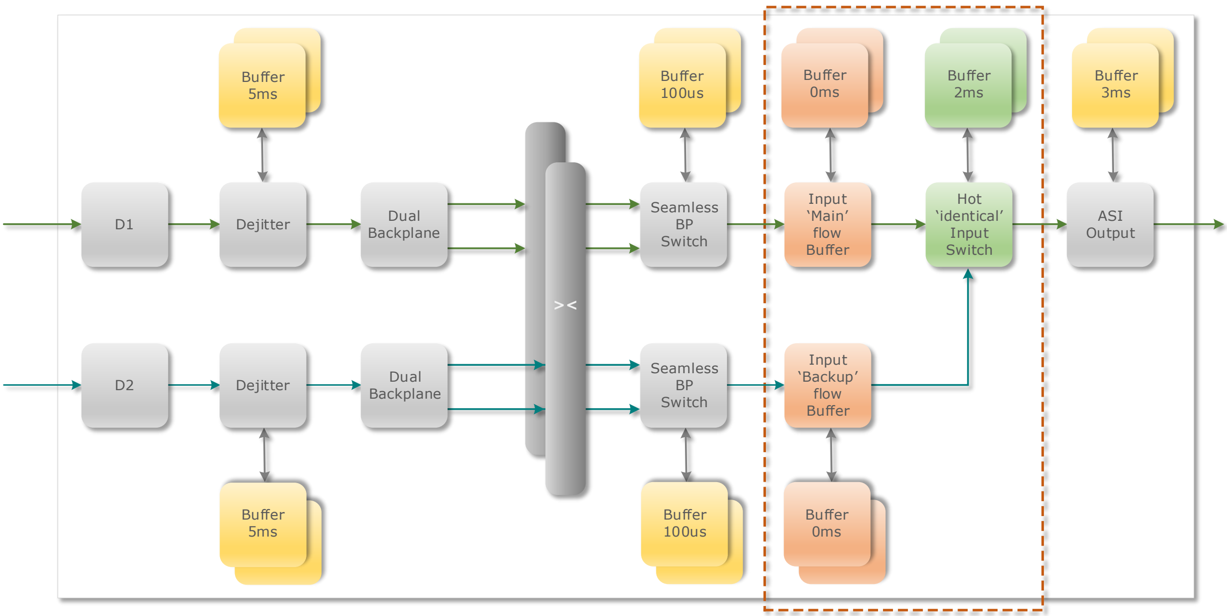

The diagram above shows the seamless backplane switch at the input to the output module, plus other key stages (e.g. De-jittering) that are critical to implement a functional system.

The stages involved with implementing INPUT REDUNDANCY have been encircled.

Basic input redundancy works by defining a main output and then defining a backup in the event the main source fails. In this mode, the following action is taken:

- When the primary source is defined, the output module will instruct the backplane to forward traffic from the primary source to the input of the output module
- As long as the output module sees healthy traffic (valid bitrate) at the input, it will use the primary source.
- If the primary source fails, the output module will instruct the backplane to deliver the backup source instead.

The logic and switching is performed in software in this mode. This means the CPU on the output card will reconfigure the output flow to be the backup service. Note: It is possible to have more than one backup service defined as a prioritised list, which the software will run through until it finds an option with a valid source.

When using INPUT REDUNDANCY in this way then the switching time will be of the order of 200ms in this mode.

To reduce the switching time noted above, the HOT STANDBY mode can be enabled. Normally, the backplane will only carry the service actually being output at any given time. When this option is enabled, the backplane will deliver all the backup options that have been defined, simultaneously. This means all sources are immediately available for the output card to switch, shaving about 50ms from the redundancy switch time and making switching performance more consistent. Note: When using this option, the available capacity from the backplane to the output module is reduced due to the extra traffic being carried to the module at all times. For MPEG_TS streams the increase of backplane traffic is not that significant so in this case it is recommended to enable this option by default.

4. Hot Standby

There is an alternative to the standard ‘software’ controlled switching mode and this is implemented by selecting the HOT STANDBY option.

In HOT STANDBY mode, the module behaves as before but the main and first backup service are routed to an FPGA for switching in hardware.

Although the FPGA switch is capable of switching any flow type, the implementation really targets MPEG_TS flows that are identical.

