


The Successful Implementation of ADDERLink INFINITY

Optimizing your network for ADDERLink INFINITY



The Successful Implementation of ADDERLink INFINITY

Optimizing your network for ADDERLink INFINITY



SINCE ITS INTRODUCTION IN 2010, the ADDERLink INFINITY system (otherwise known as ALIF), has provided new levels of flexibility in the distribution of high-performance computing power. The reliable transfer of high-resolution DVI video, audio, RS232 serial and four USB streams across a single run of network cable is no mean feat; sending all this across a network even more so.

ALIF units are fully compliant with the open Gigabit Ethernet standards, however, in order to achieve success, it is necessary to ensure that the network is structured and configured appropriately. It's about maximising throughput, banishing bottlenecks and eliminating weak links.

This guide provides an overview of the key elements, suggestions for success and assistance with potential problems.

CONTENTS

- Choosing Layer 2 Switches
- Designing the Network Layout
- Troubleshooting
- Glossary
- Switch Setups

Choosing Layer 2 Switches

This section provides various general and specific recommendations for switches to use with ALIF, however, there is no substitute for testing in real world situations. If you are in doubt about which network switch to choose, the safest approach is check the switch specs are equal to or better than the requirements.

General Recommendations

There are certain minimum features that you need to ensure:

- Gigabit (1024Mbps) or faster Ethernet ports,
- Support for IGMP v2 (or v3) snooping,
- Support for Jumbo frames (packets) up to 9216-byte size,
- High bandwidth trunk connections between switches, preferably Fibre Channel.

Additionally:

- Look specifically for switches that perform their most onerous tasks (e.g. IGMP snooping) using multiple dedicated processors - i.e. the tasks are carried out in custom ASIC hardware rather than software routines on a general processor.
- Check the maximum number of concurrent snooperable groups each switch can handle and ensure that they meet or exceed the number of ALIF transmitters that will be used to create multicast groups.
- Check the throughput speeds of the switch. Ensure that each port is full duplex (i.e. bi-directional communication) and that the up and down stream data speeds for each port are 1 Gigabit per second. Put simply the total Switching capacity (measured in Gbps) should be equal to or greater than $(n)2\text{Gbps}$ where (n) is the number of Gigabit ports. If you have your switch specs, enter them into the Suitability Calculator to see if they meet this requirement.
- Another major consideration is switching capacity (measured in millions of packets per second) This should be equal to or greater than $(n)1.448\text{mpps}$ where (n) is the number of Gigabit ports. If you have your switch specs, enter them into the Suitability Calculator to see if they meet this requirement.
- Wherever possible, to avoid conflicts, use the same switch manufacturer throughout a single subnet and, if possible, the same model of switch - this

The trouble with multicasting

Where an ALIF transmitter is required to stream video to two or more receivers, multicasting is the method used. Multicasting involves the delivery of identical data to multiple receivers simultaneously without the need to maintain individual links. When multicast data packets enter a subnet, the natural reaction of the switches that bind all the hosts together within the subnet, is to spread the multicast data to all of their ports. This is referred to as Multicast flooding and means that the hosts (or at least their network interfaces) are required to process plenty of data that they didn't request. IGMP offers a partial solution. IGMP (Internet Group Management Protocol) was developed to help prevent flooding by requiring individual hosts

to opt into multicasts. It also provided a mechanism for routers to determine whether any hosts located within their subnet still wished to receive the multicasts. However, this only has an effect at the gateway to the subnet – so if one host requests a multicast, all hosts within the subnet would also receive it.

A solution: IGMP snooping

Cue a development in the switches that glue together all the hosts within a subnet: IGMP snooping. IGMP snooping means these layer 2 switches now have the ability to take a peek at the instructions that help routers to do their job. As a result, the switches can then determine exactly which of their own hosts have requested to receive a multicast – and only pass on multicast data to those hosts.

will simplify configuration.

- When choosing Layer 3 switches for the network, at least one must be capable of operating as an IGMP Querier.

IGMP and internal switch design

In recent years, the number of Layer 2 switches that support IGMP snooping has proliferated; however, there is a wide variance in performance between the most effective and the least. In order to take a peek at (snoop) IGMP messages, Layer 2 switches are required to do something they were not originally designed for: Deciphering every data packet at Layer 3 in order to read the logical addressing and multicast instructions. This requires considerably more processing horsepower than their normal day-job of reading physical MAC addresses at Layer 2.

Implementing IGMP snooping on a low end switch with a slow processor can cause severe performance problems when data is transmitted at high data rates and/or there are multiple IGMP groups to be monitored. If a switch cannot keep pace, it will cause backlogs where large numbers of data packets are arbitrarily discarded and/or it resorts to sending all multicasts to all ports - causing multicast flooding. Either way, this results in slow video updates and a poor user experience.

