IP Multicast in a Nutshell

The basics of implementing IP Multicast for video streaming on the UEN network.

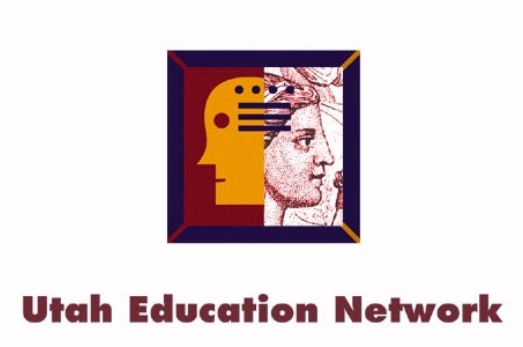

Prepared by
Michael Downie
Team Leader, UEN Field Operations
October 20, 2000

IP Multicast in a Nutshell

Traditional IP allows hosts to communicate in a variety of ways. The most common schemes are:

Unicast- one host communicating with one other host (one to one).

Broadcast- one host communicating with all other hosts (one to all).

IP Multicast provides a third scheme:

Multicast- one host communicating with a group of hosts (one to some).

Packets delivered to a group are identified by a single multicast group address. Multicast packets are UDP. The multicast environment consists of senders and receivers. Any host can send to a group, but only members of the group will receive the message.

Cisco Systems has built multicast support into their IOS's for quite awhile now. There is nothing that we need to do to our present routers, except to turn it on. (Sounds deceptively simple!) The main point to remember is that traditional IP is concerned with where a packet is going. In IP multicast the emphasis is on where the packet *came from*.

Basic terms

There are some basic terms used in IP Multicast, especially in Cisco System's implementation of multicast:

Group- a number of hosts that form a "membership" to receive an IP stream. Membership is dynamic; hosts can join or leave at anytime. There is no restriction to the location or number of hosts. A multicast address is chosen for the receivers in a multicast group. Senders that use that address as the destination address of a datagram reach all members of a group.

Source- a router that has a video server on a directly connected subnet, or the server itself.

Rendezvous Point- in PIM sparse mode, a router on the network where senders and receivers meet. The Rendezvous Point router keeps a table of available sources and directs requests for services to those sources. It is not important where the RP is physically located, only that it has a reliable connection.

Multicast addresses

A multicast address is chosen for the receivers in a multicast group. Senders (sources) use that address as the destination address of a datagram to reach all members of a group.

Multicast uses Class D addresses: 224.0.0.0 to 239.255.255.255.255

The high order bits are 1110. Some addresses are reserved for group broadcasts and established sources on the Internet.

Protocols

Described below are some of the protocols used in Cisco's implementation of IP multicast. There are many others that are used to cross autonomous systems and on the Internet backbone. The protocols described here are the ones that will apply directly to UEN's implementation.

IGMP- Internet Group Management Protocol, used between hosts on a LAN and the routers on that LAN to track the multicast groups of which hosts are members.

PIM- Protocol-Independent Multicast, versions 1 and 2. Version two is the preferred version. If a router detects version 1 it will make all devices use version 1 until version 1 can be replaced. PIM is used between routers so that they can track which multicast packets to forward to each other and to their directly connected LANs.

CGMP- Cisco Group Management Protocol, used on routers that are connected to Catalyst switches. Very similar to IGMP.

IGMP

IGMP is used for hosts to tell directly connected routers about group membership. The information is stored in tables on the router. The process begins when a host sends an IGMP Report saying that it is joining a group. After the initial Join message the router will send periodic queries to the subnet to see if any hosts are active. If a host leaves it will send a Leave message to the router. The router will query to see if there are any other active hosts on the subnet. If it does not receive replies to its queries then the group will time out.

PIM

There are two modes of operation for PIM- dense mode and sparse mode.

PIM-DM
In dense mode the source router assumes that all other routers want to forward multicast packets and it floods the network. If a router in the network has no directly connected group members or PIM neighbors it sends a Prune message back towards the source. Multicast packets are no longer sent down that segment. The pruning continues until all unnecessary branches are pruned. The flooding and pruning process is repeated every three minutes.

