

White Paper

Optimizing Your Network for Agility™ KVM Switch



ACR1002A-T

Successful Agility Implementation

SINCE ITS INTRODUCTION IN 2010, the Black Box Agility system has provided new levels of flexibility in the distribution of high-performance computing power. The reliable transfer of high resolution DVI video, audio, RS-232 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.

Agility units are fully compliant with the open Gigabit Ethernet standards; however, in order to achieve success it is necessary to ensure that various key elements are in place. Such improvements to your network will benefit all high throughput devices used on it, not just Agility. It's about maximizing throughput, banishing bottlenecks, and eliminating weak links.

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

Contents

- Choosing layer 2 switches (page 2)
- Designing the network layout (page 3)
- Configuring switches and Agility (page 5)
- Troubleshooting (page 6)
- Glossary (page 8)

Black Box Agility

ACR1000A-T-R2	Agility DVI, USB, and Audio Extenders over IP, Transmitter
ACR1002A-T	Agility Dual DVI, USB, and Audio KVM Extender over IP, Dual-Head or Dual-Link, Transmitter
ACR1012A-T	Agility Dual-Head Transmitter with VNC Port
ACR1020A-T	Agility Dual-Head, Single-Link Transmitter
ACR1000A-R-R2	Agility DVI, USB, and Audio Extenders over IP, Receiver
ACR1002A-R	Agility Dual DVI, USB, and Audio KVM Extender over IP, Dual-Head or Dual-Link, Receiver
ACR1020A-R	Agility Dual-Head, Single-Link Receiver
ACR1000A-CTL	ServSwitch Agility Controller Unit
ACR1000A-CTL-8	Agility iPATH Controller Unit with License for (8) Endpoint Devices
ACR1000A-CTL-16	Agility iPATH Controller Unit with License for (16) Endpoint Devices
ACR1000A-CTL-24	Agility iPATH Controller Unit with License for (24) Endpoint Devices
ACR1000A-CTL-48	Agility iPATH Controller Unit with License for (48) Endpoint Devices
ACR1000A-CTL-96	Agility iPATH Controller Unit with License for (96) Endpoint Devices
ACR1000A-CTL-192	Agility iPATH Controller Unit with License for (192) Endpoint Devices
ACR1000A-CTL-288	Agility iPATH Controller Unit with License for (288) Endpoint Devices
ACR1000A-CTL-ULC	Agility iPATH Controller Unit with License for Unlimited Endpoint Devices

Choosing layer 2 switches

This section provides various general and specific recommendations for switches to use with Agility; however, there is no substitute for testing in real-world situations.

General recommendations

There are certain minimum features that you need to ensure:

- Gigabit (1000 Mbps) or faster Ethernet ports,
- Support for IGMP v2 (or v3) snooping,
- Support for Jumbo frames (packets) up to 9216-byte size,
- High-bandwidth connections between switches, preferably Fiber 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 Agility hardware rather than software routines on a general processor.
- Check the maximum number of concurrent “snoopable groups” each switch can handle and ensure that they meet or exceed the number of Agility 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 downstream data speeds for each port are 1 Gigabit per second.
- Wherever possible, use the same switch manufacturer throughout a single subnet and, if possible, the same model of switch—this will simplify configuration and lessen the chances of compatibility issues.
- 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.

The trouble with multicasting.

Where an Agility 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.

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 Layer 2 switches known to work with Agility

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

- Cisco 2960
- Cisco 3750
- Cisco 4500
- Cisco 6500
- Extreme Networks X480
- HP Procurve 2810
- HP Procurve 2910

For specific details on configuring each of these switches for efficient operation with Agility, refer to the Switch setups section.

Designing the network layout

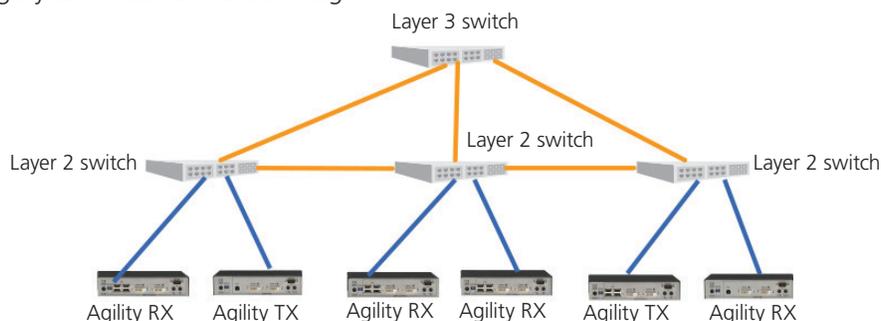
Selecting the right kind of switches is the first stage; ensuring they are laid out correctly is the next.

