TLXpress™ Network Guide

TechLogix Networx

Table of Contents

TechLogix Networx's implementation of BlueRiver technology

TechLogix Networx uses AptoVision's BlueRiver NT+ chipsets in the TL-IPFO series of products. These chipsets support up to 18 Gbps HDMI video with HDR and HDCP 2.2. While there are copper-based 10 Gb Ethernet switches, the TL-IPFO series utilize duplex fiber optics, which, in turn, means the Ethernet switches deployed in an installation must also be fiber-based.

The TL-IPTP-CI runs the TLXpress™ Server and web GUI software, which includes the Aptovision Server.

What is BlueRiver and what is it for?

BlueRiver NT+ is an FPGA-based chipset used for packetizing digital audio/video signals and transporting them across standard 10 Gb Ethernet networks. This chipset is deployed into "end-point" devices which, generally, serve as either transmitter/encoder (video to Ethernet) devices or receiver/decoder (Ethernet to video) devices. The AptoVision Server is a piece of code that runs on a general-purpose computer (*e.g.*, a Windows PC, a Raspberry Pi-class device, or other CPU as chosen by the OEM manufacturer). The purpose of the AptoVision Server is to provide an API layer that presents a simplified, Telnet-based interface to higher level control software, and translate commands into direct low-level control of the endpoint devices. This document will refer to the entire system of BlueRiver NT+ and AptoVision Server collectively as "TLXpress™."

TLXpress™ uses little or no compression on high resolution video signals, and therefore can place significant demand on an Ethernet network. Several key requirements of the TLXpress™ system must be considered in the design and commissioning of the 10 Gb network. This document's purpose is to explain those requirements, and help to ensure a smooth system bring up, even in large and complex Ethernet environments.

Note that this document provides a lot of *what*, but very little *how*. The target audience for this documentation is the IT specialist or network architect charged with designing and configuring the Ethernet network. No piece of TLXpress™ documentation can replace having the right level of expertise in configuring and operating the Ethernet switches you select, and there is no point trying to duplicate those switches' operations manuals, either. This document therefore presumes a high level of expertise in network infrastructure design and configuration, and serves only to explain the requirements that TLXpress™ has for such a network.

Basic elements of TLXpress™ communication

TLXpress™ uses layer 2 / layer 3 IPV4 protocols.

Audio and video (multicast)

The focus of TLXpress™ is transporting high resolution audio and video data across 10 Gb Ethernet networks. A typical video data stream consumes three to nine gigabits per second, depending on its resolution format. In order to manage this bandwidth, the basis of TLXpress™'s AV transmission protocol is multicasting. This way, these high data rate streams are only sent through ports across links where they are needed. Some of the most serious considerations for network deployment – especially in multi-switch environments – are around ensuring that the network is setup to handle this properly.

Bandwidth management is a key design consideration. Any significant loss of data (due to oversubscription or other reason) will result in visible on-screen problems. Ensuring that high data rate streams are routed *only* where they are needed is critical.

Control: IR, RS-232 (unicast with some optional multicast and broadcast)

TLXpress™ carries various low-speed control communications as well. These signals are typically triggered by a user-facing control system and are used to control things like turning on and off a display. Infrared signals, RS-232 signals, and consumer electronics control (CEC) signals are all possibly included here. These packets are typically sparse, and data rates are in kilobits per second. Most commonly the data packets are unicast between the TLXpress™ Server and TLXpress™ endpoints. Sometimes the communication is directly between endpoints. It is possible to encounter broadcast traffic here as well, although this should be rare in larger installations. The purpose of broadcast is to, for example, send a single infrared command to all TLXpress™ endpoint devices simultaneously. Finally, a multicast path can be opened up to transmit a command to "all receivers" or "all transmitters." Use of this feature is entirely optional.

Discovery (Broadcast)

TLXpress™ runs its own auto-discovery mechanism. It relies on broadcast communication between all endpoints and the TLXpress™ Server.

The 1 Gb port

The TLXpress™ chipsets include built-in Ethernet switching capability, in order to provide 1 Gb connectivity that can be piped through the 10 Gb link.

Switch selection and network topology

TLXpress™ is compatible with any 10 Gb Ethernet switch that has layer 3 switching capabilities. Support for multicast and IGMP v2 with IGMP snooping is required.

