AZURE DEVELOPMENT SOLUTIONS

NetworkTV

Network Considerations for IPTV

Head Office: The Livery, Millhaugh, Dunning PH2 0DW
Phone: +44 (0) 1865 522774
E-Mail: info@azure-ds.com Web: www.azure-ds.com
Registered Office: 1 Vincent Square, London, SW1P 2PN
Release 2.3

TABLE OF CONTENTS

Introduction

Internet Protocol TV (IPTV) delivers digital video and audio signals over a network to provide multimedia content. Many applications for IPTV are for distributing video content within an organization to displays in common areas or to individual users on their computer. That content can either be live TV or pre-recorded video content stored on a server.

When distributing IPTV content, the content is distributed over the same network that employees use to send emails, files, and accessing other corporate information. For this reason, every organization considering an IPTV solution must take into account the impact to their data network to make sure they have enough bandwidth for both their current user workload and their planned IPTV uses. Otherwise, there are risks that an organization will overwhelm their network when users start accessing the IPTV content. It is important that the IT group within an organization is included in the early discussions of an IPTV deployment so they can plan for its network requirements.

“IPTV” is defined as digital audio and video signals streamed over a network. It is important to understand the details behind this definition:

“digital”- Any audio or video signal transmitted over IP must be digital, meaning it must first be converted from electrical signals to computer data. This process is called encoding. Also, if the same content will be played back on a display (television) that requires either analogue signals or digital signals in a format different than the one you are streaming in, the signal may need to be decoded to play it back.

“audio”- Digital audio comes in many formats, including MP3, Windows Media Audio (.wma), etc.

“video”- Digital video comes in a variety of formats, including MPEG-2, MPEG-4, Flash video (.flv), and Windows Media Video (.wmv).

“streamed”- When a media file is “streamed” rather than downloaded, the server storing the media file transmits the file over the network a portion at a time while it is played. The complete file is never actually saved on the local computer.

“network”- The network is the collection of hardware and software that connects a server that stores multimedia to the workstation that plays it. Networks are often large structures that do more than merely handle data streaming. Because of this, organizations must always take into account the impact of video streaming on the rest of the larger network.

An Overview of Active and Passive Components used to create an IP Network

A Wired Computer Network (LAN) is basically a combination of various Active and Passive Network Components. In this article, we explore the salient points on the important Active and Passive Components that are required for building a basic wired computer network.

Wired Computer Network – Architecture Diagram:

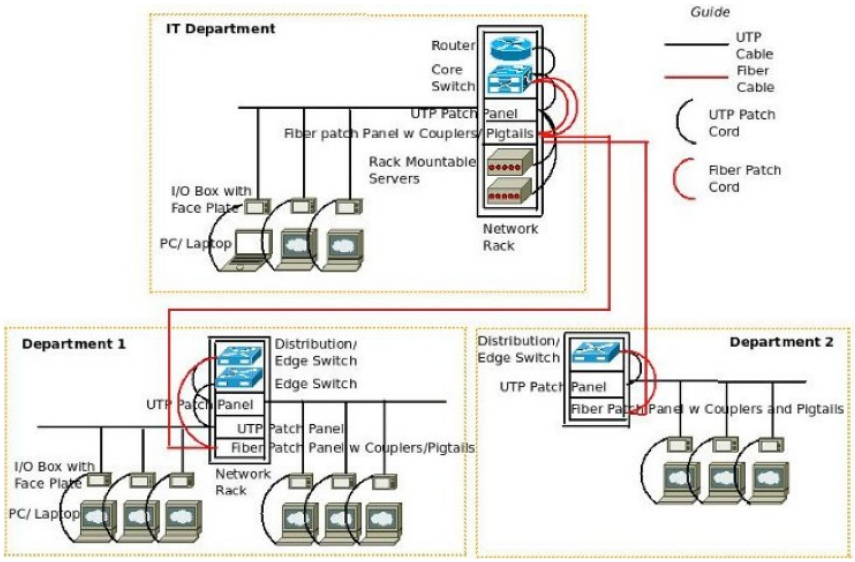

Architecture Diagram – Active and Passive Components in an IP Network (excitingip.net)

In the above diagram, let us assume that there are basically three departments in an organization that wants to have a LAN across all the departments – IT Department, Department 1, Department 2. So, if we are to plan for the network components department wise, for the IT department, we could plan for:

- Network rack,
- Router, Core switch
- Edge Switches (if required)
- UTP Patch panel, UTP Patch Cords
- Fiber Patch Panel, Fiber Patch Cords
- Cat 6/ Cat6A UTP cables
- I/O Box with Face Plate, UTP Patch Cords
- PVC Channel – Casing Caping/ Conduits
- Fiber Cables (Single Mode or Multi Mode)

The components required in the other two departments would also be similar with the exception of router/ core switch being replaced by distribution/ edge switches.