General recommendations

- The general rule is: Keep it flat. This means adopting a basic line cascade structure with a flat structure rather than a pyramid or tree arrangement.
- Keep the distances between the switches as short as possible.
- Ensure sufficient bandwidth between switches to eliminate bottlenecks.
- Where the controller box is used to administer multiple Agility transceivers, you must ensure that the Agility Controller box and all the Agility units under its control are located within the same subnet.
- Do not use VGA-to-DVI converters; 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 Agility transmitter would be forced to send all of every video frame, thus creating large amounts of unnecessary network traffic.
- Where more than ten Agility units will be used on a subnet, create a private network, i.e., no competing traffic from other network devices.

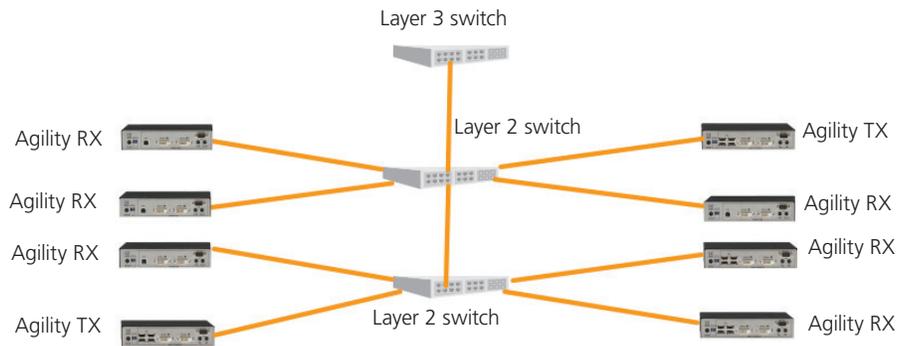
Layouts

The main problem with a pyramid or tree layout is that essential IGMP traffic issued by the Querier device (the multicast router) at the top of the pyramid will only travel down each leg of the pyramid. What this means is that an Agility unit located in one leg cannot communicate with an Agility unit situated in another leg.

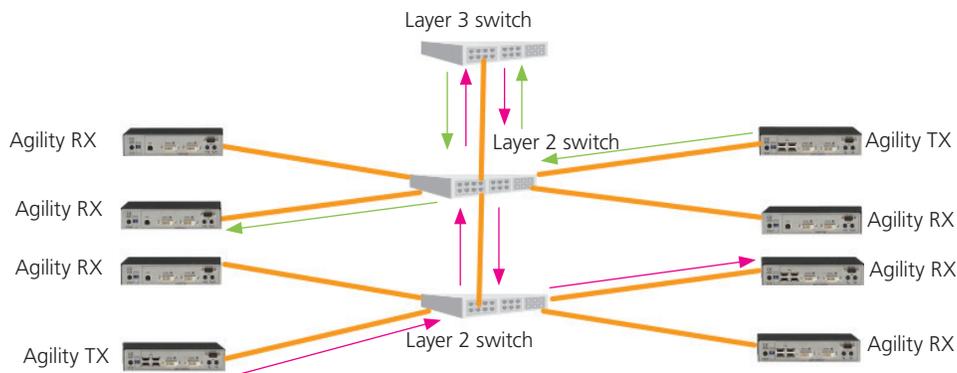


Optimizing Your Network for Agility KVM Switch

The advantage of arranging switches as a line cascade is that the communication links between them effectively form a simple bus structure that avoids the IGMP traps encountered by the pyramid arrangement.



A significant requirement in line cascades is the need for sizable bandwidth between the switches in order to provide redundancy and in so doing prevent data bottlenecks. When linking 48 port switches, generally a 20 GB link bandwidth is required between the multicast router and the first switch. A 10 GB link bandwidth is then required between the first and second switches. The reason for this becomes clearer when data paths are added to the diagram.



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 with the L3 switch.

The links between switches could either be multiple Gigabit Ethernet connections or, more favorably, fibre channel connections. Due to the interlink bandwidths required to prevent bottlenecks, we strongly recommend that there be no more than two cascade levels in any single line cascade arrangement.