PIM-SM

In sparse mode the source router assumes that other routers do not want to forward multicast packets unless there is a specific request to do so. Sparse mode uses a Rendezvous Point (RP). Senders and receivers meet at this point to learn of each other's existence. When a host joins a multicast group the directly connected router sends a Join message towards the RP. The RP keeps track of groups. The RP then sends Join messages towards the source. At this point the source sends packets to the RP and the RP forwards them to the destination router. This forms what is called a shared distribution tree. Once traffic is established, the router that is directly connected to the host will seek a more direct route to the source, if available. It will consult its routing table and send a Join message on the more direct route. Traffic then begins flowing down this segment. Prune messages are then sent to prune back the link through the RP. This results in a source based distribution tree.

The PIM-SM protocol is the preferred protocol for Cisco networks and should be the model implemented on the UEN network.

To Be Considered

To decrease the amount of headaches when multicasting is implemented, Cisco recommends that all routers be running version two of PIM. Version two is the default used in IOS version 11.3.T and above. This means upgrading most of the 2500 routers on the network.

Cisco highly recommends turning on multicasting on all routers in the network. In their words: "JUST DO IT! Enable multicast on all interfaces in all routers in a given network. Failure to do so will require complex (and typically unnecessary) Multicast Traffic Engineering." Enough said.

To reduce bandwidth problems, keep the server as close to the source router as possible (as in directly connected).

Rate limiting is possible on a per-interface or per-individual flow basis. We can also rate limit video and whiteboard streams.

UEN's Implementation of IP Multicast

Traditionally, UEN video services have been available in a fully interactive environment. Instructors and students at each location are able to see and speak to each other in real time. It has become apparent that many courses do not require full interactivity. In order to expand the resources available on the UEN network we are making these less interactive courses available via video streaming. In order to make our streaming content as flexible and available as possible, we are multicasting this content.

Multicasting must be implemented carefully as it has the ability to overwhelm networks. As mentioned earlier, PIM Dense Mode floods a network before pruning itself back. PIM Sparse Mode solves this problem. Also, streaming video servers use unicast SAPs to advertise their services. These advertisements in themselves can kill a network. In delivering multicast video we must carefully separate the multicast traffic.

Attachment 1 *UEN Hubs/Rendezvous Points* diagrams the logical connections on the UEN backbone and locations of video streaming servers. At the present time servers are installed at Southern Utah University, Eccles Broadcast Center, and Davis Applied Technology Center. In the near future servers will be installed at Utah Valley State College, Salt Lake Community College, and Snow College South (SVATC). Although each hub site router could be configured as a Rendezvous Point, we will initially only configure Core 1 and Core 3 because of their central location.

Attachment 2 *UEN Hub Site Configuration* diagrams the equipment installed at each UEN hub site. Today's unicast video services are supplied by providing MPEG codecs with audio and video from the BTS switch. The codecs have multicast capabilities. The video streaming servers will be attached to the Cisco Catalyst Switch or the Cisco 7xxx routers. The stream will be delivered to the end sites via the Cisco MGX Edge Concentrator and Inverse Multiplex over ATM (IMA) links.

Attachment 3 *UEN Video Services Endsite* shows where the critical segmentation of multicast traffic will take place. Video data and LAN data are received by the Cisco 3640 router and delivered to the Cisco Catalyst 2924XL switch via a dot1q encapsulated trunk. The VLANs that are set up on the switch segment the traffic. LAN data is delivered on one segment, traditional EdNet video data on another segment, and new multicast services are delivered on a third. The multicast data will be delivered directly to the EdNet classroom. It will not traverse the local LAN.

This initial deployment will be very controlled until multicast is well tested. In the future local school administrators and facilitators may choose to deliver multicast traffic on their local LAN. At that time the local topology must be evaluated to see if it can handle the demands of IP Multicast. Special consideration must be given to whether there is Layer 3 routing available on the LAN. A flat topology with only Layer 1 or Layer 2 devices will not be able to segment the multicast traffic, resulting in every node being flooded. In this scenario, one student pulling down a video stream will bring the whole network to a standstill.

Conclusion

This is a very simplified introduction to IP Multicasting. As we move forward more details will be provided to those who will be involved in implementation and configuration. Cisco's implementation is very straightforward. IP Multicast is being aggressively developed globally and UEN is excited to be able to expand our services with this technology.

