Video Internet: The Next Wave of Massive Disruption to the U.S. Peering Ecosystem (v1.5)

William B. Norton <wbn@equinix.com>

Abstract

In previous research we documented three significant disruptions to the U.S. Peering Ecosystem as the Cable Companies, Large Scale Network Savvy Content Companies, and Tier 2 ISPs started peering openly. By peering content directly with eyeballs, they effectively bypassed the Tier 1 ISPs resulting in improved performance, greater control over the end-user experience, and overall lower operating costs.

This paper predicts a new wave of disruption that potentially dwarfs this previous redirection of Internet traffic. Short video clip web sites, full length motion pictures, and television shows are now available via streaming to on-line devices and via downloading to iPods. More sites are coming on-line High quality movies from independent producers are being distributed via peer-to-peer methods. We observe these flash crowd effects and the larger movie file sizes as the crest of the first wave of significant incremental load on the Internet.

The majority of this paper details four models for Internet Video Distribution (Transit, Content Delivery Networks, Transit/Peering/DIY CDN, Peer2Peer) across three load models. The cost models include network and server equipment along with pricing models for various distribution methods. Over one hundred walkthroughs of this paper have led to stepwise refinements of the models and insights into why one would prefer or not prefer one model over the other.

The summary of the paper is a comparison of these video distribution techniques in terms of $-per-video units from the Video Service Provider perspective. We highlight cascading obstacles preventing large scale delivery of video traffic using commodity transit in a single location. The CDN solution and the multi-site Transit with Peering solution bypass some of these obstacles, while the peer-2-peer solution, while controversial, yields (by far) the lowest cost solution from the video service provider perspective.

Previous Wave of Evolution of the U.S. Peering Ecosystem

The U.S. Internet Peering Ecosystem went through three significant disruptions in or about 2001[1]:

1. **Cable Companies Peer.** The North American cable companies' Internet transit provider (@Home) went bankrupt in 2001, forcing the cable companies to build out and manage their own multi-gigabit-per-second Internet infrastructure with only 30 days notice. With peer-2-peer traffic representing 40% to 60% of their transit bill, they quickly recognized the benefits of peering[2] that traffic directly with each other.

2. **The Large Scale Network Savvy Content Providers** entered into the Peering Ecosystem as their traffic volume grew into the ten's of gigabits-per-second. By engaging in peering directly with the Tier 2 ISPs, both groups were able to improve performance and lower their transit expenses while enhancing and increasing control over the end-user experience.

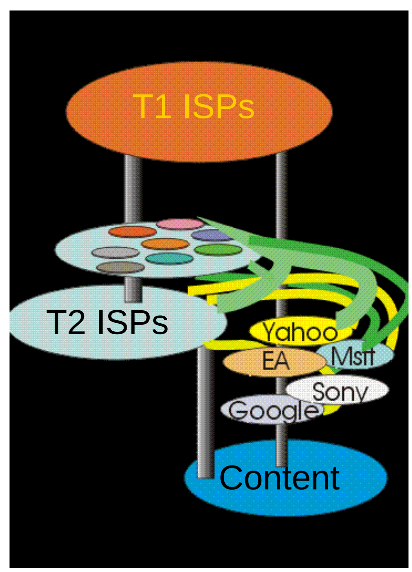

Figure 1 – 2001 U.S. Internet Peering Ecosystem Evolution

3. **MSOs peer with Content.** Since then, the cable companies peer directly with each other and with the large scale content companies. Several of these content companies have

[1] See "The Evolution of the U.S. Internet Peering Ecosystem" for a more detailed discussion of this.

[2] The term "Peering" is the reciprocal (and usually free) exchange of access to each others customers.

evolved their peering infrastructure to include global backbones and hundreds of peering sessions[3].

Peering directly saved these companies at least $1 million dollars a year each and it is estimated that this early wave of disruption to the Peering Ecosystem resulted in at least 100Gbps of peered traffic.

Internet Video Traffic: The Next Wave of Massive Disruption

In 2006 we saw early indicators of an emerging and massive growth in the scale of potentially peerable traffic that may dwarf the previous wave of disruption:

YouTube, a one-year-old community-based short video sharing service, disclosed in February 2006 that they were purchasing transit for 20Gbps of video traffic[4]! Their projected growth rate, shown graphically below, was documented as 20% compounded monthly, suggesting 100Gbps of transit expenses before the end of 2006!