Configuring switches and Agility

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 set-ups 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 Agility 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 Agility units, standard Ethernet frame sizes are used to transfer video data quickly and efficiently. However, certain video resolutions cause issues that require Agility 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 Agility 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 iPATH controller, whereby Agility 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 Agility 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 Agility transmitter video settings, if necessary

Within each Agility 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 iPATH or directly via the Agility browser-based configuration utility. Refer to the relevant user guides for details.

- If color quality is important, then leave Color 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 Color 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 > Agility transmitter video settings.

Optimizing Your Network for Agility KVM Switch

Ensure that all Agility units are fully updated

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

Troubleshooting

If you encounter problems during installation or operation, refer to the Troubleshooting section, or contact Black Box Technical Support at 724-746-5500 or info@blackbox.com.

Problem: The video image of the Agility 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 Agility 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 Agility units may be outputting Jumbo frames due to the video resolution (2048 horizontal pixels) being used. If jumbo frames are output by an Agility 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 Agility 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 Agility 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 Agility 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 Agility unit is version 2.1 or greater.
- Try adjusting the transmitter settings on each Agility to make the output data stream as efficient as possible. See Alter Agility transmitter video settings if necessary for details.

Problem: The video image of the Agility receiver is slow to respond and the mouse pointer jumps 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 color depth of images by diffusing or altering the color of pixels between video frames. This practice is commonly used on Apple Mac systems that are equipped with Nvidia graphics cards. VGA-to-DVI converters unwittingly produce a similar issue by creating high levels of pixel background noise.

Agility 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 Agility transmitter to send the whole of every frame: resulting in greatly increased network traffic.

Remedies:

- Within the video settings of the sending computer(s), disable the “Dithering” option.
- Replace old VGA adapters on host computers with DVI video cards.

Problem: The audio output of the Agility 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 left.

Problem: iPATH cannot locate working Agility units.

There are a few possible causes:

- The Agility units must be reset back to their zero config IP addresses for iPATH discovery. If you have a working network of Agility units without iPATH and then add iPATH to the network, iPATH will not discover the Agility units 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 Agility units are connected. Without portfast enabled, Agility units will all be assigned the same zero config IP address at reboot and iPATH 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 Agility 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 Agility units and the iPATH box are located within the same subnet. iPATH cannot cross subnet boundaries.
- Manually reset the Agility units to their zero config IP addresses. Please refer to the Agility user guide for details.
- Enable portfast on all switch ports that have Agility units attached to them or try temporarily disabling STP on the switches while iPATH is attempting to locate Agility units.

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.

Agility 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 2—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)

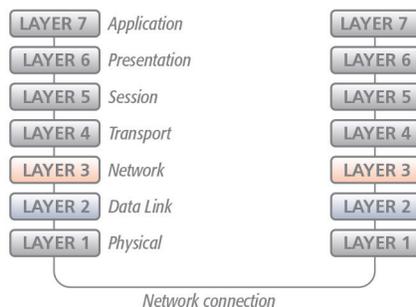
When a device/host no longer wishes to receive a multicast transmission, it can issue an IGMP Leave Group message as mentioned above. This causes the switch to issue an IGMP Group—Specific Query message on the port (that the Leave Group was received on) to check that no other receivers exist on that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

Where Agility 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.

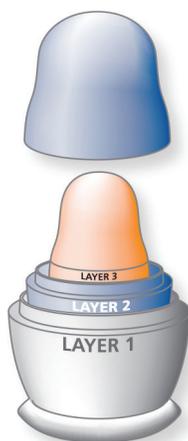
Layer 2 and Layer 3: The OSI model

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 standardized way to categorize 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. Although this analogy only goes so far, imagine each wrapper (with its payload of data) as a Russian Doll being placed into a slightly larger doll (along with a few instructions), ready for transport. Once transmission has been made to the intended destination, the reverse occurs: Each wrapper is stripped away (the Russian Dolls are opened) and the instructions examined until finally only the original data are left.



So why are Layer 2 and Layer 3 important when discussing Agility? 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 Russian Dolls), 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 when manufactured. 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. Because of 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.

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 10 Mbps to a current maximum of 100 Gbps.

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 (10 Mbps) 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.

Optimizing Your Network for Agility KVM Switch

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 Agility units.

Jumbo frames (or Jumbo packets) offer advantages for Agility 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)

To build a robust network, the design must 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 neighboring 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.

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 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.