Within the FPGA, the fabric includes a sophisticated and configurable buffer to allow relative delays between the main and backup feeds to be compensated. Since the FPGA switch is of the seamless MPEG_TS type, main and backup sources which are identical within the TS domain can be switched seamlessly.

5. Identical Sources

The need to compensate for relative presentation delays between the main and backup feed in IDENTICAL SOURCES mode requires the implementation of a delay compensation buffer. Appear TV has provided an AUTO mode to make it easy for most customers to have this task managed automatically; the only requirement is to define the maximum relative network delay that needs to be compensated and set AUTO mode. However, manual options also exist for more advanced users wishing to achieve compensation for very long relative delays and therefore optimise buffer utilisation and / or achieve a consistent deterministic delay through the switching stage.

The following diagram shows how the IDENTICAL SOURCES option has been designed with regards to the buffers mentioned above:

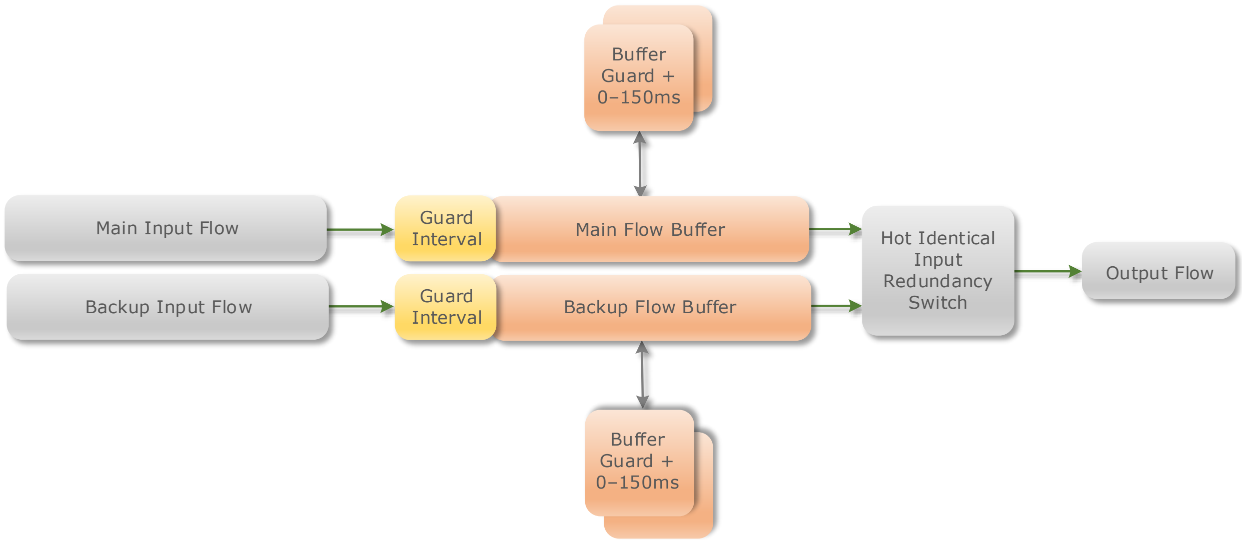

The Hot Identical Input Redundancy Switch logic contains two buffers, the main flow buffer and the backup flow buffer. Each flow buffer has a configurable minimum initial buffer filling referred to as the 'Guard Interval'. The Guard Interval is required to compensate for drift/variations in time offset between the flows. Stable inputs, e.g. ASI inputs, can have a low Guard Interval but inputs with high jitter (eg. IP inputs) may need a higher Guard Interval to compensate for swings in the CBR/PCR dejitter control loop or changes in network latency.

An example of the configuration for the HOT STANDBY option with IDENTICAL SOURCES is provided below with details on what settings are available:

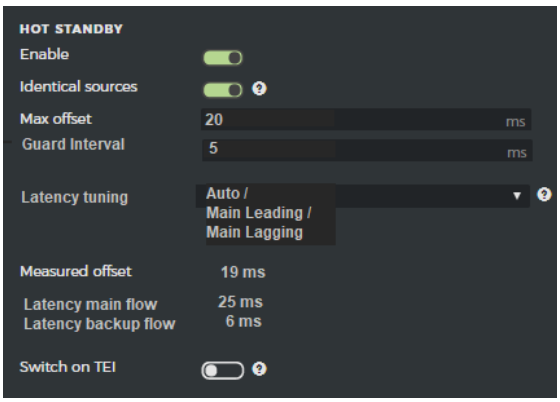

The MAX OFFSET field defines the maximum expected arrival time offset between the flows (also known as differential transmission path delay). It can be set from 0ms (streams arrive at the same time) up to 150ms.