The above mentioned network components can be broadly divided in to two categories – Active Components and Passive Components. Active Components are those devices which required to be supplied with external power (AC/DC/POE etc) in order to function. They also boost the power of the signals. Passive components do not require to be provided with any electrical power to work – they just plug on to active components and transmit/carry the information (electrical/optical signals).

Active Network Components:

Network Switches:

Network Switches are the basic components of an IP Network. All the network endpoints (like PC's, Laptops, Printers, etc) connect to these switches. As the name goes, they switch (distribute) the data received from one node to any other node in the network. The network switches come in a variety of configurations, and the popular ones are mentioned below:

8 Port – 10/100/1000 Mbps
16 Port – 10/100/1000 Mbps
24 Port – 10/100/1000 Mbps
48 Port – 10/100/1000 Mbps

Network switches could also have 10/100 Mbps and POE/Non-POE Port combinations. They could also have variations in terms of functionalities – Manageable, Semi-Manageable and Unmanaged Switches. There are even 24/ 48 Port Optical Switches which connect as many optical connections in addition to the 2/4 ports of the optical connections that normal switches have.

The numbers (8, 16 etc.) in the above list refers to the number of Copper UTP Connectors the switch has, and the switches can connect to as many network devices. Each such port supports a maximum speed of 10(or) 100 / 10(or)100(or)1000 Mbps depending on the end-point connecting to it (it can auto negotiate to the highest speed supported by the endpoint). Some ports support POE (Power over Ethernet) which is a technology to carry the power as well as data to the endpoints, so that the endpoints need not connect to a separate power source (In the case of Wireless Access Points, IP Phones etc).

Some network switches are of Un-managed type – You can just connect the computers to them, connect them to neighboring switches and extend the network, but beyond that function, not much functionalities/ management features are supported by them. The advantages of unmanaged switches are their cost – they are inexpensive.

Some network switches are of Semi-Managed type – They come with a web browser-based management interface, limited QoS configurations, VLAN configuration, 802.1x support and other such limited management features required for the management of the critical functionalities of the network. But these management features are limited to what is determined by the manufacturer. These switches are slightly more expensive than the unmanaged variety but less expensive than fully manageable switches.

Some network switches are Fully Manageable – They allow the configurations of VLAN's per port, allow VLAN trunking, support web-based management functionalities, support SNMP/RMON protocols so that each port can be monitored by an SNMP based network management system, support RSTP (Rapid Spanning Tree Protocol) so that alternates cabling paths can be created for uplinking, support Link Aggregation so that couple of cables from individual ports can connect to the uplink switch with double the speed, support port mirroring for management/ call recording, support stacking and many other such useful features which help in maintaining a network.

Layer 3 Switches:

As the network becomes bigger and bigger, it becomes difficult to manage all the nodes using a single layer 2 network segment. One of the main problem with such unsegmented networks are broadcasts which can create performance bottlenecks on large networks. Another issue is the spreading of virus and botnets – with a segmented network, these remain mostly within their segments. That's why VLAN's are advocated on large networks which segment the network based on the location/ department/ application etc.

But the devices in one segment of the network would need to invariably communicate with the other segments – Especially in centralized networks where all the servers are designed to be in a common VLAN and the nodes communicating with them are from different VLAN's. In such cases, there needs to be a Layer 3 network device that performs seamless Inter-VLAN routing without affecting the performance of the network – This is exactly the reason why Layer-3 switches are required. They are capable of performing both the Layer 2 Switching and Layer 3 Routing at Line Rate. They also allow to configure flexible network wide security policies and perform Layer 3 QoS functionalities which are critical in converged networks which carry a substantial amount of real-time traffic that require low latency.

Passive Network Components:

Structured Cabling has become quite common for inter-connecting the various active devices in an IP network. So the following passive components are commonly utilized in an IP Network for Structured Cabling:

- Cat 6 UTP (Un-shielded Twisted Pair) Copper Cables – These are the network cables that connect a PC/ endpoint to a network switch. Some times, they are also used to provide inter-connectivity between switches as long as the distance is not greater than 90 meters, which is the distance they support for transmitting data without using any repeater (repeater function is provided by using network switches).