A quick note about Layer 3 switches and multicast routing

A key component of subnets used for multicast distribution is the multicast router. This device plays a vital coordinating role in ensuring that network traffic is delivered to the correct Layer 2 switches and the hosts connected to them. Increasingly, the role of router is being fulfilled by Layer 3 switches, particularly for private networks that do not require links to wider external networks. When selecting a Layer 3 switch for your network, ensure that it can operate as an IGMP Querier and also has sufficient capacity for the size of your subnet.

Specific switches known to work well with ADDERLink INFINITY

The following models have been tested and found to produce favourable results:

Layer 2 switches

- Cisco 2960
- Cisco 3750 (Superseded by Cisco 3850)
- Cisco 3850

- Cisco 4500
- Cisco 6500
- Extreme Networks X480
- HP Procurve/Aruba 2530
- HP Procurve/Aruba 2810 (Superseded by 2530)
- HP Procurve/Aruba 2910 (Superseded by 2920)
- HP Procurve/Aruba 2920
- H3C 5120

Layer 3 switches

- Huawei Quidway s5328c-EI

For specific details on configuring some of these switches for correct operation with ADDERLink INFINITY please refer to the Switch setups section.

Designing the network layout

Selecting the right kind of switches is the first stage; ensuring they are laid out correctly is the next.


General recommendations

- Choosing your topology is a start of life procedure (factory reset of the system may be required in order to make significant changes later).
- Keep the distances between the switches as short as possible.
- Sufficient bandwidth must be provisioned between switches to prevent bottlenecks.
- Avoid using VGA-to-DVI converters, where possible, instead replace VGA video cards in older systems with suitable DVI replacements. This is because VGA-to-DVI converters create sufficient background noise (even in static images) that an ALIF transmitter would be forced to send all of every video frame, thus creating large amounts of unnecessary network traffic. ALIF1002 and ALIF2000's models have a feature called Magic Eye that helps to reduce the noise.
- Where practical, dedicated network hardware should be used for Infinity. Alternatively, as a minimum, dedicated VLAN's should be created minimise competing traffic from other network devices.


Layouts

Tree or Pyramid layout is now supported, it's worth considering, when


provisioning inter-switch bandwidth, that Each channel connection between switches will require up to 1 Gbps. It is also now possible to configure a multi-subnet system, where each subnet or VLAN contains an independent DHCP server and is capable of managing that network segment in the event it becomes isolated from the primary subnet.



Line or Flat topology is also a viable option but, bearing in mind the requirement for each connection requiring a Gb connection, the interconnecting bandwidth requirement can soon become very high.



Stacked topologies are common and beneficial as the stacking cables between switches often have high bandwidth capacity available. The same rules apply regarding 1Gbps bandwidth required for each connection between switches.



The streaming output from any transmitting device must travel up to the Querier, even if a recipient is located on the same switch as the transmitter. Given that the top layer 2 switch must also pass up and down all data from the lower switch, this is the reason for its double bandwidth link to the top level switch. The links between switches could either be multiple Gigabit Ethernet connections or, more favourably, fibre channel connections. Due to the interlink bandwidths required to prevent bottlenecks, it is recommended that there are no more than two cascade levels in any single line cascade arrangement.

Configuring switches and ALIFs

The correct configuration of each switch can have a significant impact on overall performance, so in this section we'll take a look at the various issues. The precise configuration details for each switch vary with make and model, so please also check the Switch setups section where a selection of specific configuration details are made available for a limited number of common switches.

Enable IGMP snooping

IGMP snooping is an essential and effective tool to prevent multicast data from flooding the switches within a subnet. Ensure that IGMP snooping is enabled on every switch. Ensure that the IGMP Fast-Leave option is enabled on all switches that have ALIF units connected directly to them. This will allow the switches to respond quicker to changes in multicast arrangements. As an alternative to IGMP, some Cisco switches also support a proprietary standard called CGMP (Cisco Group Management Protocol) which is not directly compatible with IGMP but achieves a similar outcome. CGMP sends multicast group messages at layer 2 and so switches can natively read these without needing to perform IGMP snooping. The use of CGMP is only possible if all of the switches within your network are Cisco units that support this standard.

Enable Jumbo frames (Jumbo packets)

For the majority of video resolutions that are transmitted by ALIF units, standard Ethernet frame sizes are used to transfer video data quickly and efficiently. However, certain video resolutions cause issues which require ALIF units to instead output Jumbo frames in order to correctly convey the necessary data. To efficiently handle these enlarged frames, the switches within the subnet need to have their Jumbo frames options enabled. If any of the computers attached

to your ALIF transmitters are using/or will use any resolution that uses 2048 horizontal pixels (e.g. 2048 x 1152), ensure that Jumbo frames are enabled on all of the switches within the subnet.