UEN HUBS/RENDEZVOUS POINTS

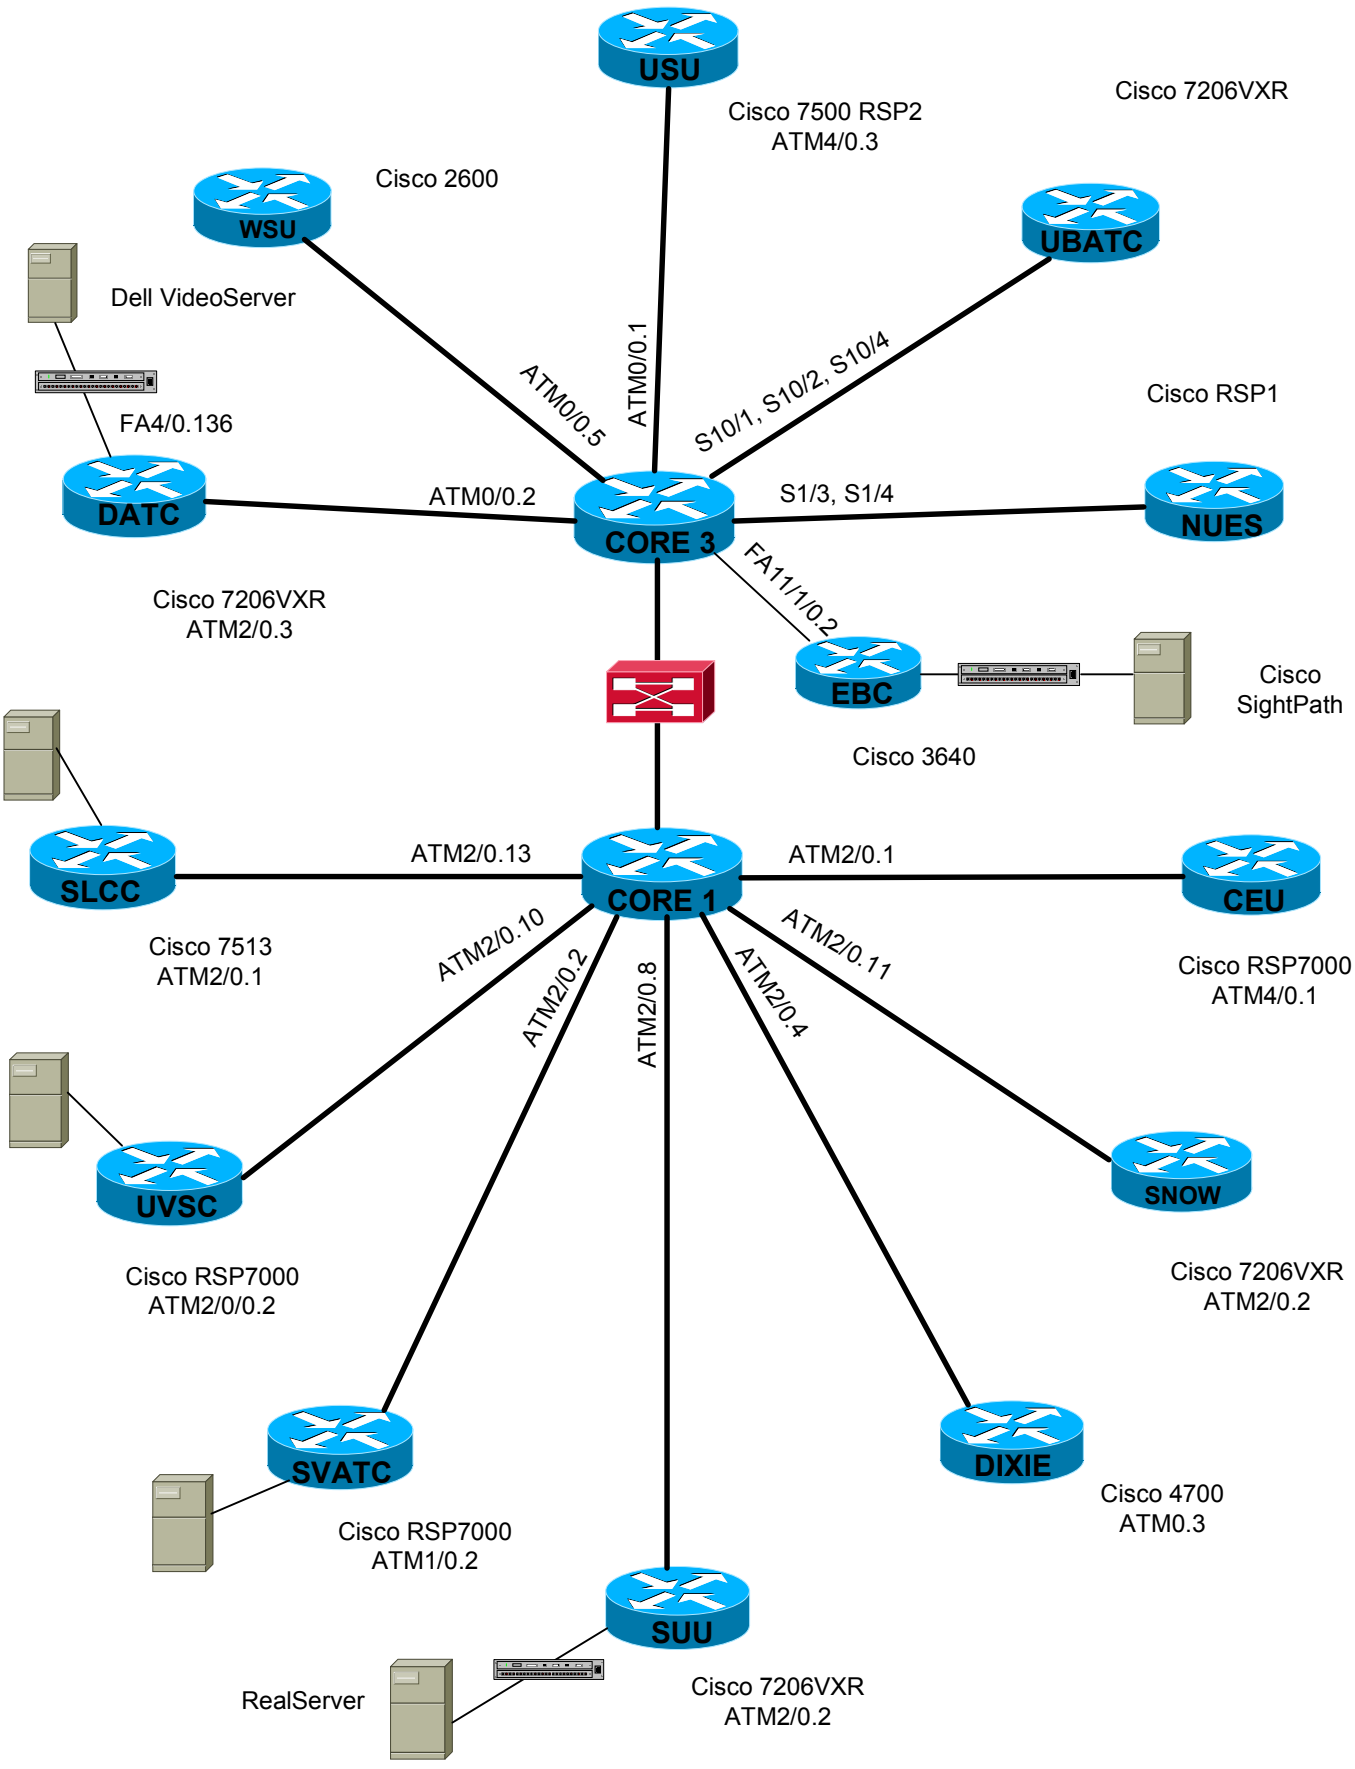