In general, complex multi-switch systems will rely on switches from reputable vendors with good documentation. At times, it will be beneficial to know what is really going on in a switch (*e.g.*, how does it process multiple streams into an aggregated link?), and having well understood behavior will be important.

One example of bad behavior we have observed from an "off-brand" switch: When oversubscription to a particular port is detected, rather than throttling traffic (dropping packets) at that egress port, the switch throttled traffic at its input port – causing that traffic to be lost to all destinations – even those which were not oversubscribed.

Another simple "gotcha" to avoid: most switches' default behavior is to broadcast multicast packets. Watch out for this and make sure to enable multicast and IGMP before trying to use TLXpress™.

Regarding trunk ports

QSFP ports can usually be configured as single 40 Gb trunk port or as four independent 10 Gb ports. If you intend to use the port as a high bandwidth link to another switch, be sure to configure it as a 40 Gb trunk port. Otherwise you're just adding extra 10 Gb ports to your switch, and if you connect them all to the same other switch, you made a loop, and 3 ports will get shut down, leaving you with only 10 Gb between the switches.

Regarding "stackable switches"

A stackable switch works with other stackable switches to present themselves as one cohesive "single switch." The entire system can be easily configured from a single IP address. This type of system is compatible with TLXpress™, but note that bandwidth must still be managed. Typically, a system of stackable switches is *not* fully non-blocking, meaning that there will be bottlenecks (often 40 Gb links between switches with 24, 48, or more 10 Gb ports). Bandwidth demands of the TLXpress™ system must be compared against the user requirements (how much video must be routed over the stacking ports to meet the user requirements) and all this must be considered in the design of the system's connectivity.

Regarding "leaf-spine"

Leaf-spine comes into play when number of endpoints grows beyond 100-200 end-points. TLXpress™ devices connect to leaf switches and leaf switches connect to spine switches. Leafs mesh into spines meaning that every leaf is connected to every spine. This is compatible with TLXpress™ but like for stackable switches, bandwidth has to be managed where the limiting factor will be the trunk bandwidth between leaves and spines.

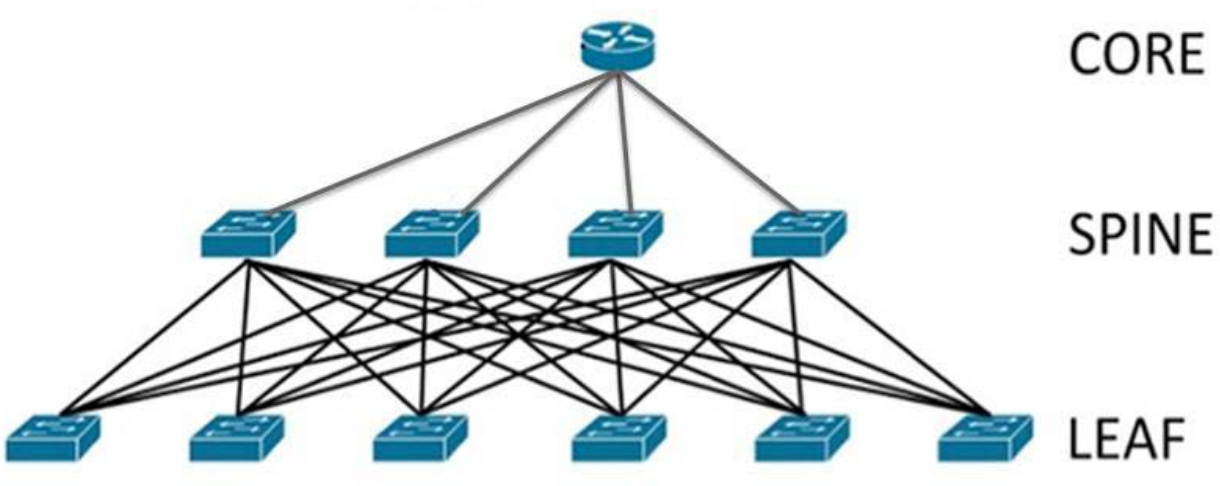

Figure 1: An illustration of leaf-spine topology

Multicast management

In order to minimize overall bandwidth consumption, TLXpress™ relies on multicast routing for distributing audio and video data. The basic idea is to *only* send AV data through switch ports where it is needed. As few as two video streams at 4K can oversubscribe a 10 Gb link, so getting multicast right is critical.