Enable Spanning Tree Protocol

Multiple and redundant links between switches are important for both transfer speeds and network resilience. However, unless carefully managed, such multiple links can cause bridge loops to occur and that is what the Spanning Tree Protocol (STP) is designed to prevent. However, that's not the end of the story because STP can cause a new issue. It temporarily blocks newly found network links for tens of seconds to ascertain their function, often causing problematic timeouts. This also causes a problem for device discovery when using AIM, whereby ALIF units cannot be correctly configured. Therefore, it is important to enable STP on every switch and also enable a technique known as portfast on every switch port that is not a link to another switch.

Choose the most appropriate switch forwarding mode

Every switch is built to transfer data from any one of its ports to any other as quickly as possible. Of the numerous factors involved to make this possible is the forwarding mode: the precise manner in which data packets are deciphered, checked and transferred. For use with ALIF units, the Cut-through method generally produces the fastest results. However, if the network produces numerous data errors then it may be necessary to use a Store and forward mode instead. On higher grade switches, the latter should not impact performance too greatly.

Alter ALIF transmitter video settings, if necessary

Within each ALIF transmitter, there are various settings that affect the manner in which video is sent. Below is a brief summary of the common combinations. These settings can be accessed either through the transmitter tab of AIM or directly via the ADDERLink INFINITY browser-based configuration utility. Please refer to the relevant Adder user guides for details.

- If colour quality is important, then leave **Colour depth** at 24 bits and adjust other controls,
- If moving video images are being shown frequently, then leave **Frame Skipping** at a low percentage and instead reduce the **Peak bandwidth limiter** and **Colour Depth**.

- Where screens are quite static, try increasing the **Background Refresh** interval and/or increasing the **Frame skipping** percentage setting. Make changes one at a time, in small steps, and view typical video images so that you can attribute positive or negative results to the appropriate control. Please refer to Glossary > ALIF transmitter video settings

Ensure that all ALIF units are fully updated

Early versions of ALIF firmware (pre v2.1) exhibited an issue with the timing of IGMP join and leave commands that caused multicast flooding (i.e. network congestion) in certain configurations. Ensure that all ALIF units are fully updated to the latest firmware version.

Troubleshooting

Problem: The video image of the ALIF receiver shows horizontal lines across the screen.

This issue is known as Blinding because the resulting video image looks as though you're viewing it through a venetian blind. When video is transmitted by ALIF units, the various lines of each screen are divided up and transmitted as separate data packets. If the reception of those packets is disturbed, then blinding is caused. The lines are displayed in place of the missing video data packets. There are several possible causes for the loss of data packets:

- Incorrect switch configuration. The problem could be caused by multicast flooding, which causes unnecessary network traffic. This is what IGMP snooping is designed to combat, however, there can be numerous causes of the flooding.
- Speed/memory bandwidth issues within one or more switches. The speed and capabilities of different switch models varies greatly. If a switch cannot maintain pace with the quantity of data being sent through it, then it will inevitably start dropping packets.
- One or more ALIF units may be outputting Jumbo frames due to the video resolution (2048 horizontal pixels) being used. If jumbo frames are output by an ALIF unit, but the network switches have not been configured to use jumbo frames, the switches will attempt to break the large packets down into standard packets. This process introduces a certain latency and could be a cause for dropped packets.
- One or more ALIF units may be using an old firmware version. Firmware versions prior to v2.1 exhibited an issue with the timing of IGMP join and

leave commands that caused multicast flooding in certain configurations.

Remedies:

- Ensure that IGMP snooping is enabled on all switches within the subnet.
- Where each ALIF unit is connected as the sole device on a port connection to a switch, enable IGMP Fast-Leave (aka Immediate Leave) to reduce unnecessary processing on each switch.
- Check the video resolution(s) being fed into the ALIF transmitters. If video resolutions with 2048 horizontal pixels are unavoidable then ensure that Jumbo frames are enabled on all switches.
- Check the forwarding mode on the switches. If Store and forward is being used, try selecting Cut-through as this mode causes less latency on lesser switch designs.
- Ensure that one device within the subnet is correctly configured as an IGMP Querier, usually a multicast router.
- Ensure that the firmware in every ALIF unit is version 2.1 or greater.
- Try adjusting the transmitter settings on each ALIF to make the output data stream as efficient as possible. See Alter ALIF transmitter video settings if necessary for details.

Problem: The audio output of the ALIF receiver sounds like a scratched record.

This issue is called Audio crackle and is a symptom of the same problem that produces blinding (see left). The issue is related to missing data packets.

Remedies:

As per blinding discussed above.

Problem: AIM cannot locate working ALIF units.