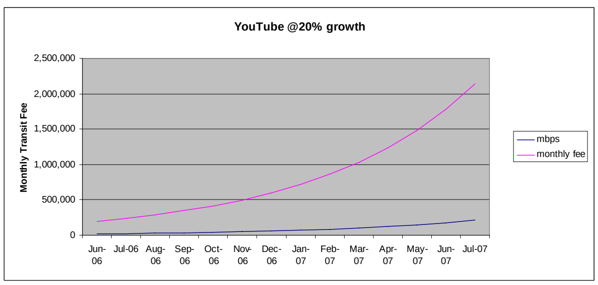

Figure 2 - YouTube Internet Traffic Projection

The author consulted several venture capitalists in the silicon valley area who suggested that, with the popularity and press that YouTube was receiving, there were likely hundreds of direct and indirect competitors seeking to provide variations of the YouTube theme of video and social networking. If we assume that the top 30 competitors eventually reach an average of 20Gbps of load, we see a potential market of up to 600Gbps of incremental video traffic! This alone would represent six times the 100Gbps previous wave.

[3] Brokaw Price, Sydney Peering Forum indicated that Yahoo! Had over 640 peering sessions and was addressing the challenge of determining which Asian countries were most conducive to regional traffic distribution.

[4] NANOG 37, San Jose, Peering BOF, http://www.nanog.org/mtg-0606/pdf/bill.norton.3.pdf

ABC Television delivers the show "Desperate Housewives" over the Internet. Each one hour episode is a 210MB file[5] for a 320x240 H.264[6] image suitable for display on a video iPod or a small window on a laptop PC. Following the Cringely logic[7], with an audience of 20,000,000 people in 10,000,000 households, we can project a maximum video traffic distribution of 2.1 petabytes per episode. If we assume that the distribution will not be all at once but rather distributed over a three day period, we observe that this single episode of a single show would require a 64Gbps pipe filled for 3 days non-stop!

Large Web Properties Predict 1000 fold increase in video traffic by 2010. The author participated in a small private large-scale content provider meeting in which projections for the year 2010 were made. From the largest content providers in the world perspective, based on their collective current and historical vantage points, they predicted that video traffic will represent at least 80 percent of all Internet traffic. David Filo (Co-founder of Yahoo!) suggested that video traffic will more likely represent closer to 90 percent of all Internet traffic. These sentiments were echoed at the June 2006 NANOG meeting in San Jose.

High Quality Video. Some providers are focusing on a slightly different niche – high-quality full-length independent movies published by the producers over the Internet. Jason Holloway (CEO, DoveTail) suggested that the video file size will grow even larger as thousands of hi-def full feature independent film producers vie for public attention as their works are distributed over the Internet[8]. Cringely suggests further multipliers in demand include 180 channels of hi-definition video with ability to time shift. All of these have a multiplicative effect on the Internet traffic load.

As with all predictions, there will be disagreement on the details; nonetheless, this wave of Internet video traffic has the potential to dwarf the previous

[5] This analysis was taken from an on-line by Robert Cringely: http://www.pbs.org/cringely/pulpit/pulpit20060302.html

[6] a digital video codec standard which is noted for achieving very high data compression. http://en.wikipedia.org/wiki/H.264

[7] Ibid.

[8] Red Herring "DoveTail showcases HD Films" http://www.redherring.com/Article.aspx?a=18373

(100Gbps) disruption.

Modeling the Video Service Provider Distribution Networks

How can Video Service Providers distribute this much Internet Video traffic? The rest of this paper documents and compares four models for Internet Video Distribution:

Model 1: Simple Commodity Transit

Model 2: Content Delivery Network (CDN)

Model 3: Hybrid Transit/Peering/DIY CDN

Model 4: Peer-to-Peer Networking

David Cheriton (Sun Microsystems) makes the observation that this is really the question "Where does it make sense to cache the video content? Transit is caching at the seed of the tree, CDNs are caching at a mid-point in the hierarchy (at the peering points), and p2p is really caching at the leaf ("at grandma's house").

We are applying these four models across three load models:

Load Model A – Light Load: Every 5 minutes, **10** customers each start to download a 1.5 GB movie, resulting in an average 15GB five minute load. Whether completed in the five minute sample or spread across many 5-minute samples, the average load on the server will remain 15GB (see Appendix A for a discussion of this laddering effect)[9].

The 15 GB 5-minute load can be converted to GBps by dividing the 5-minute load by the number of seconds in each 5 minute sample:

$$60\frac{\sec s}{\min}*5\frac{\min}{sample}=300\frac{\sec s}{sample}$$

and then converted to Mbps by multiplying by 8 Gigabits-per-second per Gigabyte-per-second and multiplied by 1000 Mbps per Gbps:

$$\frac{15GB}{300\sec s}=0.05\frac{\text{GB}}{\sec}*8\frac{\text{Gbits}}{\text{Gbyte}}=0.4Gbps=400Mbps$$