- Cat 6 UTP Patch Cords – These are one meter/ 2 meter factory crimped cables with RJ-45 connectors installed at both ends. Actually, the Cat 6 Cables are not recommended to be directly terminated in either the network switch or the PC/endpoint. Only the patch cords terminate on both devices and connect to the Cat 6 Network cable through an I/O Box and UTP patch panel.

- Network Rack – Network Racks are either wall mounted or Floor Standing types depending upon their size. Common sizes of network racks range from 6U to 42U. All the network equipments are designed in multiples of 1U so as to be accommodated in to these racks with standard fittings. They generally have a width of 19". The network racks come with a glass door, lock and key, fans required for cooling, trays, power supplies, cable managers and all other accessories.

- I/O Box and Face Plate: The I/O Box and Face Plate are kept near the computers and a UTP patch cord is used to connect the Face Plate with the network port in the PC. The Cat 6 UTP cable which comes from the switch terminates in to a permanent connection behind the I/O Box.

- UTP Patch Panel: The UTP Patch Panel is used for terminating all the Cat 6 Cables that come from various PC's/endpoints in the network (Actually I/O Box) to the rack. These Cables are permanently connected behind the UTP Patch Panel and UTP Patch Cords connect from the respective ports in front to the network switches. This allows for flexible moves, adds and changes without disturbing the switch ports. All the ports in the patch panel are labelled for easy identification of which node they are connected to.

- Optical Fiber Cables: For carrying data over 90 meters, Optical Fiber Cables are used. These cables use light for transmission of data instead of the electrical signals used by the UTP cables. They can carry data for longer distances – even to a few kilo meters without having to repeat the signals in between. There are two types of cables – Single Mode (Used for higher bandwidth requirements for longer distances) and Multi Mode (Used for shorter distances). They connect directly to the Fiber Patch Panel at either end. Usually they come in multiples of 6 Cores – 6 Core, 12 Core, 24 Core being common. For each connection, two cores are used – one for transmit and another for receive.

- Fiber Patch Panel/Patch Cords: The Optical Fiber Cables are terminated on either end using the Fiber Patch Panel, Pigtails and Coupler assembly. Actually each core of the Fiber Cable is spliced to fit in to the Fiber Patch Panel. A Fiber Patch Cord connects to the Patch Panel and the Fiber interface of the Network Switch. The Fiber interface is usually an SFP Port over which a Fiber Module is inserted (Mini-Gbic interface). This Fiber Module can connect to the fiber patch cord directly.

Different File Types Use Various Bandwidths

In the chart below, there are several different encoding formats and the bandwidth usages for standard and high-definition video streams, all measured in megabits per second, or Mbps. Figures are included for Windows Media, MPEG-2, and H.264, which is a type of MPEG-4 encoding.

Type of Encoding	SDTV	HDTV
Windows Media	1.0 -1.5 Mbps	5-12 Mbps
MPEG-2	2-6 Mbps	18-20 Mbps
MPEG-4 (H.264)	1-2 Mbps	5-8 Mbps

Encoding is the format that the analogue signal has been saved into when it was converted from an analogue to a digital signal. The type of encoding has a major impact on the bandwidth usage. Resolution, whether it is standard or high-definition, also has a large impact. A Standard Definition Windows Media file may stream at only 1 Mbps, but a High-Def video in MPEG-2 could stream at 20 Mbps.

These figures refer to the amount of data that is being sent over the network at any one time. IPTV uses a very high level of bandwidth when compared to voice or email. One hour of IPTV uses a minimum of about 5 Gigabytes of bandwidth and a single video stream can use up between 1 and 20 Mbps. So companies planning to use IPTV must determine how many simultaneous streams of video, and in what format, they will be using in order to insure their network bandwidth can support it.

There are three main ways to stream media over a network: Unicast, Multicast, and Broadcast.

Unicast creates a one-to-one connection between an individual user and a server. Each additional connection between a user and the server is a separate unicast connection, and takes up its own amount of bandwidth. Multicast creates a one-to-many connection between several users and a server. Only the users wanting to receive the signal will receive it.

Broadcast creates a one-to-all connection between all users and a server. Because it transmits data to all the workstations on a network, broadcast is bandwidth-intensive.