Multicast source addresses

The AptoVision Server is responsible for assigning multicast source addresses to transmitters. Each transmitter may use one or more source addresses. Separate source addresses are used for:

- Video, including the embedded audio from an HDMI source
- Audio that has been extracted from the HDMI stream
- Audio brought into the transmitter via I2S (*e.g.*, analog audio input)

Note: for a receiver in fast switch mode to output audio, it must receive that audio from one of the secondary channels. Therefore, a system using fast switch mode will generally need two multicast source addresses per transmitter.

TLXpress™ Server may be configured to distribute multicast source addresses automatically. In this mode, TLXpress™ Server assigns addresses in the range of 224.1.1.1 to 224.1.1.255. This range must not be used by any other devices on the network.

Regardless of whether the addresses are assigned automatically by TLXpress™ Server, or under control of higher level control software, TLXpress™ transmitters may only be configured with addresses in the range of 224.1.1.1 to 224.1.1.255. The last three addresses are reserved. This means that the maximum number of addresses that can be configured for streaming audio or video content is 252.

Important note: the switch must be configured to drop any packets from a multicast stream with no subscribers. Some switches could treat such packets as broadcast and impact the bandwidth and performance of the entire network.

Multicast routing management

In a single switch environment, IGMP suffices to enable the switch to understand which traffic must be routed to which ports. In a TLXpress™ system, the switch will only route a transmitted video to those ports where a TLXpress™ receiver has requested a particular stream. In a multi-switch environment, higher-level protocols must be employed so that "switch X" can understand which streams are demanded by receivers connected to "switch Y." If this is not carefully managed, trunk links can easily be oversubscribed by multiple (dozens or more) streams of 8 Gb video.

The basic communication flow is:

- Switch X is connected to switch Y via a 40 Gb trunk link. No video traffic is currently passing between them.
- Transmitter 1 is connected directly to switch X, and sending stream A into that switch.
- Switch X needs to notify switch Y that stream A exists. Stream A still does not traverse the trunk link.
- Receiver 1 is also connected directly to switch X, and makes an IGMP request to switch X for stream A.
- Switch X begins routing stream A to receiver 1, who displays the video signal. Still there is NO video traffic across the trunk port between switches.
- TLXpress™ receiver 2 is connected directly to switch Y.
- TLXpress™ receiver 2 stream A by issuing an IGMP request to switch Y (which receiver 1 is connected to directly).
- Switch Y realizes that it does not have access to stream A, but switch X has notified switch Y about stream A.
- Switch Y sends a request to switch X to provide stream A across the trunk link.
- Switch X complies, begins routing stream A across the trunk link, and switch X begins providing stream A to receiver 2.
- Later, if receiver 2 releases his IGMP subscription to stream A (and no other receivers on switch Y have requested stream A), then stream A should be removed from the trunk link.

The implementation of these higher-level protocols is vendor specific. One example of such a protocol is Protocol Independent Multicast (PIM), implemented by Extreme Networks and others.

Broadcast management

TLXpress™ implements a device discovery mechanism that relies on broadcast traffic to the broadcast address of the zero network (255.255.255.255). These discovery packets must reach all ports in the TLXpress™ system for the device discovery system to function.

The mechanism itself is very simple. Each device broadcasts a hello packet every two seconds. Any device on the network may collect a list of active devices and populate a list of who is available.

Optionally, some IR and RS-232 traffic is also routed with this broadcast mechanism.

Network performance issues

Latency and network jitter

In general, TLXpress™ is not sensitive to network latency. However, TLXpress™ can be sensitive to latency variation (jitter). An endpoint in genlock mode is more sensitive than an endpoint in fast switching mode. Fast switching mode relies on a frame buffer, which smooths out any network jitter of up to around a frame of video (16 milliseconds for 60 fps video). Genlock mode can compensate for the jitter induced by a single buffer in a single Ethernet switch, but a genlock mode receiver's HDMI output jitter will increase with network jitter in multi-switch configurations. Practically speaking, a TLXpress™ receiver should always be set to fast-switch mode if there is more than one network switch separating it from its paired transmitter.

Future plans for BlueRiver NT+ include the use of network timing protocol gPTP, which will dramatically improve the system's sensitivity to network jitter and allow genlock mode to work across multiple hops. This new feature will not place any new demands on the switches.