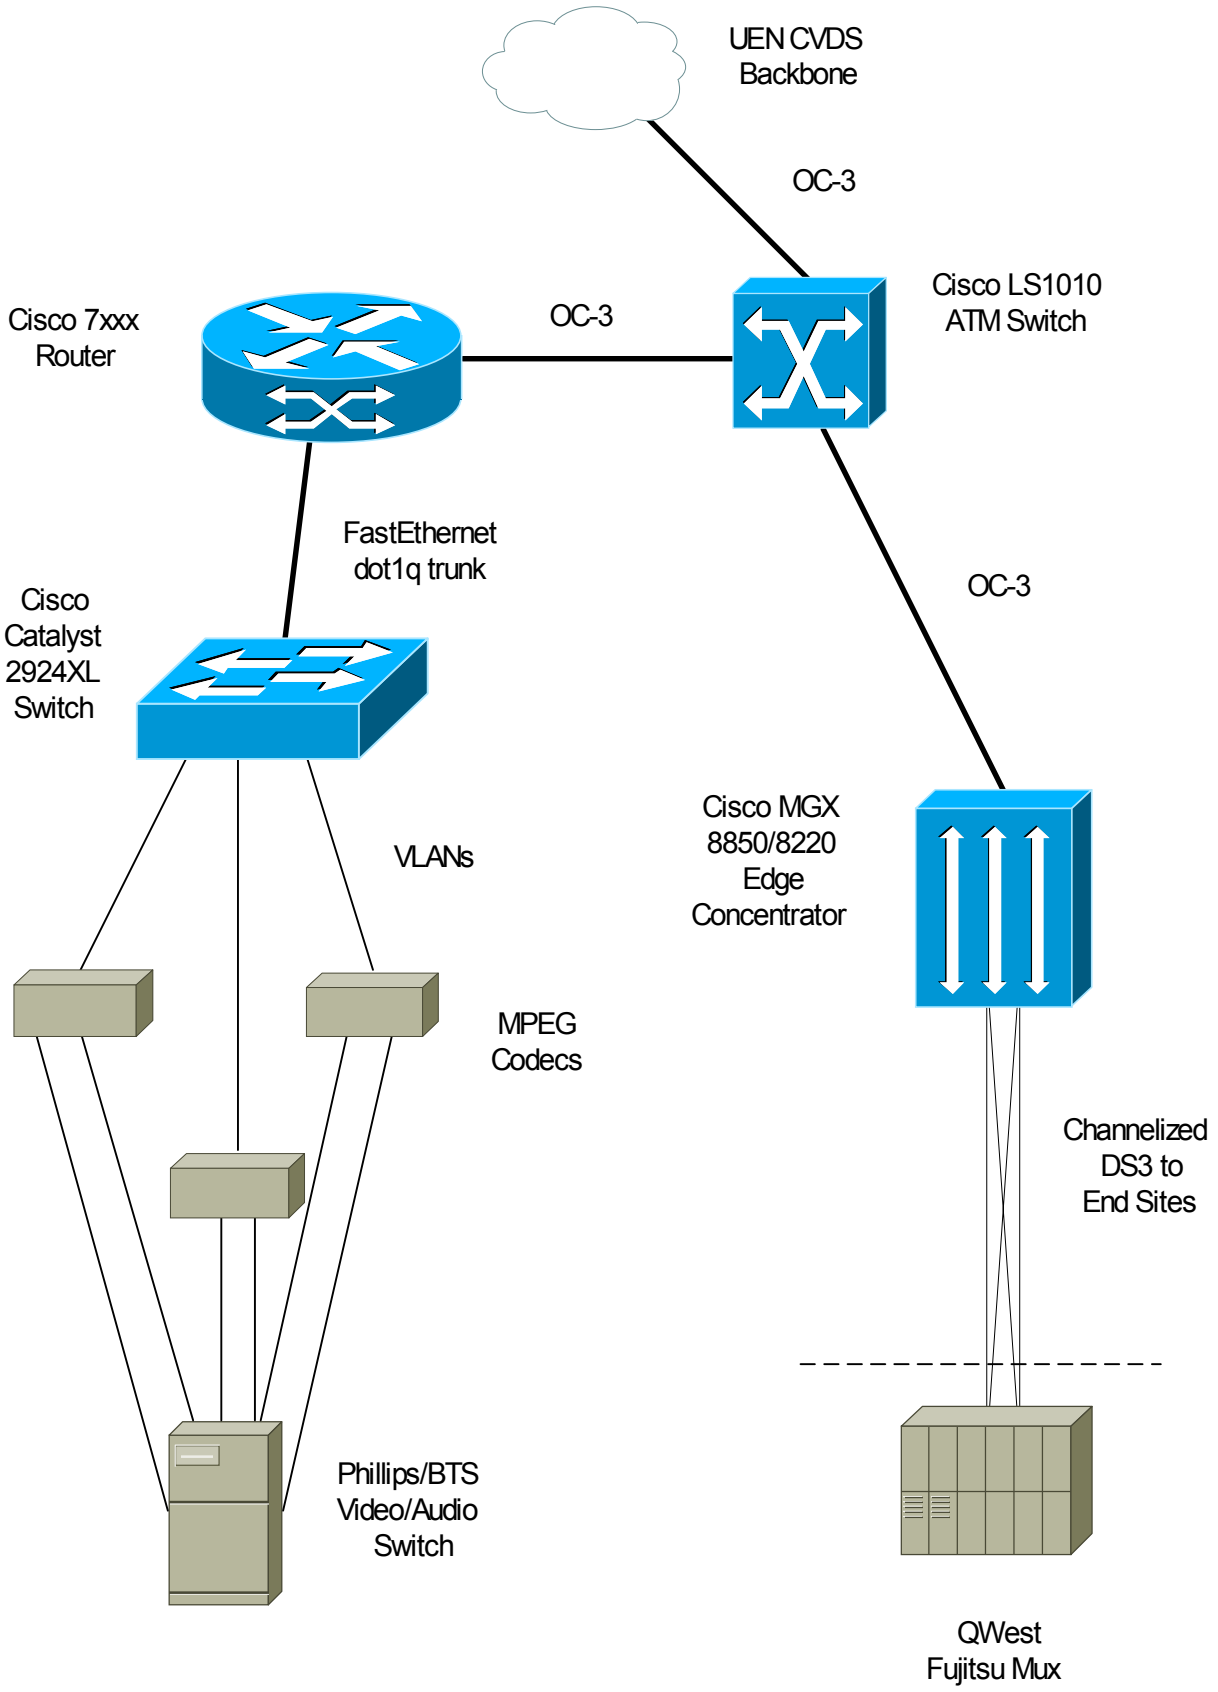

UEN Hub Site