Adjust average load to sinusoidal demand curve. We have assumed a flat load across the month, but user Internet traffic load historically follows a sinusoidal curve pattern, with peaks during the waking hours and valleys in the middle of the night. The ISPs have a rule of thumb that the 95th percentile measure is typically about double the average for Internet traffic. However, Jeff Turner (InterStream) points out that video traffic demonstrates a peak-to-mean ratio of 6.6:1. This has two significant modeling implications:

1) VSPs need much larger server and network infrastructure to accommodate the peak of 6.6 times the mean. This means that the equipment for Load Model A for example needs to be specified to be able to handle a peak load of 6.6 times the average load, or about 2,640 Mbps if the mean is 400Mbps.
2) the 95th-to-mean ratio for video traffic is probably higher than the 95th percentile to mean 2:1 ratio for other Internet web. As a first pass guesstimate, we will therefore assume a 95th-to-mean ratio of 4:1. This means that the Model A upstream ISP will be charging transit fees on 1,600 Mbps.

Load Model B – Medium Load: Every 5 minutes users start to download **100***1.5GB movies, resulting in an average 150GB five minute load. Applying the same math as above we need to plan to pay transit on 16,000 Mbps and offload about 26,400 Mbps of video traffic.

Load Model C – Heavy Load: Every 5 minutes **1000***1.5GB movies, resulting in an average 1500GB five minute load. Applying the same math as above, we need to pay transit on 160,000 Mbps and plan to offload about 264,000 Mbps of video traffic.

Further Assumptions

The VSP distributes 1000 full-length (1.5GB) feature films. We will assume that the VSP has a library of 1000 videos to distribute, each 1.5GB in size.

Equipment: Richard Steenbergen suggested that we model the servers using an inexpensive 2U Dual Core Opteron systems with 4GB of RAM, 4-500GB Disks, and 2 gigabit Ethernet uplinks to an aggregation switch. Each server can deliver up to 1Gbps aggregated across the two gigabit Ethernet links. The disks will hold 1000 videos at 1.5GB each with the remaining space for formatting overhead, OS, LAMP SW, miscellaneous scripts, etc.

There was some concern raised by reviewers as to whether the server as specified can indeed hit 1Gbps, given the likelihood of the request going to disk. The more simultaneous sessions, the more the disk I/O has the performance characteristics of random disk I/O. The alternatives suggested include: a) expect each

[9] We assume stream homogeneity – specifically, that each stream has an equal share of available bandwidth so they start and stop roughly at the same time.

server to deliver 400-500Mbps, and/or b) deploy SCSI disks and/or shared SANs, and/or c) increase the server size to 3U and use disk striping for increased performance, and/or massive memory configuration with substantially higher cost but increase chance of the I/O hitting the memory cache.

Aggregation Switch. The aggregation switch will handle up to 24 servers using its two ports to handle peaks of up to 24Gbps. We will accept only 20 servers however (across 40 ports) and therefore up to 20Gbps of traffic per switch so we will not oversubscribe to the upstream routers two 10G Ethernet cards.

Router. The router used is a Cisco 6509 class router with 10G ingress and 10G egress port(s) added as load demands. The model is shown in the diagram below.

Sylvie Laperriere suggested that the equipment model should include a 15% maintenance contract so we have factored that into the model.

(It should be noted for all the models that there is great variability in network and server architectures and the corresponding equipment and vendor selection. We have modeled for simplicity instead of completeness and have selected **a** set of equipment, not necessarily the perfect fit for every type of traffic a VSP may be distributing.)

Software costs are zero with LAMP. It is assumed that the Video Service Provider is deploying open source software (as is the case with 60% of all web servers on the net[10]). We will model assuming LAMP (Linux, Apache, mySQL, and PHP) system components so this software is essentially free. We assume no customization or scripting costs.

Colocation Costs: We will assume that each model requires collocation at an exchange point (IX) for two reasons. First, the servers are assumed to be hosted somewhere, so there is a assumed to be some corresponding cost of housing. Second, we will assume an open market for transit, and in some of the models the ability to peer. In the U.S. these are typically accomplished at a collocation center that also operates a public peering fabric and/or private peering facilities.

Multi-homed: We are assuming that the VSP will be multi-homed to more broadly distribute the traffic to handle spikes, and provide some redundancy. Note that this implies:

1. VSP have some level of network expertise to configure, manage, test, etc. the multi-homed router configuration, and
2. the IX has multiple networks capable of handling the multi-homed load.

Disk. Disk is inexpensive so we assume that each server is configured with at least 2TB, enough to hold the videos (10*1.5GB=1.5TB), the OS and support software. More important than the size, is the number and configuration of the disks. This has a material affect on the ability for the system to fetch and deliver the large video content. The system will essential grab data from one slow I/O system and sending it across the Internet through another slow I/O system. These are likely the two bottlenecks to be managed.