Multicast is generally preferred to broadcast, because it only transmits data to users that are actually accessing the information at a given time. For purposes of this paper, we will not explore broadcast.

When a user requests a media stream that employs unicast, there is an individual connection between the server and the user. For unicast, the server doesn't take into account the other streams that are taking place at the same. In a unicast system, the demand on the server increases as the number of users increases. However, each unicast stream is independent allowing each user to independently control their content.

In systems that use unicast, it is important to limit the number of users accessing content at any one time, because the impact of multiple unicast streams can quickly increase. Take our chart from earlier. If you have a high-definition MPEG-2 stream at 20 megabits per second, and have 100 users accessing that same content, the total amount of bandwidth used is 2000 megabits per second, or 2 gigabits. A large gigabit network would fail quickly.

One way to significantly reduce the impact media streaming has on a network is by using multicast, which groups all the different streams into a single stream.

With multicast, the server, rather than sending the media content to the IP address of a single workstation, sends the signal to an IP address that is dedicated to that particular multicast stream. Multicast streams have an IP address between 224.0.0.0 - 239.255.255.255.

Users subscribe to the stream by tying into the IP address for the stream. Each user accesses the same stream, and because the server is sending to only one IP address, the stream stops and starts at the same point for all users, much like a television program which starts on the hour.

There are benefits of using multicast and unicast, depending on the application.

Streaming Method	Benefit	Disadvantage	When to Use
Unicast	Allows for individual control of media stop and start	Bandwidth increases as usage increases	On-Demand
Multicast	One signal no matter how many users connect to the stream	Does not allow for individual control of media stop and start	Streaming

Methods to Manage Multicasting

Multicast Overview:

Multicast services are used to distribute streaming media, such as audio and video traffic, over the network. These video applications typically generate large amounts of traffic taking up network bandwidth. Both multicast and broadcast allow a network device to send single packets to multiple destinations. The difference is broadcast is designed to forward packets to all nodes on the VLAN, whereas multicast is designed to only forward to nodes in a multicast group. Multicast is therefore a better solution when trying to conserve network bandwidth. Also, since all nodes must process every frame they receive, multicast saves processing cycles on nodes that do not need to receive the multicast frames.

The multicast model consists of groups, transmitters, and receivers. A typical multicast group contains one transmitter (or sender) and one or more receivers. The receiver sends a “group join” message to network switches and routers via the Internet Group Multicast Protocol (IGMP). The switches and routers then forward multicast traffic only to the receivers that have joined the multicast group.

Generally, switches by default will forward multicast frames to all interfaces in a VLAN. No additional configuration is required to pass this traffic; however, the default behavior provides no benefit over using broadcast. To limit unnecessary traffic on the network, an IGMP Querier (layer 2) or PIM router (layer 3) must be seen on the network for multicast to stop broadcasting to all ports. IGMP snooping is enabled on all switches;

however, this does not keep the multicast traffic from broadcasting until the switches can see a Querier or PIM router.

Note: By creating an IGMP Querier or PIM router on the network, switches are able to stop multicast frames from being broadcasted and can forward those frames to transmitters and receivers in a specified multicast group.

Switches can use IGMP snooping to automatically collect information about which interfaces are participating in multicast groups. The switches then use this information to direct multicast traffic away from devices that are not interested in the IP multicast traffic.

Switches that are not “multicast aware” can use static multicast groups to manually select the interfaces that will pass multicast traffic.

IGMP snooping and static multicast groups accomplish the same task. One is dynamic; requiring less setup and maintenance, while the other is static; requiring more setup and maintenance time.

IGMP & PIM

Users that want to connect to a Multicast stream use a special protocol called Internet Group Management Protocol (IGMP) to join the multicast group, which is the list of all the users listening to the stream. While the users join the multicast group, a tree of the group is formed using a protocol called Protocol Independent Multicast (PIM). Using the tree generated by the PIM protocol, the routers send the signal to all users in the multicast group. The signal is sent only once over each branch of the tree, thus limiting the impact the stream has on network bandwidth.

IGMP is the protocol that manages multicast group membership. However, there are a few questions that need to be asked to an organization’s IT group if a customer is planning to use it.

First, not all networks support IGMP. Make sure your network supports IGMP and if IGMP is enabled on your network. Also, if you plan on distributing any IPTV content over the internet, be aware that the internet does not support multicast.