There are a few possible causes:

- The ALIF units must be reset back to their zero config IP addresses for AIM discovery. If you have a working network of ALIF's without AIM and then add AIM to the network, AIM will not discover the ALIFs until they are reset to the zero config IP addresses.
- This could be caused by Layer 2 Cisco switches that have Spanning Tree Protocol (STP) enabled but do not also have portfast enabled on the ports to which ALIF units are connected. Without portfast enabled, ALIF units will all be assigned the same zero config IP address at reboot and AIM will only

acquire them one at a time on a random basis.

- You can easily tell whether portfast is enabled on a switch that is running STP: When you plug the link cable from a working ALIF unit into the switch port, check how long it takes for the port indicator to change from orange to green. If it takes roughly one second, portfast is on; if it takes roughly thirty seconds then portfast is disabled.

Remedies:

- Ensure that the ALIF units and the AIM box are located within the same subnet. AIM cannot cross subnet boundaries.
- Manually reset the ALIF units to their zero config IP addresses. Please refer to the ALIF user guide for details.
- Enable portfast on all switch ports that have ALIF units attached to them or try temporarily disabling STP on the switches while AIM is attempting to locate ALIF units.

Problem: The mouse pointer of the ALIF receiver is slow or sluggish when moved across the screen.

This issue is often related to either using dithering on the video output of one or more transmitting computers or using VGA-to-DVI video converters. Dithering is used to improve the perceived quality and colour depth of images by diffusing or altering the colour of pixels between video frames. This practice is commonly used on Apple Mac computers using ATI or Nvidia graphics cards. VGA to-DVI converters unwittingly produce a similar issue by creating high levels of pixel background noise. ALIF units attempt to considerably reduce network traffic by transmitting only the pixels that change between successive video frames. When dithering is enabled and/or VGA-to-DVI converters are used, this can have the effect of changing almost every pixel between each frame, thus forcing the ALIF transmitter to send the whole of every frame: resulting in greatly increased network traffic and what's perceived as sluggish performance.

Remedies:

- Linux PCs Check the video settings on the PC. If the Dither video box option is enabled, disable it.
- Apple Mac with Nvidia graphics. Use the Adder utility for Mac's – Contact technical support.
- Apple Mac with ATI graphics. Use the ALIF 2000 series unit with Magic Eye dither removal feature.
- Windows PCs. If you suspect these issues with PC's, contact technical support for assistance.

Glossary

Internet Group Management Protocol

The Internet Group Management Protocol (IGMP) is designed to prevent multicast flooding by allowing Layer 3 switches to check whether host computers within their care are interested in receiving particular multicast transmissions. They can then direct multicast data only to those points that require it and can shut off a multicast stream if the subnet has no recipients. There are currently three IGMP versions: 1, 2 and 3, with each version building upon the capabilities of the previous one:

IGMPv1 allows host computers to opt into a multicast transmission using a Join Group message, it is then incumbent on the router to discover when they no longer wish to receive; this is achieved by polling them (see IGMP Querier below) until they no longer respond.

IGMPv2 includes the means for hosts to opt out as well as in, using a Leave Group message.

IGMPv3 encompasses the abilities of versions 1 and 2 but also adds the ability for hosts to specify particular sources of multicast data. ADDERLink INFINITY units make use of IGMPv2 when performing multicasts to ensure that no unnecessary congestion is caused.

IGMP Snooping

The IGMP messages are effective but only operate at layer 3 - intended for routers to determine whether multicast data should enter a subnet. A relatively recent development has taken place within the switches that glue together all of the hosts within each subnet: IGMP Snooping. IGMP snooping means these layer 2 devices now have the ability to take a peek at the IGMP messages. As a result, the switches can then determine exactly which of their own hosts have requested to receive a multicast – and only pass on multicast data to those hosts.

IGMP Querier

When IGMP is used, each subnet requires one Layer 3 switch to act as a Querier. In this lead role, the switch periodically sends out IGMP Query messages and in response all hosts report which multicast streams they wish to receive. The Querier device and all snooping Layer 2 switches, then update their lists accordingly (the lists are also updated when Join Group and Leave Group (IGMPv2) messages are received).

IGMP Fast-Leave (aka Immediate Leave)

When a device/host no longer wishes to receive a multicast transmission, it can issue an IGMP Leave Group message as mentioned above. This causes the switch to issue an IGMP Group-Specific Query message on the port (that the Leave Group was received on) to check no other receivers exist on that connection that wish to remain a part of the multicast. This process has a cost in terms of switch processor activity and time.

Where ALIF units are connected directly to the switch (with no other devices on the same port) then enabling IGMP Fast-Leave mode means that switches can immediately remove receivers without going through a full checking procedure. Where multiple units are regularly joining and leaving multicasts, this can speed up performance considerably.