Digital Rights Management: We are ignoring the Digital Rights Management (DRM) issues, assuming that this will be settled by the VSP and Content Owners (if they aren't one and the same).

Analysis. The end result of this analysis will be a summary per-video cost comparison table showing the four models across the three load models as shown by the matrix below.

Models	**A:10 videos**	**B: 100**	**C: 1000**
1: Transit	Model 1A	Model 1B	Model 1C
2: CDN	Model 2A	Model 2B	Model 2C
3: Hybrid	Model 3A	Model 3B	Model 3C
4: P2P	Model 4A	Model 4B	Model 4C

Model 1: Simple Commodity Transit for Video Distribution

Business Premise: Video Service Providers outsource Internet video distribution to two or more transit providers in a single collocation center. The VSPs instead focus on the Content **Server** Operations, marketing, sales, etc. and oversight of the entire supply chain for video content. This is usually based on the following assumptions:

a) Transit Providers can handle the Internet traffic better and cheaper because they can leverage economies of scale and aggregation efficiencies.
b) Transit Providers have the network expertise, the peering arrangements, the billing engines, etc. required to maintain and scale an Internet network.

[10] Red Herring, 08.21.06, page 26 "More than 60 percent of the World Wide Web servers, for example, run open-source software."

Definition: **Transit** is a business relationship whereby an entity sells access to the Internet to a customer.

Transit is best considered a pipe in the wall that says "Internet This way". The customer sends its packets out to the Internet, and the transit provider announces reachability of the customers' network to the rest of the Internet.

Transit in the U.S. is a metered service, charged on a megabits-per-second basis, measured at the 95th percentile. Traffic is sampled every 5-minuutes and the deltas are stacked lowest to highest every month. The 95th percentile value is used to determine the volume on which transit is charged.

Wholesale transit prices in 2006 vary widely but appear to hover around $10-$20/mbps with a 1 gigabit-per-second commit as sampled at NANOG in February 2006 and as shown in the graph below[11].

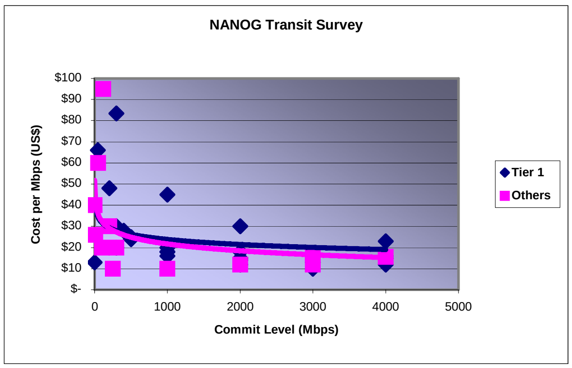

Figure 3 – NANOG 36 Transit price Survey

However, these *wholesale* transit prices are not generally available to content providers. The traffic flow is almost entirely outbound which has the effect of adversely affecting the ISP's traffic ratios[12]. Further, we are assuming multi-homing so the bandwidth commit to any single ISP is split among multiple ISPs. Without the full volume commitment to a single transit provider, pricing may be a bit higher. To properly set VSP expectations, Jeffrey Papen (Peak Web Consulting) suggests we model Video Service Providers (VSPs) using slightly higher (approximately retail) transit prices. Based on walkthroughs of this paper we have settled on the following transit pricing[13]:

- Model A: 1,600Mbps @ $25 /Mbps
- Model B: 16,000Mbps @ $18 /Mbps
- Model C:160,000Mbps@ $10 /Mbps

Note: the pricing here has generated a lot of controversy. About half the reviewers say the prices are about right or a bit lower than the market price, and others stating that these prices are higher than the current market price.

Pro: Transit is Simple. One virtue of transit is that it is simple and the unit price decreases as the commitment increases. There is also a contractual obligation for the transit provider to deliver the content to the rest of the Internet, a feature that we will see later is not a feature of peer-2-peer systems. The transit provider takes the packets and delivers them to the rest of the Internet. This simplicity allows the content provider to focus on content creation rather than network engineering.

Con: Video Transit can be expensive. One potential downside of transit for video distribution is that the traffic volume can grow very large. At the same time, some video content (i.e. YouTube hits, new popular TV shows) exhibit viral qualities so the volume can also spike unpredictably. As a result, the cost can grow large and unpredictably. Consider YouTube for example; founded in 2005 and within a year it was already pushing 20 Gbps of video traffic[14] across its transit providers.

To illustrate the costs involved, let's assume YouTube was able to acquire transit at $10/mbps. With 20% monthly traffic growth, as shown in the graph below, the resulting expense would approach $1 million per month within the next year[15]! This is pretty expensive for a startup company, although we now know with the subsequent purchase by Google for

11 Dave Wodelet (Shaw / BigPipe) shared the results of his transit survey at the previous NANOG 36 in Dallas: http://www.nanog.org/mtg-0606/pdf/bill.norton.2.pdf

12 See "The Folly of Peering Ratios" white paper by the author for a more complete discussion of the ISP's desire to maintain roughly symmetric peering ratios.