Packet ordering

TLXpress™ relies on packets being delivered in-order. This is especially critical for 10-bit and 12-bit video formats. There is some tolerance built in for out of order packets of 8-bit video, and it is therefore recommended to mostly use 8-bit video in complex multiple switch networks where reordering is possible.

Methods to ensure in-order delivery across aggregated links or multiple (different) routes will vary by switch manufacturer and network architecture. TLXpress™ will in the future offer improved resilience for 8-bit modes to out of order packets.

Bandwidth management

The TLXpress™ video system has no tolerance for link oversubscription. Lost packets translate to lost pixels on screen. Video data is transmitted via UDP, so there is no retransmission (since retransmitted packets would be too late anyway). There is intelligence built in to mask small errors (by filling in surrounding or previous pixels, etc.), but any significant loss of data will result in significant image problems on screen.

The easiest solution is to design a fully non-blocking network. This is relatively straightforward for smaller systems, with affordable "top of rack" switches ranging to around 100 non-blocking ports. Beyond that size, blade-based systems of non-blocking switches exist, but can be expensive. A spine and leaf architecture may be more affordable, and also may fit more neatly in the physical layout of the network.

To design a cost-efficient system, the video routing use cases must be carefully considered and weighed against bandwidth availability. If true "any source to any destination at any time" functionality is required, then a non-blocking architecture is required. Realistically,

most video needs tend to be local (inside a room), with fewer streams traversing between rooms or buildings at a time.

The following table gives a summary of how much data a TLXpress™ video stream consumes. The data rate does scale up and down with video format (resolution, frame rate, etc). It is critical that the network is designed to handle the worst-case routing scenario demanded by the use cases. Special attention must be paid to the bottlenecks – the 40 Gb trunk ports between 10 Gb switches.

Resolution	Frame rate	Bit depth	Chroma	Ethernet bandwidth consumed (Gbps)	Note
640x480	60p	8-bit	4:4:4	0.50	
1280x720	60p	8-bit	4:4:4	1.43	
1920x1080	60p	8-bit	4:4:4	3.21	
1920x1080	60p	8-bit	4:2:2	2.14	
3840x2160	30p	8-bit	4:4:4	6.42	
3840x2160	60p	8-bit	4:4:4	8.57	Compression is enabled
3840x2160	60p	10-bit	4:2:0	8.03	
3840x2160	60p	10-bit	4:2:2	8.57	Compression is enabled
4096x2160	60p	8-bit	4:4:4	9.13	Compression is enabled

Table 1: Ethernet bandwidth consumed by different video formats

The TLXpress™ 1 Gb port

The TL-IPFO products include a 1 Gb “courtesy port” whose traffic is piped back through the main 10 Gb AV port of the endpoint. The BlueRiver NT+ chipsets include an Ethernet switch built in, which is how this traffic is connected back to the 10 Gb network.

A few notes on this port:

- Due to a limitation of the Realtek switch IC used, the 1 Gb port does not support jumbo packets.
- In current implementations, there is no VLAN or priority assigned to the 1 Gb port
- TLXpress™ does not implement any form of STP or loop protection. The 1 Gb port must never be looped back to the same switch as the 10 Gb port.

Because there is no priority assigned to this port, keep in mind that traffic from this port may trigger an oversubscription condition and cause video failures. This can be especially bad through trunk links. Consider a 48-port 10 Gb switch with a single 40 Gb uplink. Potentially 48 TLXpress™ endpoints could be connected, each with some Ethernet device on the 1 Gb port. If all of these devices were to maximize their bandwidth consumption (1 Gb each), that would be 48 Gb consumed without a single video link in place. Use of the 1 Gb port must be very carefully considered in complex multi-switch Ethernet networks.

Required switch characteristics

Switch Speed

TLXpress™ requires the switch to be a 10 GbE switch.

BlueRiver NT+ technology is used to transmit uncompressed video up to 4K along with other AV signals such as audio and control signals. For video alone, it means raw bandwidth of about 4 Gb/sec for HD and over 8 Gb/sec for 4K 60Hz.

Packets Routing

To enable the transmission of a source to multiple destination, TLXpress™ devices make use of Multicast. The default behavior of layer 2 Ethernet switch is to broadcast those packets which means that every packet is transmitted to all possible destinations. This is why any network switch used with TLXpress™ has to support IGMP Snooping. TLXpress™ end points use IGMP protocol to assign the end points into multicast groups and the router uses IGMP snooping to efficiently route multicast packets only to receivers that want to receive them.