Each flow buffer has a configurable minimum initial buffer filling, defined by the Guard Interval. The Guard Interval is needed to compensate for fluctuations in the time offset between the flows.

The "Latency tuning" field is a drop-down menu defining;

- Auto mode
    - The first flow activated will be initialized with the configured buffer setting: Max-Offset + Guard-interval (configurable Max-Offset value 150ms).
- Main Leading mode
    - The Main flow buffer will be set to the value configured in the "Max offset" + Guard-interval (configurable Max-Offset value 300ms).
- Main Lagging mode
    - The Backup flow buffer will be set to the value configured in the "Max offset" + Guard-interval (configurable Max-Offset value 300ms).

When the input flows are detected as identical the MEASURED OFFSET will be indicated (the offset in time between Main and Backup flows). If the Backup flow is leading, the value will be negative. Please note that the offset is measured at the input of the input flow buffers.

The measured latency through the Hot Input Redundancy system for the main path, including main flow buffer, will be indicated as the LATENCY: MAIN FLOW.

The measured latency through the Hot Input Redundancy system for the backup path, including backup flow buffer will be indicated as the LATENCY: BACKUP FLOW

Please note that the new MAIN LEADING and MAIN LAGGING modes will be introduced in the new 3.10 release.

6. Input Redundancy Supported Modules

The input redundancy feature is currently supported by select modules only although it will eventually be implemented on all OUTPUT modules.

Currently, it is supported on:

- ASI output module
- IP output module

It is also due to be released soon on the following modules:

- AVC/HEVC decoder module (to make main / backup decode sources possible)
- AVC/HEVC module in TRANSCODE mode where it will provide a main / backup source capability into the decoder stage of the transcoder. Additionally, n+m redundancy will also be introduced for transcoder (these two features are required to support each other).
- IP to SDI gateway module

7. Further Redundancy and Feature Options

Appear TV's X platform is a highly flexible solution that provides various other options with regards to redundancy and features. Included in this document is an overview of some of the other features that are available.

- **IP Input Policing** can be set to define a bitrate cap per IP input flow. This allows the user to cut the IP flow if the bitrate is exceeded. The feature can be useful if feeds are being processed through the chassis from multiple providers where an error on one of the feeds could cause failure to occur on a given output feed.

- **OSPF Output Redundancy** is available on the IP Output of the Management/Switch or Dual IP modules. The OSPF redundancy implements part of the OSPF protocol on the IP Output which can be added to any multicast output from the nodule. In response to alarm conditions or if a failure on the port occurs then OSPF is used to switch the IP multicast streams individually within a layer 3 router.
  The OSPF redundancy can be used alongside the Cloned IP Output.

- **Cloned IP Output** duplicates the same content over 2 IP Output ports from the same Management/Switch or Dual IP module. The concept behind this is to provide exactly the same content over 2 network paths using RTP to enable any receiving chassis to receive a seamless input. Unless duplicate traffic is carried over the network, it is combined with either the OSPF Output redundancy feature or 'Monitor In Out' redundancy.

- **Seamless IP Input** works on the receiving chassis and works at IP packet level. If the requirements for seamless operation are met, the system takes valid IP packets from either redundant distribution networks. As with the Cloned IP Output feature it is available on the Management/Switch or the Dual IP modules. The solution works by bringing in both network paths linked to the cloned output into the module. Before passing the content over the backplane in the chassis the stream is rebuilt using the two inputs. Any packets missing from one stream due to network losses will be added using the packets from the other stream. Using this method allows the user to reduce the overhead required if FEC is used.
  Note: this mechanism operates at IP packet level only. It is possible to combine this method with INPUT REDUNDANCY to provide protection at both the IP and flow domains (especially if the flow is MPEG_TS based and can benefit from the seamless MPEG_TS switch mode that INPUT REDUNDANCY now offers)

- **Mute on Error** is a way to mute failing outputs if there is a problem on the incoming stream. The feature is triggered by the chassis alarms system which is aware of the status of each service. If a critical alarm is raised against a service or its components, or the service is missing the output module will stop the affected flow so that no packets are outputted. This can be useful in triggering downstream redundancy mechanisms. An alarm will be raised to identify that this has happened.
  For the case where SPTS is used, this is a simple process but if an MPTS is used then a priority should be set on the services included in the stream. Priority levels available are High (100), Medium (50), Low (25) and None (0). For the MPTS to be muted then an aggregated priority of 100 or more is required allowing the user to still provide high value content even if the low value content has errors or goes missing.