13 The nice thing about models is you can change the numbers in a spreadsheet and see the effect. Feel free to request a copy of the spreadsheet from the author.

14 YouTube's Peering Personals slide set at NANOG 37 in San Jose, June 2006: http://www.nanog.org/mtg-0606/pdf/bill.norton.3.pdf

15 Rumors from the field is that in September 2006 YouTube already faces the $1M/month transit expense! –anonymous citation.

$1.65B it turned out to be a good investment..

Month	mbps	monthly fee
Jun-06	20,000	$200,000
Jul-06	24,000	$240,000
Aug-06	28,800	$288,000
Sep-06	34,560	$345,600
Oct-06	41,472	$414,720
Nov-06	49,766	$497,664
Dec-06	59,720	$597,197
Jan-07	71,664	$716,636
Feb-07	85,996	$859,963
Mar-07	103,196	$1,031,956
Apr-07	123,835	$1,238,347
May-07	148,602	$1,486,017
Jun-07	178,322	$1,783,220
Jul-07	213,986	$2,139,864

Figure 4 – Projected YouTube Transit Expenses

At least one source cites YouTube paying one million dollars per month already in September 2006[16], and in March 2007 rumors have their traffic volume exceeding 200Gbps[17].

Now we will model the three video loads across each distribution model, starting with a basic commodity transit service.

Model 1A: Light Load

Here we need to construct a simple system to pass 2,640 Mbps peak worth of videos to distribute to a transit provider.

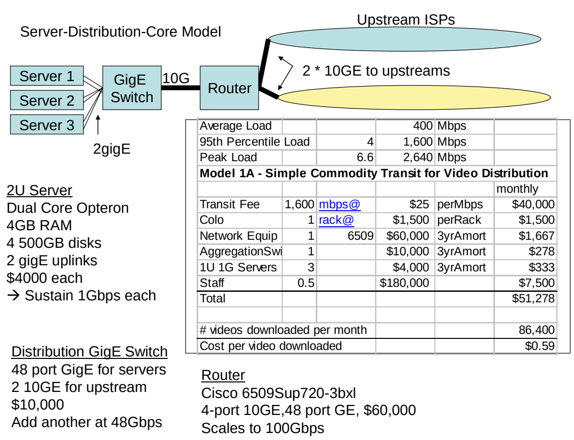

If we examine the monthly cost to the VSP we see that the greatest expense is in the transit fee. All other costs except staff are relatively insignificant. We assume for all models that the fully allocated cost of an experienced single Full-Time Equivalent (FTE) is $180,000 per year.

Average Load			400	Mbps	
95th Percentile Load		4	1,600	Mbps	
Peak Load		6.6	2,640	Mbps	
Model 1A - Simple Commodity Transit for Video Distribution					
					monthly
Transit Fee	1,600	mbps@	$25	perMbps	$40,000
Colo	1	rack@	$1,500	perRack	$1,500
Network Equip	1	6509	$60,000	3yrAmort	$1,667
AggregationSwi	1		$10,000	3yrAmort	$278
1U 1G Servers	3		$4,000	3yrAmort	$333
Maintenance				15%	$342
Staff	0.5		$180,000		$7,500
Total					$51,619
# videos downloaded per month					86,400
Cost per video downloaded					$0.60

As one reviewer noted, a smaller router could have been selected, and no aggregation switch is required for this load model as the servers could be directly attached to a router Ethernet blade. We are assuming however that this VSP is planning to scale the infrastructure as the load increases. The larger router and aggregation switch assumes this scaling will occur at some point in the future.

Model 1B: Medium Load

We can scale Model1A to handle ten times more video traffic if we beef up the server and network infrastructure. At this point we exceed the capacity of the first distribution switch so we spread the load onto a second switch. We have also filled the first switch's uplink capacity (the max 24Gbps from the servers can't all fit across the two 10G Ethernet ports) so we will probably spread the server load more evenly across the switches.