Jumbo frames (Jumbo packets)

Since its commercial introduction in 1980, the Ethernet standard has been successfully extended and adapted to keep pace with the ever-improving capabilities of computer systems. The achievable data rates, for instance, have risen in ten-fold leaps from the original 10Mbit/s to a current maximum of 100Gbit/s. While data speeds have increased massively, the standard defining the number of bytes (known as the Payload) placed into each data packet has remained resolutely stuck at its original level of 1500 bytes. This standard was set during the original speed era (10Mbits/s) and offered the best compromise at that speed between the time taken to process each packet and the time required to resend faulty packets due to transmission errors. But now networks are much faster and files/data streams are much larger; so time for a change? Unfortunately, a wholesale change to the packet size is not straightforward as it is a fundamental standard and changing it would mean a loss of backward compatibility with older systems. Larger payload options have been around for a while, however, they have often been vendor specific and at present, they remain outside the official standard. There is, however, increased consensus on an optional 'Jumbo' payload size of 9000 bytes and this is fully supported by the ADDERLink INFINITY (ALIF) units. Jumbo frames (or Jumbo packets) offer advantages for ALIF units when transmitting certain high-resolution video signals across a network. This is because the increased data in each packet reduces the number of packets that need to be transferred and dealt with - thus reducing latency times. The main problem is that for jumbo frames to be possible on a network, all of the devices on the network must support them.

Spanning Tree Protocol (STP)

In order to build a robust network, it is necessary to include certain levels of redundancy within the interconnections between switches. This will help to ensure that a failure of one link does not lead to a complete failure of the whole network. The danger of multiple links is that data packets, especially multicast packets, become involved in continual loops as neighbouring switches use the duplicated links to send and resend them to each other. To prevent such bridging loops from occurring, the Spanning Tree Protocol (STP), operating at layer 2, is used within each switch. STP encourages all switches to communicate and learn about each other.

It prevents bridging loops by blocking newly discovered links until it can discover the nature of the link: is it a new host or a new switch? The problem with this is that the discovery process can take up to 50 seconds before the block is lifted, causing problematic timeouts. The answer to this issue is to enable the portfast variable for all host links on a switch. This will cause any new connection to go immediately into forwarding mode.

However, take particular care not to enable portfast on any switch to switch connections as this can result in bridging loops.

Layer 2 and Layer 3: The OSI model

When discussing network switches, the terms Layer 2 and Layer 3 are very often used. These refer to parts of the Open System Interconnection (OSI) model, a standardised way to categorise the necessary functions of any standard network. There are seven layers in the OSI model and these define the steps needed to get the data created by you (imagine that you are Layer 8) reliably down onto the transmission medium (the cable, optical fibre, radio wave, etc.) that carries the data to another user; to complete the picture, consider the transmission medium is Layer 0. In general, think of the functions carried out by the layers at the top as being complex, becoming less complex as you go lower down. As your data travel down from you towards the transmission medium (the cable), they are successively encapsulated at each layer within a new wrapper (plus instructions), ready for transport. Once transmission has been made to the intended destination, the reverse occurs: Each wrapper is stripped away and the instructions examined until finally only the original data are left.

So why are Layer 2 and Layer 3 of particular importance when discussing ADDERLink INFINITY? Because the successful transmission of data relies upon fast and reliable passage through network switches – and most of these operate at either Layer 2 or Layer 3. The job of any network switch is to receive each incoming network packet (the collection of layered data), strip away only

the first few wrappers to discover the intended destination then rewrap the packet and send it in the correct direction. In simplified terms, the wrapper that is added at Layer 2 (by the sending system) includes the physical address of the intended recipient system, i.e. the unique MAC address (for example, 09:f8:33:d7:66:12) that is assigned to every networking device at manufacture. Deciphering recipients at this level is more straightforward than at Layer 3, where the address of the recipient is represented by a logical IP address (e.g. 192.168.0.10) and requires more knowledge of the surrounding network structure. Due to their more complex circuitry, Layer 3 switches are more expensive than Layer 2 switches of a similar build quality and are used more sparingly within installations.

Protocols and ports

In order to achieve the feat of sending high resolution DVI video, audio, RS232 serial and four USB streams across a standard Gigabit Ethernet network, ALIF units rely upon a combination of industry standard protocols. These protocols operate at Layer 4 (the Transport layer) in the OSI model, i.e. another level of sophistication above the Layer 2 and 3 techniques enacted by the switches and routers through which it travels.

TCP

TCP (Transmission Control Protocol) is a fundamental internetworking standard that allows a reliable data delivery route to be established between two hosts or devices. To ensure reliability of data transfer, TCP employs various techniques: Flow control to regulate data flow to suit the receiver, error detection to locate and replace corrupted packets, and congestion control to avoid swamping a busy network. By their nature, the techniques employed for TCP connections impose a certain latency to the connections. For this reason, TCP is used by ALIF to handle the slightly less time-sensitive, but highly accuracysensitive RS232 serial and USB data links.

