Video Bandwidth

Will MPEG Video Kill Your Network?

The thought that more bandwidth will cure network ills is an illusion like the thought that more money will ensure human happiness. Certainly more is better. But when "There is not enough of bandwidth" is stated as quickly as "We can't afford it", either or both statements may have been offered to dismiss a request because of misunderstanding.

This paper attempts to explain bandwidth perceptions, issues and solutions as it relates to sending MPEG video over modern networks.

Bandwidth Rule #1:
Applications will grow to fill available bandwidth

Bandwidth Rule #2:
Network bottleneck points limit the *apparent* bandwidth

Bandwidth Rule #3:
Multicasting saves an enormous amount of bandwidth

Bandwidth Rule #4
Quality of Service affects both real and apparent bandwidth

Broadband?

In recent years, the term "broadband" has been used so widely it's beginning to lose its meaning. Cable companies and DSL providers would have you believe that "broadband" is anything better than a dial-up connection.

But the industry has long recognized three bandwidth segments, defined as follows:

- Narrowband – 0 to about 56 Kbps.
- Wideband – 56 Kbps to 2 Mbps.
- Broadband – Above 2 Mbps.

Narrowband defines the speeds provided by analog modems, wideband is the T1 and E1 data range, and above T1/E1 we have broadband. Without a qualifier (such as "broadband-*access*"), the terms imply a sustained and continuous throughput capability. For example, saying you have a T1 connection implies you have 1.536 Mbps of connectivity. But if your T1 connects you to a Frame Relay network that delivers just 512 Kbps to you, saying you have a "T1" can be misleading[1].

Let's say an Internet service provider has a DS3 (45 Mbps) connection to the Internet. If that provider had 45 subscribers, they might fairly state that each user has 1 Mbps of bandwidth. But if they have twice that number (2:1 over-subscription), can they still make the claim? How about at 100:1 over-subscription? 1000:1?

The answer lies in subscriber usage patterns, and the expected nature of the data. Subscribers do not generally send 1 Mbps of data all of the time, and this fact makes room for statistical gain. In other words, it does not matter if there is a million:1 over-subscription as long only a few subscribers are actually using the network at any given instant. As networks gain more subscribers, and as those users become more dependent on the network, usage goes up which drives performance down. For example, a DSL or cable modem provider may claim high-speed local connectivity, but at some point all of the users squeeze through the provider's Internet access pipe that is typically much smaller than the sum total of bandwidth available to all users.

The same is true for our 10 or 100 Mbps private Local Area Networks (LAN) too, but now we are talking about true broadband that delivers your 1's and 0's at blistering speeds.

The Fire Hose Principle

Modern networks deliver a 10 or 100 Mbps connection to each and every computer, and often 1000 Mbps to high

---

[1] Another example would be fractional T1, which is sold in increments of 64 Kbps. Like Frame Relay, you would have a full T1 connected between you and your provider, but less than the full T1 is actually available for use.

capacity shared devices like file servers. Inside our buildings, the garden hoses that once interconnected our computers have given way to fire hoses.

This has happened because the cost of high-speed local area networking has dropped from over $1,000 per connection just five years ago, to under $100 per connection today. This 10X cost reduction also brought a dramatic increase in other network capability, including a move to fully switched Ethernet, data priority mechanisms, better management techniques, and much more.

But in most cases, Wide Area Network (WAN) bandwidth did not see the same dramatic cost reduction. Hence the gap between WAN and LAN bandwidth has only widened. While 5 years ago we may have had a 10 Mbps Ethernet network connected to a 256 Kbps private WAN via Frame Relay, today we often have a 1000 Mbps networks connected via T1 (1.536 Mbps). And the post-Internet bubble demise of promising new native LAN wide area carriers has not helped matters. At the same time, our traffic patterns have changed. Not long ago, the primary destination for data traffic was "inside" our networks: file servers, printers, mail servers, etc. But today, the "outside" World Wide Web has become a dominant traffic destination, putting even more stress on the WAN.

"WAN" once meant a point-to-point connection in a private network, but today it usually means "a connection to the Internet". This change in meaning is not trivial because the behavior and capabilities of the Internet are quite different from the behavior and capabilities of a private network. Moreover, "off network" traffic (that is, data that originates in your LAN but is destined for the Internet) can easily saturate expensive WAN bandwidth.

> "It would be easy to conclude you are bandwidth-challenged everywhere when in fact you only have one bottleneck point"

So today we are faced with the Fire Hose Principle: Connecting a LAN to a WAN is like drinking from a fire hose, *and the user evaluates the performance of the LAN based on the performance of the WAN*. It would be easy to conclude you are bandwidth-challenged everywhere, when in fact you only have one bottleneck point.