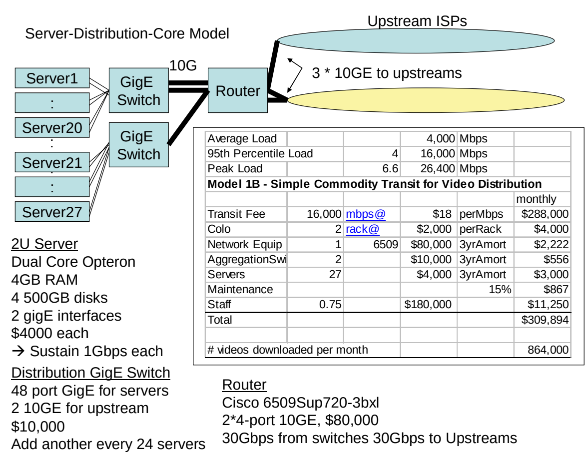

16 http://www.datacenterknowledge.com/archives/2006/Sep/12/youtube_gets_bandwidth_boost_from_level_3.html

17 Industry rumors from a variety of sources.

Average Load		4,000	Mbps
95th Percentile Load	4	16,000	Mbps
Peak Load	6.6	26,400	Mbps

Model 1B - Simple Commodity Transit for Video Distribution

						monthly
Transit Fee	16,000	mbps@		$18	perMbps	$288,000
Colo	2	rack@		$2,000	perRack	$4,000
Network Equip	1		6509	$80,000	3yrAmort	$2,222
AggregationSwi	2			$10,000	3yrAmort	$556
Servers	27			$4,000	3yrAmort	$3,000
Maintenance					15%	$867
Staff	0.75			$180,000		$11,250
Total						$309,894
# videos downloaded per month						864,000
Cost per video downloaded						$0.36

We observe here again that the transit fees substantially dominate the cost of video distribution in this model.

Model 1C: Large Load

With this next model we stress the equipment but can still fit the traffic across four routers each connected to eight upstream ISPs at 10Gbps capacity.

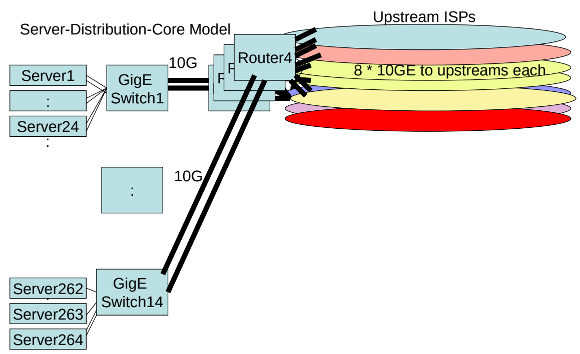

Distribution GigE Switch
48 port GigE for servers
2 10GE for upstream
$10,000
Add another every 24 servers

Routers
Cisco 6509Sup720-3bxl
w/4*4-port 10GE, $150,000
80Gbps from switches, 80Gbps to upstreams

Here we clearly see that the primary cost component for distributing video is the transit expense. It is not even close. We could double or triple the server and networking equipment without making a dent in the overall cost model.

Average Load		40,000	Mbps
95th Percentile Load	4	160,000	Mbps
Peak Load	6.6	264,000	Mbps

Model 1C - Simple Commodity Transit for Video Distribution

						monthly
Transit Fee	160,000	mbps@		$10	perMbps	$1,600,000
Colo	14	rack@		$2,000	perRack	$28,000
Network Equip	4		6509	$150,000	3yrAmort	$16,667
AggregationSw	14			$10,000	3yrAmort	$3,889
Servers	264			$4,000	3yrAmort	$29,333
Maintenance					15%	$7,483
Staff	3			$180,000		$45,000
Total						$1,730,372
# videos downloaded per month						8,640,000
Cost per video downloaded						$0.20

.(Note: There were voices suggesting that the Cisco 6509s may be able to handle this load in the lab but may not work well in this configuration in a production environment. These reviewers suggested spreading the load across more routers and servers.

Cascading Failures. By far the biggest problem with this model is that there is no single location where the Tier 1 ISPs can all provision four ten gigabit per second links to the Video Service Provider. Several Tier 1 ISPs in the U.S. shared anonymously that they all face the following issues with respect to distributing video today:

1) Colocation providers are out of space and power for equipment capable of handling this next wave of video traffic. They are all turning away or queuing dozens of 10gbps orders at places where they house their core peering routers.

2) If space and power were available at the IX, they would need to convince their finance people to allocate budget to replace gear they had not paid for yet. The new gear would be very expensive and support the emerging group of VSPs that expect prices in the $10/mbps range. This may be a tough case to make.

3) Even if a Tier 1 ISP upgraded their equipment, this is not sufficient. Their peers would need to upgrade their peering equipment at the IX as well to handle the much larger peering capacity.

4) Even if the Tier 1 ISPs all upgraded their peering gear this is not sufficient. Since the video traffic does not terminate at the IX, each Tier 1 would need to upgrade their backbones.

5) The final parts of the chain involve the last mile providers who need to have the ability to carry all of this incremental video traffic.

A separate category of problem is that there is no standard for interfaces greater than 10Gbps, and those non-standard solutions that do exist involve trunking together multiple 10Gbps into 40Gbps aggregate links. These links do no benefit from the typical scaling characteristics exhibited before: that twice the price gets you four times the bandwidth. Here you get four times the bandwidth for about four times the cost. A 100Gbps solution appears to be 4-6 years in the future.