Second, make sure your switches support IGMP Snooping. IGMP Snooping is a feature in certain switches where the switch listens in on the conversations between hosts and routers having an IGMP conversation. Because the switch knows whether or not the workstations connected to it are part of the multicast group, it can prune the multicast traffic at the switch level, thus decreasing traffic for the workstations connected to the switch. This will also prevent the switch from being swamped by data traffic.

IGMP & PIM cont.

Sometimes on older switches such as "very old" HP procurve switches, multicast traffic is flooded. The best explanations are usually found in their user manuals such as the HP multicast routing guide.

On switches that do not support Data-Driven IGMP, unregistered multicast groups are flooded to the VLAN rather than pruned. In this scenario, Fast-Leave IGMP can actually increase the problem of multicast flooding by removing the IGMP group filter before the Querier has recognized the IGMP leave. The Querier will continue to transmit the multicast group during this short time, and because the group is no longer registered the switch will then flood the multicast group to all ports.

On ProCurve switches that do support Data-Driven IGMP (“Smart” IGMP), when unregistered multicasts are received the switch automatically filters (drops) them. Thus, the sooner the IGMP Leave is processed, the sooner this multicast traffic stops flowing.

Below are HP switches WITH problems (i.e. NOT supporting data driven igmp):
Switch 2600
Switch 2600-PWR
Switch 4100gl
Switch 6108

So if you have one of the above switches this is "normal". The workaround is to make multicasts join the multicast group. For this put multicast_auto_join=1 in your configuration file.

Other Network Technologies to Manage Bandwidth Impacts

A Virtual Local Area Network (VLAN) is a software-based grouping of workstations that are on different Local Area Networks (LANs). A VLAN can also be configured to be a subgroup of workstations on a single LAN. The software groups the workstations and treats them as a separate LAN. Packet routing for the VLAN is then separate from the LAN or LANs to which the workstations are physically connected.

There are several features of Virtual LANs that make them beneficial for media distribution. Because they are software-based, VLANS are not dependent on physical proximity like traditional LANs. Thus, a workstation on a physical LAN in Japan, a workstation on a physical LAN in England, and two workstations on a physical LAN in the USA can all be part of the same VLAN.

VLANs are also “plug and play,” meaning there is no need to pull wires to physically connect the workstations into a separate LAN. The workstation can be connected to whatever LAN is nearby, and then have its location updated in the software.

The final and most important benefit of VLANs is the workstations in a VLAN cannot see the traffic of the other nodes on the other VLANs without Inter-VLAN routing. This prevents bandwidth-intensive applications on one VLAN from impacting the performance of applications running on another VLAN. The other good piece about VLANs is that they introduce virtual barriers between segments of the network where firms can implement security controls.

These features make VLANs very beneficial when streaming multimedia. For example, a group of reception area displays playing a schedule of videos can be grouped in a dedicated VLAN, which will lower the impact on the rest of the network. The grouped recipients can be located anywhere on the network and still be part of the same VLAN, and their positions can even change without additional wires needing to be pulled. This results in systems that are easier to install and maintain, and that have less adverse effect on existing networks.

TCP and UDP are two protocols that control how data packets are sent over a network.

TCP guarantees reliable data transmission by providing additional checks that all packets are received and received in order. TCP automatically slows down the data transmission when it appears that data is being sent too fast to ensure all the data is transmitted.

UDP offers higher-speed data transmission but at a loss of quality. Ensures packets received contain the same data as the packets that were sent, but does not track packet order or if all the packets were sent.

TCP has high reliability and high bandwidth usage. UDP has high speed and low bandwidth usage.

It is recommended that UDP be used when streaming multimedia. When streaming audio and video, a single packet contains only a small amount of data. For example, a packet may contain data about a small portion of an image for a few milliseconds of video. Such a loss is usually not noticeable to the end user.

TCP, because of its error checking, slows the transmission speed when a packet is lost. This causes lag in the signal, especially for large streams, such as HD signals. UDP is preferred because when streaming, it is more important that the bulk of the packets are received in a timely manner than maintaining 100% signal integrity.

It is important that organizations deploying IPTV have a managed network that has Quality-of-Service (QoS) for the video packets. Because most firms will have other data being transmitted over the same network, it is important to have QoS to minimize disruption to the video packets. QoS solutions can prioritize video packet data and separate it from other traffic to minimize interruption of the video picture.