It is a common misperception that our local networks are saturated. In fact, most networks are far from saturated. Like a sports stadium with only one entrance, getting through the gate can be a problem: once you are inside there is plenty of room. Live and stored DVD-quality MPEG video typically originates and terminates within our true broadband networks (i.e. our LAN's), which are more than able to carry the traffic.

Do The Math

Local area networks are built using Ethernet switches and a Category 5 wire connecting each port of a switch to each computer. If a switch has 16 ports, and each port is operating at 100 Mbps, then the switch would need to support 1.6 Gbps (100 Mbps x 16 = 1600 Mbps) for it to be "non-blocking"[2]. Happily, modern Ethernet switches are fully non-blocking, and a 1.6G switching capacity for a 16 port switch is today as common as 2.4G capacity is for a 24 port switch.

But non-blocking switching is really only meaningful if there is a higher speed port that can accept all of the wireline speed data from all of the other ports, or with somewhat artificial traffic patterns: port 1 sends to port 2, port 3 sends to port 4, and so on. The reality is that normal traffic patterns typically require ports 1 to 15 to <u>all</u> send data to port 16 -- because port 16 may be the port that connects

---

[2] A non-blocking switch has enough switching capacity for all ports to transfer data at wireline speed to any other port. Most modern Ethernet switches are non-blocking, although older switches that are not non-blocking are still in use.

the workgroup to the corporate backbone and then on to the Internet.

One could easily conclude that sending 1.5 Gbps (15 ports at 100 Mbps each) to a single 100 Mbps port would be a huge issue, but surprisingly it is not. This is because while each computer may send data at the 100 Mbps *rate*, they don’t send very much data! For example, consider what happens when you download a 1 MB file over a 100 Mbps network:

1 MB file = 1,048,576 x 8 bits = 8,388,608 bit file

100 Mbps = 1/100M = 0.00000001 seconds per bit

8388608 x 0.000000001 = 0.08388608

It will take only 83.8 *thousands of a second* to download your file (it actually takes much longer than this because of computer disk operations and other factors). The point is that you are not using the network at all for the vast majority of the time, leaving time for others to use it. The sharing of an uplink from your workgroup switch, like the sharing of your WAN connection, is possible because of the bursty nature of most data sources and the statistical nature of the network usage. As long as there is not too much data, all is well.

But as traffic increases, the likelihood of multiple users contending for the same network port increases. At some point, typically about 80% of network capacity, there is so much contention that the network seriously slows down, leading to complaints. Therefore, higher speed uplinks such as Gigabit Ethernet (1000 Mbps) are a superior solution.

Considering the above discussion, one might conclude it would be a bad idea to send an 8 Mbps video stream from one port to every other port of an Ethernet switch. That would require the source port to provide 120 Mbps (15 x 8

Mbps), well in excess of a 100 Mbps port's capacity, right? Wouldn't this more than saturate a 100M uplink? Wouldn't all mission-critical applications slow to a crawl? Not when sending well-regulated video and when multicasting techniques are employed.

Multicasting To The Rescue

While conventional packet data is normally sent from one source to one destination, multicast traffic is sent from one source to multiple destinations but *without* using more bandwidth.

With multicast, the source delivers only one packet stream to the switch (for example, at exactly 5 Mbps), and the switch replicates the packets and delivers them to anyone connected to that switch that requests them. In this local Ethernet switch environment, it is rather pointless to worry about bandwidth when everything is happening at wireline speed.

Modern Ethernet switches replicate multicast packets locally without using *any* additional uplink bandwidth. As a result, sending 5 Mbps to every user will have the same network load as sending 5 Mbps to one user.

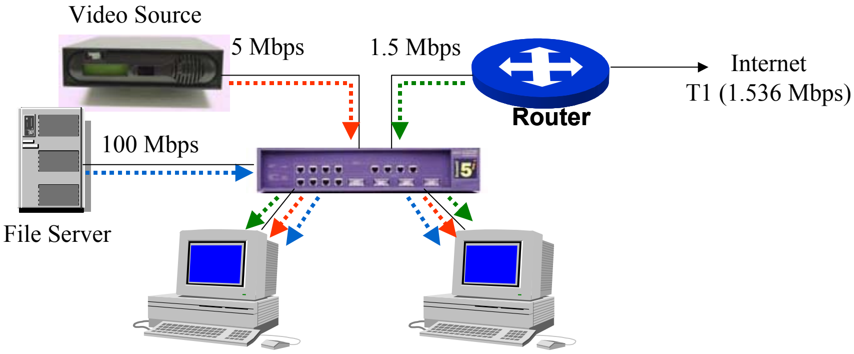

10/100 Mbps Ethernet
- 5 Mbps video to every desktop
- Internet Access via T1
- 95% of Ethernet capacity shared access to Internet & File Servers

But if one Ethernet switch has 16 ports, and one port is connected to the router and 15 ports are connected to users who wish to each view the 5 Mbps video, wouldn't it require 75 Mbps (15 x 5 Mbps), dangerously close to maximum uplink capacity? The answer is no, because there is only one stream coming from the video source and it is delivered via multicast.