The solution to this problem may require the Tier 1 ISPs and IX Operators to synchronize the creation of more space and power and the ordering and deployment of new peering and backbone infrastructure. During the meantime, the Content Delivery Network solution looks like a good solution to explore for large scale video delivery.

Model 2: Content Delivery Networks (CDNs) for the Distribution of Video Content

Business Premise: Single-site transit traffic may not be capable of handling the transit load. Further, traffic traverses potentially many network devices, increasing latency and the potential of packet loss. The CDN promises improvement:

a) By spreading web objects closer to the eyeball networks latency is reduced

b) Fewer network elements are traversed so reliability is improved

c) Congestion points in the core of the Internet are avoided

d) CDNs have the expertise[18], deployed infrastructure, economies of scale from aggregation efficiencies.

Pro: Remove the transit provider(s) bottleneck(s). What is important is the location of the *bandwidth* bottleneck. As David Cheriton says "You always want the content to be on the other side of the bottleneck (e.g. on the same side as the eyeballs)." The bottleneck will affect how large the TCP window size can grow before packet loss and poor performance occurs.

So the CDN distributes the content to the mid point in the hierarchy at the IX or within the last mile network.

Modeling VSPs using CDNs

To properly set expectations, Jeffrey Papen (Peak Web Consulting) suggests we use retail prices for modeling of VSPs using a CDN:

- Model A: 1,600Mbps @ $35 /Mbps
- Model B: 16,000Mbps @ $23 /Mbps
- Model C:160,000Mbps@ $13 /Mbps

In all three CDN models, we do not need as large a server/network infrastructure since the CDN will pull the video objects once to seed the content across the CDN. The CDN then stores the videos at or near the edge, and distributes the videos to the end users from the appropriate "pod".

Model 2A: Light Load

There is some debate surrounding the benefits of a CDN for distributing these large video files.

Barrett Lyon (BitGravity) argues that video files are so large that the TCP congestion algorithm will grow its window size (amount of data transmitted without previous data delivery acknowledgement) to the maximum delay-bandwidth product. The further (latency-wise) away the eyeballs are from the data source, the more packets are "mid-air". Let's assume a loss-free network and no worse bottlenecks between the eyeballs and source. In both the close and distant source cases, once this maximum window size steady state is reached, the packet latency between the end points is less of a material issue.

The counter view is that the window size for the current era OSes has a maximum window size of 64K which is far from enough to fill the delay-bandwidth product pipe with packets.

Others argue the key problem with CDNs like Akamai is that the thousands of servers deployed are small servers incapable of holding the 1000 videos we assume in the model. Therefore, most of the requests will not be served from the edge but rather from a source further into the network. So, what is the incremental benefit of the video stream being served from somewhere in the middle of the Internet as opposed to from the origin server?

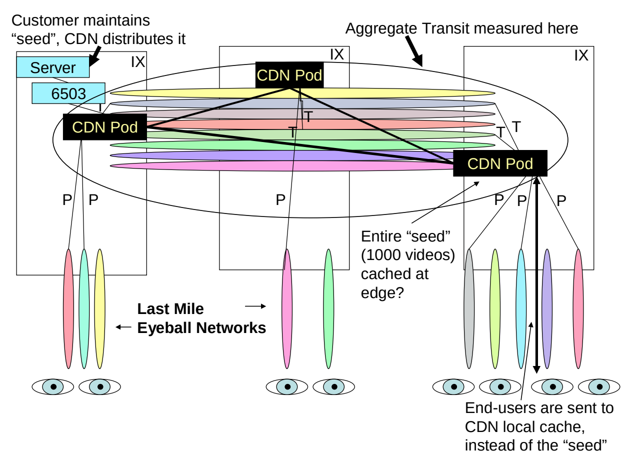

[18] One reviewer considers the specification, operation and tuning of CDN deployments to be a highly specialized skill.

Average Load			400	Mbps	
95th Percentile Load		4	1,600	Mbps	
Peak Load		6.6	2,640	Mbps	
Model 2A - Content Delivery Network for Video Distribution					
					monthly
Transit Fee	1,600	mbps@	$35	perMbps	$56,000
Colo	1	rack@	$1,500	perRack	$1,500
Network Equip	1	6503	$30,000	3yrAmort	$833
Servers	1		$4,000	3yrAmort	$111
Maintenance				15%	$367
Staff	0.5		$180,000		$7,500
Total					$66,311
# videos downloaded per month					86,400
Cost per video downloaded					$0.77

In any case, with model 2A the load is spread across the CDN and measured in aggregate across the CDN. Note that, as with model 1, video traffic distribution (called transit here) is the primary cost factor when outsourcing to a CDN, dwarfing all other costs. Since we pay a premium for CDN services, we have effectively increased the most expensive (top line) costs and reduced the insignificant costs.