Mobile and WiFi Connectivity

With regard to the networking configuration and WiFi setup for organisations who are considering providing the capability for mobile users to access video content streams, there are a number of factors that need to be taken into consideration. Note though, traditionally, WiFi networks are not multicast enabled and all video traffic is delivered as unicasts (some network manufacturers are now starting to employ multicasting over WiFi) so therefore we should base our calculations and technology on providing all streams as unicasts:

1. The number of mobile devices (eg. laptops, iPads etc) that will need to pick up unicast streams: this will dictate the number of wireless access points required within the room, floor or building and how the signal may be affected by other wireless networks in the vicinity. Generally it should be remembered that by its nature, a single access point cannot be expected to provide sustained and reliable service for large amounts of traffic. Typically, an access point can reasonably support around 25 or so users with a 'bursty' traffic profile (e.g. opening typical web pages, reading email etc). Performance of wireless in a particular area will decrease dramatically as the number of users increases, or if those users maintain sustained flows of traffic for a period of time (for example streaming audio or video content).
2. Reliability: However, the actual range of communication can differ significantly depending on a number of variables such as placement, height above ground, nearby obstructions, other electronic devices that might actively interfere with the signal by broadcasting on the same frequency, and the power output of devices. So the WAP will generally support about 25 users who are located within a radius of about 100m. Wireless devices can "listen" for data traffic on other frequencies, and can rapidly switch from one frequency to another to achieve better reception. However, the limited number of frequencies becomes problematic in areas where many wireless networks are operating over multiple WAPs. In such an

environment, signal overlap becomes an issue causing interference, which results in signal droppage and data errors.

3. The encoding bit rate: For live streaming there are a couple of factors to take into consideration. The bit rate will be dictated by the desired quality of video output versus the bandwidth available (see above). It might be useful for solutions to accommodate a dual output encoder, producing a low bit rate stream and also a high bit rate stream of the same content. This could also be recorded in the better quality (or both variants) and stored in a library archive thereby offering the capability to authenticated users to review it at a later date.
4. WAP capability: Wireless networking lags wired networking in terms of increasing bandwidth and throughput. Whilst typical wireless devices for the consumer market can reach speeds of 300 Mbit/s (megabits per second) (IEEE 802.11n) or 54 Mbit/s (IEEE 802.11g), wired hardware of similar cost reaches 1000 Mbit/s (Gigabit Ethernet). One impediment to increasing the speed of wireless communications comes from Wi-Fi's use of a shared communications medium, so a WAP is only able to use somewhat less than half the actual over-the-air rate for data throughput. Thus a typical 54 MBit/s wireless connection actually carries TCP/IP data at 20 to 25 Mbit/s.
5. Authentication: For authentication purposes, wireless network traffic passes through one of a number of gateway devices, and these may cause congestion problems or even failure under error conditions or high or malicious traffic.

So in summary, the crucial areas to consider are choosing the right wireless network equipment, making sure it is broadcasting on a specific frequency/channel, encoding at the bit rate that gives the accepted quality against bandwidth availability but making it simple for the authenticated users to pick up the streams.

A simple example of bandwidth usage of delivering streams to iPads is if we assume there are 20 users, we are creating 20 x unicast streams at a bit rate of 500Kbps, this requires constant bandwidth availability of 10Mbits.

Conclusions

When planning for the deployment of an IPTV solution, it is critical that an organization include their IT department in the planning stages for deployment to be sure the planned IPTV solution will be supported by an organization’s network. Since the amount of bandwidth used to support an IPTV deployment can vary widely depending on the type of multimedia content and how it will be distributed, there are no simple guidelines on how much network bandwidth an organization will require to deploy IPTV. Instead, it is important that the organizations IT group be included in the planning so they can optimize their network to support the specific IPTV deployment.

Below are the preferred minimum specifications for the network switches used for an IPTV implementation to operate correctly, all job and consulting specifications need to clarify this with the end users and dealers.

Required:

- Gigabit layer 2/3 Switch with IGMP Snooping/Querier support.
- 10Base-T/100Base-TX/1000Base-T.
- 8 - 16 - 24 or 48 Ports
- Layer 2/3 switching
- IGMP Snooping V1/V2 (Not router dependant for operation)
- IGMP Querier
- Jumbo frame support
- QOS
- Switching Capacity Non Blocking 32Gbps (16Port) - 140Gbps (48Port)
- Forward rate 24Mpps (16Port) - 96Mpps (48Port)
- Storm Control Broadcast and Multicast
- Spanning Tree
- IPv4 and IPv6

Optional:

- VLAN support
- Web User Interface