UDP

Like TCP, UDP (User Datagram Protocol) is a fundamental standard that provides a data route between hosts and devices. However, unlike TCP, it does not carefully regulate the link. Once setup, data is sent to receiver(s) without knowledge of either their status or that of sent data. For this reason, UDP is totally unsuited for important control signals, however, it is ideal for data streams that can easily cope with the odd lost frame: video and audio. The great

advantage of UDP is that its latency is minimal and it also allows one sender to communicate with more than one receiver. As such, UDP is a key component of multicasting.

Ports used

ALIF units make use of various ports in order to achieve their various parallel data transfers:

Data	Protocol	Port
RS232 serial	TCP	3030
USB	TCP	3040 to 3050
Control	TCP	3000
Video	UDP	1237
Audio	UDP	3020

If ALIF communications are required to cross a firewall within the network, these ports must be opened to allow access.

ALIF transmitter video settings

Each ALIF transmitter includes controls to help you customise how video data is transmitted. When configured correctly for the application, these can help to increase data efficiency.

Background Refresh

The transmitter sends portions of the video image only when they change. In order to give the best user experience, the transmitter also sends the whole video image, at a lower frame rate, in the background. The Background Refresh parameter controls the rate at which this background image is sent. The default value is 'every 32 frames', meaning that a full frame is sent in the background every 32 frames. Reducing this to 'every 64 frames' or more will reduce the amount of bandwidth that the transmitter consumes. On a high-traffic network this parameter should be reduced in this way to improve overall system performance.

Colour Depth

This parameter determines the number of bits required to define the colour of every pixel. The maximum (and default) value is '24 bit'. By reducing the value

you can significantly reduce bandwidth consumption, at the cost of video colour reproduction.

Peak Bandwidth Limiter

The transmitter will employ a 'best effort' strategy in sending video and other data over the IP network. This means it will use as much of the available network bandwidth as necessary to achieve optimal data quality, although typically the transmitter will use considerably less than the maximum available. In order to prevent the transmitter from 'hogging' too much of the network capacity, you can reduce this setting to place a tighter limit on the maximum bandwidth permissible to the transmitter.

Frame Skipping

Frame Skipping involves 'missing out' video frames between those captured by the transmitter. For video sources that update only infrequently or for those that update very frequently but where high fidelity is not required, frame skipping is a good strategy for reducing the overall bandwidth consumed by the system.

Forwarding modes

In essence, the job of a layer 2 switch is to transfer as fast as possible, data packets arriving at one port out to another port as determined by the destination address. This is known as data forwarding and most switches offer a choice of methods to achieve this. Choosing the most appropriate forwarding method can often have a sizeable impact on the overall speed of switching:

- Store and forward is the original method and requires the switch to save each entire data packet to buffer memory, run an error check and then forward if no error is found (or otherwise discard it).
- Cut-through was developed to address the latency issues suffered by some store and forward switches. The switch begins interpreting each data packet as it arrives. Once the initial addressing information has been read, the switch immediately begins forwarding the data packet while the remainder is still arriving. Once all of the packet has been received, an error check is performed and, if necessary, the packet is tagged as being in error. This checking 'on-the-fly' means that cut-through switches cannot discard faulty packets themselves. However, on receipt of the marked packet, a host will carry out the discard process.
- Fragment-free is a hybrid of the above two methods. It waits until the first

64 bits have been received before beginning to forward each data packet. This way the switch is more likely to locate and discard faulty packets that are fragmented due to collisions with other data packets.

- Adaptive switches automatically choose between the above methods. Usually they start out as a cut-through switches and change to store and forward or fragment-free methods if large number of errors or collisions are detected.
- So which one to choose? The Cut-through method has the least latency so is usually the best to use with ADDERLink INFINITY units. However, if the network components and/or cabling generate a lot of errors, the Store and forward method should probably be used. On higher end store and forward switches, latency is rarely an issue.

Switch Setups

Cisco Catalyst 2960S and 3750

You need to configure the following settings:

- Assign switch IP address and default gateway
- Assign VLAN1 with an IP address
- Enable IGMP snooping
- Enable IGMP querier
- Enable Jumbo packets and choose the Max setting.
- Set portfast for STP

Assign switch IP address and default gateway

Beginning in privileged EXEC mode, follow these steps to manually assign IP information to multiple switched virtual interfaces (SVIs).

```
configure terminal
interface vlan 1
ip address
exit
ip default-gateway
end
show interfaces vlan 1
show ip redirects
copy running-config startup config
```

Assign VLAN1 with an IP address

```
configure terminal
vlan 1
end
copy running-config startup config
```

Enabling IGMP snooping

Beginning in privileged EXEC mode follow these steps to globally enable IGMP Snooping on the switch, first globally and then on VLAN1.