- **Monitor-In-Out** has been designed to work alongside the Mute On Error functionality available in the X platform. The solution is a simple and convenient way to implement a 1+1 redundancy configuration between chassis, or even a single chassis fitted with 1+1 redundant Management/Switch modules. An example of how Monitor-In-Out can be used is shown below.

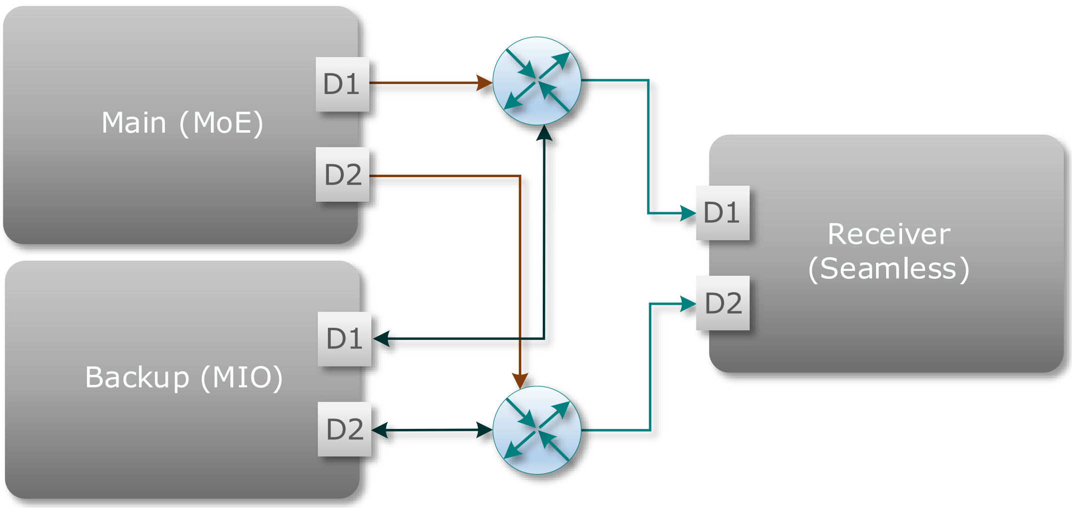

The backup chassis due to running exactly the same configuration as the main chassis will have all the output feeds muted. The Monitor-In-Out option will then monitor the output feeds from ports A and B from the main chassis and if a flow goes missing due to MoE or a network failure then the backup chassis will lift the mute for the required flow and output the content instead. When the main flow returns then the backup chassis will mute the output so that the main feed is active again.

It is also possible to enable Mute on Error alongside the Monitor-In-Out option on the backup chassis so if there are errors on the content from both chassis the output is not passed through the network.

As can be seen Appear TV's X platform is a feature rich solution when it comes to redundancy. Having so many options available can be a little overwhelming so included below are a couple of examples of how these options can be used:

1+1 with OSPF based switching

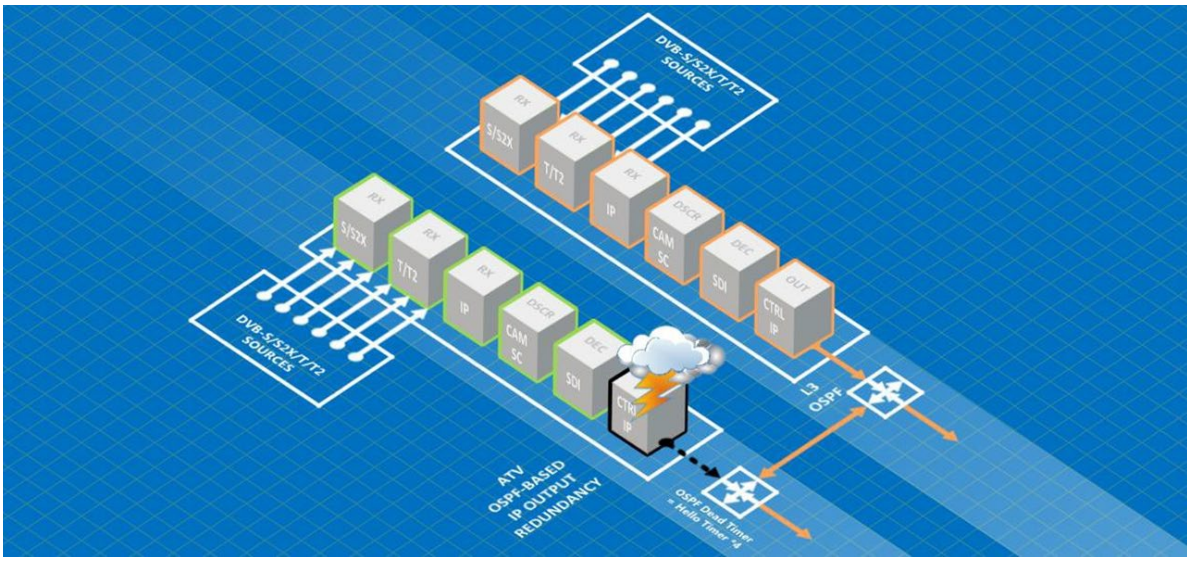

Shown above is an example of a 1+1 system that is using OSPF Output Redundancy. In this scenario both the main and backup chassis are running exactly the same content into the network. The backup chassis is configured in a way either via the chassis setting or the network infrastructure to have more OSPF cost then the main chassis. In this example failure in the main chassis of the CTRL/IP module will essentially drop the content from the network and the Layer 3 network will automatically take the content from the backup chassis.

1+1 with OSPF and Seamless Protection

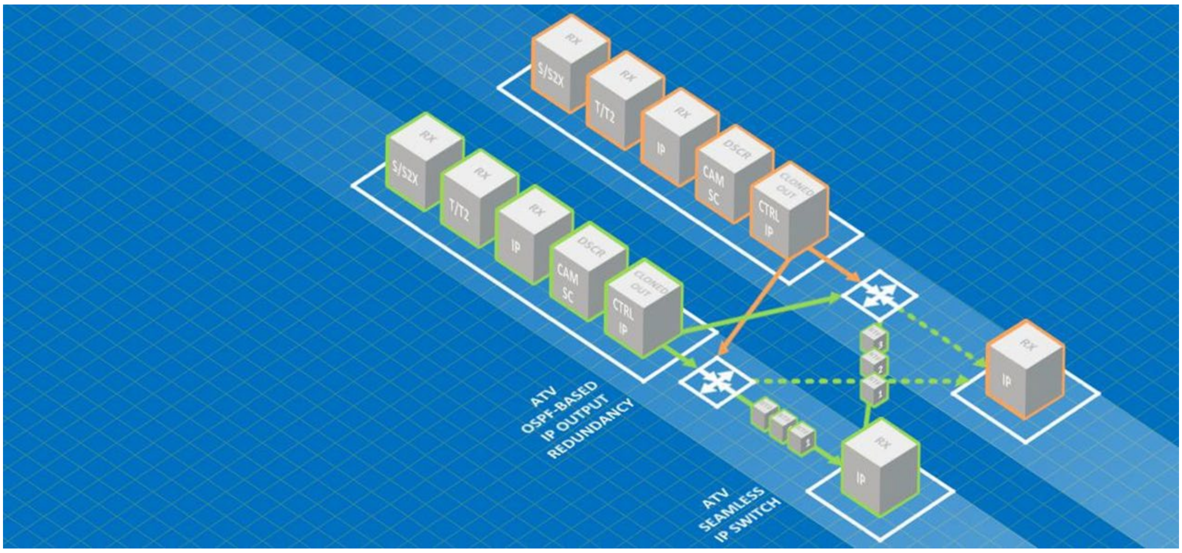

Shown above is an example of a 1+1 system that is using OSPF Output Redundancy alongside Seamless Protection. Configuration wise it is similar to the previous example with the addition of adding Cloned IP Output to both the main and backup chassis and running Seamless Input on the receiving chassis. The addition of Seamless Protection introduces a higher level of resilience in the network with regards to packet loss that may occur when the cloned content is distributed over two different networks to the receiver. As before the backup chassis should be configured the same as the main chassis but with a higher OSPF cost on the cloned outputs feeding the networks.