Bursts and Priority

In the above discussion, a 1MB file is transferred in about 83 milliseconds. It is important to understand that the intended nature of the Local Area Network is to allow everything to happen as quickly as possible. If you send a 1MB file, the network will attempt to use all of the bandwidth to complete the transfer. If you have 10 Mbps Ethernet, the network will try to send your file at 10 Mbps for as long as it takes; if you have 100 Mbps Ethernet, the network will try to use 100 Mbps for as long as it takes. In other words, you are "betting" that your file will be done before someone else needs the network[3].

Most routers support priority routing based on the "DiffServ" field in the IP header (the TOS field). You can safely prioritize VBrick video, and you can directly set the priority level in the VBrick appliances.

With this in mind, you can see why giving one network user priority over another can become complex. If one user can send data at 100 Mbps on a 100 Mbps network, and if they have priority over everyone else, that priority user could lock out everyone else each time they use the network!

However, if a device such as a VBrick encoder were given top priority, it could never use more than the rate at which it was running. For example, if a VBrick were sending video at 5 Mbps, it would use exactly 5% of the 100 Mbps Ethernet connection at all times. It could never use more because the video is a well-regulated continuous stream that does not burst, unlike conventional web, email, file transfers, and other traffic.

---

[3] There are many complex mechanisms that prevent one user from consuming all bandwidth for too long and enables bandwidth sharing, but the general idea of rapidly bursting your data in the LAN is a fundamental principal of modern networks. Many router and switch vendors have implemented policy-based features to control priority and QoS.

Exactly 95% of the Ethernet port would simply be unused. To the extent the video data were to leave the Ethernet switch via an uplink port (perhaps destined to a router), it will use exactly 5 Mbps, never more. If that uplink were 100 Mbps, 95% is available for other traffic; if that link were Gigabit Ethernet, exactly 99.5% remains available for other traffic. Using our previous example, if a 5 Mbps MPEG video stream were present, a file transfer that might otherwise require 83 milliseconds would now require 88 milliseconds – not much of a difference!

Mix It Up

For the most part, Ethernet and IP networks have grown in an unplanned way. It is a rare IT manager who actually has an up-to-date map of their network, and it is not uncommon for there to be pockets of old shared-media wiring hubs in some areas and modern switches in other areas.

Hubs do not support multicast and can be a problem for the deployment of network video. In fact, hubs do not really support unicast since all computers connected to a hub receive all traffic at all times.

Video can still be successfully deployed with hubs, but the trick is not to have *too much* high bandwidth traffic. For example, if 15 computers were connected to a hub via 10 Mbps and one 5 Mbps video data stream were present in that hub, all computers would receive the stream (whether they like it or not…just like they receive all email, web, and other traffic whether they like it or not).

Since the video is a continuous stream, the effect on a 10 Mbps Ethernet network is to reduce network capacity by 50% (5 Mbps / 10 Mbps). <u>However</u>, this fact alone may not have any practical meaning! If a hub-based network is used primarily to access the Internet via a T1, the real maximum capacity of the network is only 1.536 Mbps, meaning 8.464 Mbps (10 Mbps – 1.536 Mbps) is not used. In this case, adding 5 Mbps to the mix has no adverse affect. If two such 5 Mbps streams were added, there would not be adequate bandwidth on a 10 Mbps network…although there would be ample bandwidth on a 100 Mbps network.

In a mixed corporate network, it would be good practice to deploy multicast video in switched Ethernet segments, and to filter it out, or allow only a limited number of lower bandwidth streams to flow to hub-based segments. With hubs and other legacy devices in your network, the best practice is to go slow, and try it before committing to full-scale deployment in those areas of your network.

The Bottom Line

Very high quality video is easily deployed on modern networks, and even on networks that are not so modern. Multicasting makes it possible to practically eliminate bandwidth concerns, but for some organizations, multicasting is new[4].

Perceptions still linger that video requires more bandwidth than is available. While this can easily be true for wide area networks, it is rarely true for local area networks particularly with good network knowledge and pre-deployment planning.

With simple, straightforward and conventional network planning, an unlimited number of users connected to a broadband network can reap the benefit of DVD-quality video on desktops and TV monitors for better communications, training, and enhanced security and monitoring.

---

[4] The history of the Internet Protocol shows that multicasting has been with us longer than the world wide web!