Globally:

```
configure terminal
ip igmp snooping
end
copy running-config startup-config
```

On VLAN1:

```
configure terminal
ip igmp snooping vlan 1
end
copy running-config startup-config
```

Enable IGMP querier

Enable IGMP Querier using the following commands in CLI mode

```
#configure terminal
#ip igmp snooping querier
#end
#show ip igmp snooping vlan 1
#copy running-config startup-config
```

Enable jumbo frames

Enable Jumbo frames and set the size to the maximum (9000 bytes).

```
configure terminal
system mtu jumbo 9000
end
copy running-config startup-config
reload
```

Enable STP portfast

```
configure terminal
interface
spanning-tree portfast
end
show spanning-tree interface portfast
copy running-config startup-config
```

Cisco 6500

For Cisco 6500 switches, portfast must be enabled for the initial device discovery process.

```
Switch#
Switch#spanning-tree portfast
```

Enable ip igmp snooping fast-leave

To enable IGMPv3-snooping fast-leave processing, use the ip igmp snooping fastleave command in interface configuration mode. To disable fast-leave processing, use the no form of this command:

```
ip igmp snooping fast-leave
no ip igmp snooping fast-leave
```

Syntax Description

This command has no arguments or keywords.

Defaults:

The defaults are as follows:

- IGMP version 2-Disabled
- IGMP version 3-Enabled

Command Modes

Interface configuration

Command History

Release Modification

- 12.2(17d)SXB - Support for this command on the Supervisor Engine 2 was extended to Release 12.2(17d)SXB.
- 12.2(33)SRA - This command was integrated into Cisco IOS Release 12.2(33) SRA.

Usage Guidelines

This command is not supported on Cisco 7600 series routers that are configured with a Supervisor Engine 720. Enter this command in VLAN interface configuration mode only.

Extreme X460-24t

You need to configure the following settings:

- Assign VLAN1 with an IP address
- Enable IGMP Snooping
- Enable IGMP Fast Leave
- Enable Jumbo packets and set the size to the Max

Assign VLAN1 with an IP address

- Connect a terminal or workstation running terminal emulation software to the console port, as detailed in "Using the Console Interface".
- At your terminal, press [Return] one or more times until you see the login prompt.
- At the login prompt, enter your user name and password. The user name is not case-sensitive. The password is case-sensitive. Ensure that you have entered a user name and password with administrator privileges. If you are logging in for the first time, use the default user name admin to log in with administrator privileges. For example: login: admin. Administrator capabilities enable you to access all switch functions. The default user names have no passwords assigned. If you have been assigned a user name and password with administrator privileges, enter them at the login prompt.
- At the password prompt, enter the password and press Return. When you have successfully logged in to the switch, the command line prompt displays the name of the switch.
- Assign an IP address and subnetwork mask for the default VLAN by using the following command: `configure {vlan} <vlan_name> ipaddress ipv6-linklocal` For example: `configure vlan default ipaddress 123.45.67.8 255.255.255.0`. The changes take effect immediately. NOTE: As a general rule, when configuring any IP addresses for the switch, you can express a subnet- mask by using dotted decimal notation or by using classless inter domain routing notation (CIDR). CIDR uses a forward slash plus the number of bits in the subnet mask. Using CIDR notation, the command identical to the previous. Example is: `configure vlan default ipaddress 123.45.67.8/24`

- Configure the default route for the switch using the following command:
configure iproute add default <gateway> {<metric>} {multicast | multicast-only | unicast | unicast-only} {vr <vrname>}. For example: configure iproute add default 123.45.67.17. Save your configuration changes so that they will be in effect after the next switch reboot. If you want to save your changes to the currently booted configuration, use the following command: save
ExtremeXOS allows you to select or create a configuration file name of your choice to save the configuration to. If you want to save your changes to an existing or new configuration file, use the following command: save configuration <new-config>
- When you are finished using the facility, log out of the switch by typing: logout or quit
- Enable IGMP Snooping

To enable or disable IGMP snooping, use the following commands:

```
enable igmp snooping {forwardmcrouter-only | {vlan} |  
withproxyvr }  
disable igmp snooping {forwardmcrouter-only | with-  
proxy | vlan }
```

Enable IGMP Fast Leave

Enabling and Disabling IGMP Snooping Fast Leave

To enable the fast leave:

```
enable igmp snooping {vlan} fastleave
```

To disable the fast leave:

```
disable igmp snooping {vlan} fast-leave
```

Enable Jumbo packets and set the size to the Max

To enable jumbo frame support, enable jumbo frames on the desired ports. To set the maximum jumbo frame size, use the following command:

```
configure jumbo-frame-size amesize>
```

The jumbo frame size range is 1523 to 9216. This value describes the maximum size of the frame in transit (on the wire), and includes 4 bytes of CRC plus another 4 bytes if 802.1Q tagging is being used.