Protocols and ports

To send high-resolution DVI video, audio, RS-232 serial, and four USB streams across a standard Gigabit Ethernet network, Agility units rely upon a combination of industry standard protocols. These protocols operate at Layer 4 (the Transport layer) in the [OSI model](#), which is 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 Agility to handle the slightly less time-sensitive, but highly accuracy-sensitive RS-232 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

Agility units use various ports to achieve their various parallel data transfers.

Ports used for parallel data transfers.

Data	Protocol	Ports
RS-232 serial	TCP	3030
USB	TCP	3040 to 3050
Control	TCP	3000
Video	UDP	1237
Audio	UDP	3020

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

Agility transmitter video settings

Each Agility transmitter includes controls to help you customize 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.

Color Depth

This parameter determines the number of bits required to define the color 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 color reproduction.

Optimizing Your Network for Agility KVM Switch

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.

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.

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).

1. configure terminal
2. interface vlan 1
3. ip address <ip-address subnet-mask>
4. exit
5. ip default-gateway <ip-address>
6. end
7. show interfaces vlan 1
8. show ip redirects
9. copy running-config startup config

Assign VLAN1 with an IP address

1. configure terminal
2. vlan 1
3. end
4. 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:

1. configure terminal
2. ip igmp snooping
3. end
4. copy running-config startup-config

On VLAN1:

1. configure terminal
2. ip igmp snooping vlan 1
3. end
4. copy running-config startup-config

Enable IGMP querier

Enable IGMP Querier using the following commands in CLI mode

1. configure terminal
2. ip igmp snooping querier
3. end
4. show ip igmp snooping vlan 1
5. copy running-config startup-config

Enable jumbo frames

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

1. configure terminal
2. system mtu jumbo 9000
4. end
5. copy running-config startup-config
6. reload

Enable STP portfast

1. configure terminal
2. interface <interface-id>
3. spanning-tree portfast
4. end
5. show spanning-tree interface <interface-id> portfast
6. copy running-config startup-config

Optimizing Your Network for Agility KVM Switch

Cisco 6500

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

1. Switch<config-if>#
2. Switch<config-if>#spanning-tree portfast

ip igmp snooping fast-leave

To enable IGMPv3-snooping fast-leave processing, use the ip igmp snooping fast-leave 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

1. Connect a terminal or workstation running terminal emulation software to the console port, as detailed in "Using the Console Interface".
2. At your terminal, press [Return] one or more times until you see the login prompt.
3. 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.

4. 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.

5. Assign an IP address and subnetwork mask for the default VLAN by using the following command:

```
configure {vlan} <vlan_name> ipaddress [<ipaddress> {<ipNetmask>} | ipv6-linklocal {eui64} <ipv6_address_mask>]
```

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
```

6. 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 123.45.67.1
```

7. 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
```

This 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 [<existing-config> | <new-config>]
```

8. 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 {forward-mcrouter-only | {vlan} <name> | with-proxy vr <vrname>}
```

```
disable igmp snooping {forward-mcrouter-only | with-proxy | vlan <name>}
```

Enable IGMP Fast Leave

Enabling and Disabling IGMP Snooping Fast Leave

To enable the fast leave:

```
enable igmp snooping {vlan} <name> fast-leave
```

To disable the fast leave:

```
disable igmp snooping {vlan} <name> fast-leave
```

Optimizing Your Network for Agility KVM Switch

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 <framesize>
```

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 <mtu> vlan <vlan_name>
```

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
```


Optimizing Your Network for Agility KVM Switch

Disclaimer:

Black Box Network Services shall not be liable for damages of any kind, including, but not limited to, punitive, consequential or cost of cover damages, resulting from any errors in the product information or specifications set forth in this document and Black Box Network Services may revise this document at any time without notice.

About Black Box:

Black Box is a leading technology product solutions provider that helps customers build, manage, optimize, and secure their networks. The Black Box quality management system is ISO 9001:2008 certified, and the company has received numerous industry recognitions. Black Box provides its customers with free, 24/7 pre- and post-sales technical support. The Black Box catalog and Web site offer an extensive range of infrastructure products including Cabling, Cabinets & Racks, Data Center Cooling Solutions, Power & Surge Protection, and Environmental Monitoring.

© Copyright 2015. Black Box Corporation. All rights reserved. Black Box® and the Double Diamond logo are registered trademarks of BB Technologies, Inc. Any third-party trademarks appearing in this publication are acknowledged to be the property of their respective owners.