Many switches have the IGMP Snooping feature disabled by default and manual configuration is required. Often, a simple check mark near “Enable IGMP Snooping” is the only thing needed to enable IGMP Snooping.

However, the implementation of IGMP Snooping is vendor specific and more configuration is often needed.

An Ethernet switch can be informed that a device want to leave a multicast channel by sending it a IGMP LEAVE GROUP packet. Once received, the time it takes for the switch to apply the new configuration may vary from one switch to the other. Most switches implement and include FASTLEAVE configuration option. When enabled, it takes much less tie for a particular port to leave a multicast group to assign the port to a different multicast group. The end result is a noticeably shorter video switching time. TechLogix recommends to always enable the FASTLEAVE option when available. With FASTLEAVE option, seamless switching is possible for 4K video sources. Without FASTLEAVE option, 'seamless' switching is limited to 1080P 60 Hz video signals.

Ethernet Switch Configuration

The following list includes typical network switch configuration options that TechLogix Engineers have come across so far. Look for these or similar options when configuring your switch.

- Enable IGMP Snooping
  - Must be enabled
  - Sometimes IGMP Snooping is a global parameter that applies to all VLANs
  - More often IGMP snooping is enabled for a particular VLAN, i.e. VLAN 1.
- Filter/Drop unregistered Multicast traffic
  - If not applied, the behavior of the switch will be to broadcast multicasted packets if the switch has no known destination for that packet.
  - Must be enabled if found
- Unregistered Multicast Flooding
  - Must be disabled if found
- Enable IGMP Querier
- Enable IGMP Querier on VLAN1
  - Should be enabled if found.
- Set IGMP Version to IGMP V2
  - Must be set if found
- Enable FASTLEAVE on port X
  - Should be enabled if found

TLXpress™ Network Environment Specifications

IP Addresses and Subnets

TLXpress™ devices must be on the same IP subnet because parts of the control protocol rely on IP broadcast. Devices can have their IP address and subnet mask configured either through a DHCP server or manually.

If a device is configured for DHCP address allocation and no DHCP server can be reached, the device falls back to self-allocation using the Automatic Private IP Addressing (APIPA) scheme, also commonly called auto-IP. This scheme allocates addresses in the 169.254.0.0/16 range (i.e. 169.254.x.x with a subnet mask of 255.255.0.0). If a device is configured for manual IP address allocation, addresses are set by client software with API commands. IP address collisions must, of course, be avoided.

Clients can be on any subnet as long as they can connect to TLXpress™ Server.

Port Numbers

The table below lists port numbers used by TLXpress™ Server. These ports must not be used by other services on the same host as TLXpress™ Server and must not be blocked by firewall rules. In the case of the UDP ports, also ensure that these port numbers are not in use by other applications on your network as some data is sent to broadcast IP addresses on those ports.

Protocol	Port Number	Description
TCP	6970	Telnet/TCP connections between clients and TLXpress™ Server
UDP	6969	Used by TLXpress™ Server and TLXpress™ devices to exchange control information
UDP	10001 to 10004	Used by TLXpress™ Server and TLXpress™ devices to exchange RS-232 data

Multicast IP Addresses

Audio and video data is exchanged between TLXpress™ devices using IP multicasting. The table below specifies the multicast IP address range used by these devices. This address range is reserved and must not be used by other network applications.

Address Range	Description
224.1.1.1 to 224.1.3.255	Video and audio data

NOTE: Data traffic sent to these multicast addresses has very high bandwidth. Ensure the network switch on which TLXpress™ devices are connected is properly configured (e.g. IGMP snooping) to prevent this traffic from leaking to the rest of your network.

Connecting to the TLXpress™ Server

The TLXpress™ server is configured to operate in DHCP mode by default. To access the web GUI, navigate to http://tlxpress.local in a Google Chrome web browser.

If using a Windows operating system to connect to the TLXpress™ Server, it may be necessary to download and install Bonjour Print Services in order to access the web GUI. This is not necessary if iTunes is installed on the PC, because Bonjour is included with iTunes. Bonjour can be downloaded from Apple at http://support.apple.com/kb/DL999.