Set the MTU size for the VLAN by using the following command:

```
configure ip-mtu vlan
```


HP ProCurve 2810 and 2910

Using Command Line Interface (CLI) will provide access to the switch's full set of commands. The settings below affect multicast operation on the switch.

IGMP

IGMP configuration is disabled by default on ProCurve 2810 and 2910 switches.

- To determine the current status of IGMP use the command:

```
show ip igmp config
```

- To enable or disable IGMP use the command:

```
[no] ip igmp
```

- To enable or disable IGMP Querier function use:

```
[no] ip igmp querier
```

Jumbo frames

Jumbo packet support is disabled by default on ProCurve 2810 and 2910 switches.

- To determine the current jumbo configuration setting use the command:

```
show vlans
```

- To enable or disable jumbo packet support issue the following command in Manager level operation :

```
[no] vlan 1 jumbo
```

This lists static VLANs configured on the switch. The Jumbo column indicates whether the given VLAN is configured for Jumbo packets or not.

Spanning-Tree Protocol (STP)

STP is disabled by default on ProCurve 2810 and 2910 switches.

- To determine the current STP configuration on the switch use the command:

```
show spanning-tree
```

H3C 5120

```
<alif>dis cur
version 5.20, Release 2202P06
sysname alif
domain default enable system
telnet server enable
undo ip ttl-expires
igmp-snooping
vlan 1
igmp-snooping enable
igmp-snooping drop-unknown
igmp-snooping querier
igmp-snooping general-query source-ip
172.29.0.1

radius scheme system
server-type extended
primary authentication 127.0.0.1 1645
primary accounting 127.0.0.1 1646
user-name-format without-domain
domain system
access-limit disable
state active
idle-cut disable
self-service-url disable
user-group system
interface NULL0
interface Vlan-interface1
ip address 172.29.0.1 255.255.255.0
interface GigabitEthernet1/0/1
interface GigabitEthernet1/0/2
interface GigabitEthernet1/0/3
interface GigabitEthernet1/0/4
```

HuaWei Quidway s5328c-EI (Layer 3)

You need to configure the following:

- Assign switch IP address
- Enable IGMP snooping
- Enable IGMP snooping send-query

- Enable IGMP querier
- Enable multicast drop-unknown

Assign switch IP address

```
1 <Quidway> system-view
2 [Quidway]
3 [Quidway] interface vlanif 1
4 [Quidway -Vlanif1] ip address { mask | mask-length }
5 [Quidway -Vlanif1] quit
```

Enable IGMP snooping

```
1 [Quidway] igmp-snooping enable
2 [Quidway] vlan 1
3 [Quidway] -Vlanif1 igmp-snooping enable
4 [Quidway] -Vlanif1 quit
```

Enable IGMP snooping send-query

```
1 [Quidway] igmp-snooping send-query enable
2 [Quidway] igmp-snooping send-query source-address
{vlanif 1 address }
```

Enable IGMP querier

```
1 [Quidway] vlan 1
2 [Quidway -Vlanif1] igmp-snooping querier enable
3 [Quidway] -Vlanif1] quit
```

Enable multicast drop-unknown

```
1 [Quidway vlan 1]
2 [Quidway] -Vlanif1] multicast dropunknown
3 [Quidway] -Vlanif1] quit
4 [Quidway] quit
5 <Quidway> save
```

About Adder

Adder design and manufacture an advanced range of KVM switches, extenders and IP solutions which enable the reliable control of local, remote and global IT systems. By empowering IT professionals to securely manage technology resources anywhere in the world, Adder solutions help customers make the best use of those resources while driving down total cost of ownership. In addition, through its advanced video and audio extension solutions, Adder is enabling the next generation of digital signage.

More information about Adder and its solutions is available at www.adder.com.

ADDER TECHNOLOGY

Head Office
Tel: +44 (0)333 2079766 Fax: +44 (0)1954 780081
email: sales@adder.com www.adder.com

ADDER AMSTERDAM

Benelux, Western and Southern Europe
Tel: +31 (0)321 383501 Fax: +44 (0)1954 780081
email: sales@adder.com www.adder.com

ADDER CORPORATION

USA and Canada
Tel: +1 888 932 3337 Fax: +1 888 275 1117
email: usasales@adder.com www.adder.com

ADDER BERLIN

Central and Eastern Europe, Russia, CIS
Tel: +49 (0)30 8849 67-50 Fax: +49(0)308849 6748
email: vertrieb@adder.com www.adder.com

ADDER ASIA

Asia Pacific
Tel: +65 6288 5767 Fax: +65 6284 1150
email: asiasales@adder.com www.adder.com

ADDER STOCKHOLM

All Nordic Countries
Tel: +46 (8) 574 210 95 Fax: +46 (8) 574 211 95
email: sales@adder.com www.adder.com