Configuration

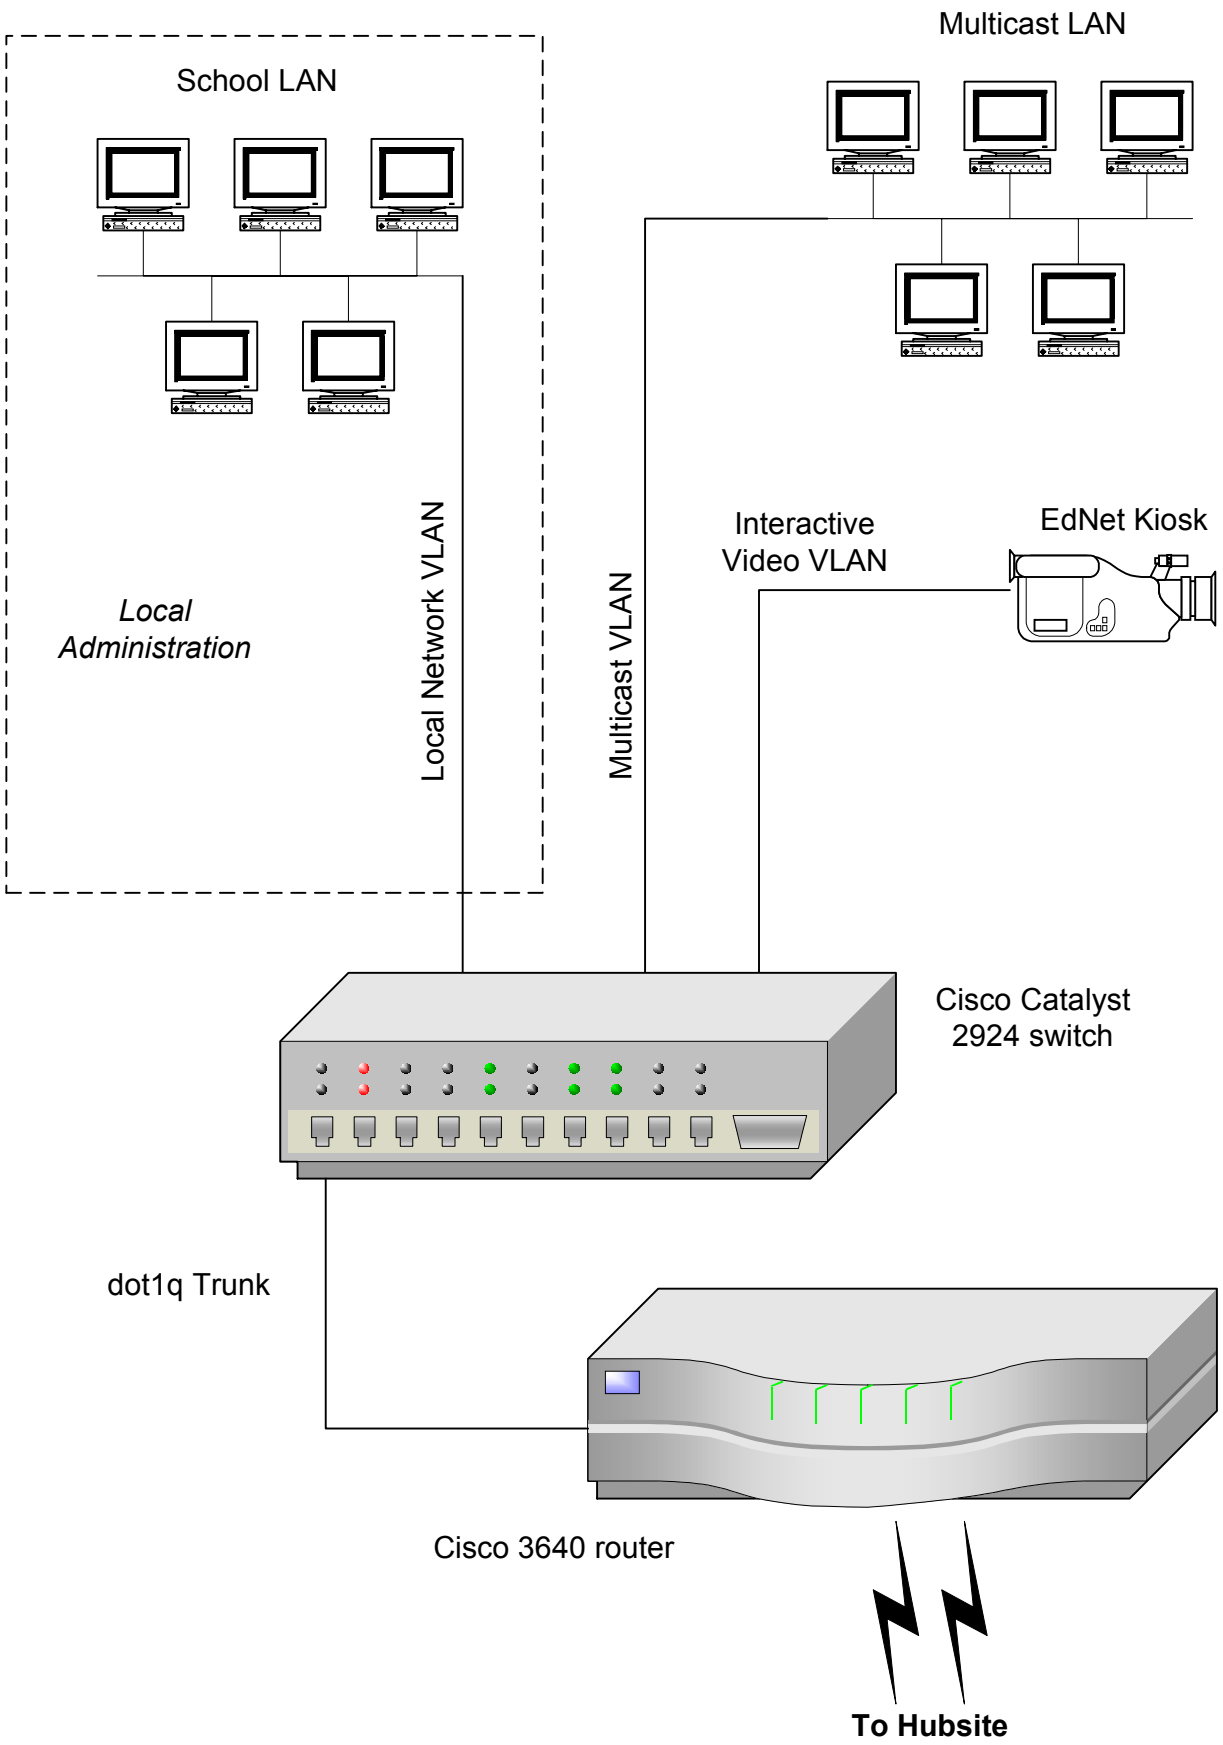

UEN Video Services Endsite