Model 2B: Medium Load

The VSP server infrastructure remains the same regardless of the demand; the CDN handles all the distribution complexity of perhaps viral and unpredictable end user demands.

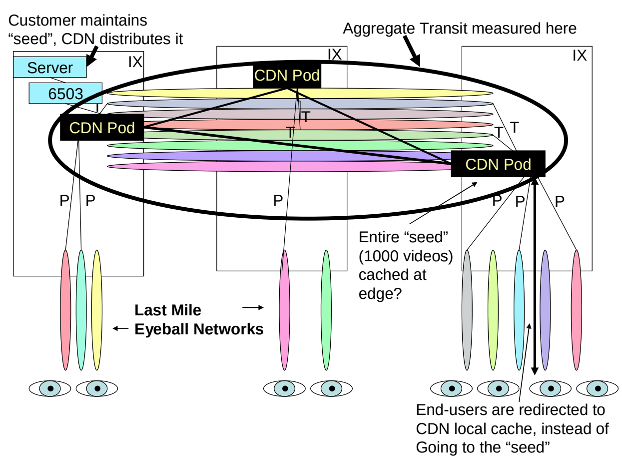

Average Load			4,000	Mbps	
95th Percentile Load		4	16,000	Mbps	
Peak Load		6.6	26,400	Mbps	
Model 2B - Content Delivery Network for Video Distribution					
					monthly
Transit Fee	16,000	mbps@	$23	perMbps	$368,000
Colo	1	rack@	$1,500	perRack	$1,500
Network Equip	1	6503	$30,000	3yrAmort	$833
Servers	1		$4,000	3yrAmort	$111
Maintenance				15%	$367
Staff	0.5		$180,000		$7,500
Total					$378,311
# videos downloaded per month					864,000
Cost per video downloaded					$0.44

As we scale up the CDN to load model C, some reviewers expressed some discomfort outsourcing such a large amount of traffic to any single entity:

First, the Video Service Provider has better and advanced information about the popularity of new releases. The VSP may be able to time and manage network resources associated with these release events. To the CDN, these files are simply customer objects.

Second, while a CDN may be able to aggregate many companies loads efficiently, there is a chance that multiple customers have coinciding spot events that stress/break the shared CDN resources.

Third, there were concerns expressed about outsourcing this key operations activity to and company; a significant spot event, or any network related problems could severely effect the end customer experience and the VSP reputation. For these reviewers, control over the end-user experience was seen as a strategic requirement.

Model 2C: Large Load

No surprises in the large load model, but the scale most clearly highlights the comparative cost of each part of the infrastructure. Again, we are increasing the most expensive component of the distribution (transit) while improving performance and reducing the internal cost of managing the video distribution.

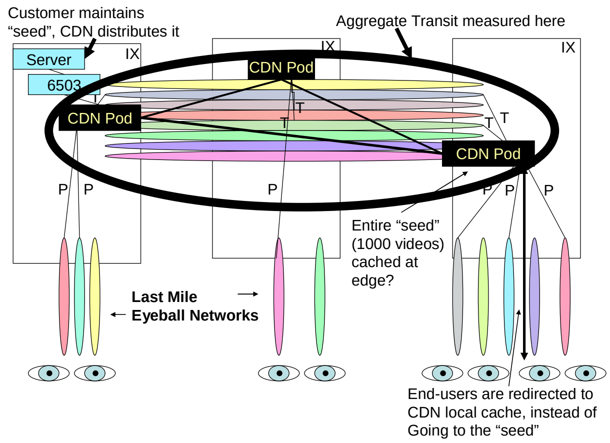

Average Load			40,000	Mbps	
95th Percentile Load		4	160,000	Mbps	
Peak Load		6.6	264,000	Mbps	
Model 2C - Content Delivery Network for Video Distribution					
					monthly
Transit Fee	160,000	mbps@	$13	perMbps	$2,080,000
Colo	1	rack@	$1,500	perRack	$1,500
Network Equip	1	6503	$30,000	3yrAmort	$833
Servers	1		$4,000	3yrAmort	$111
Maintenance				15%	$367
Staff	0.5		$180,000		$7,500
Total					$2,090,311
# videos downloaded per month					8,640,000
Cost per video downloaded					$0.24

Since so much of the cost to the VSP is in transit, some VSPs will consider taking the distribution in-house, which is discussed next.

Model 3: Blended Transit and Peering for the Distribution of Video Content (aka Di-It-Yourself CDN)

Business Model Premise: Operation of the Internet distribution is seen as strategic to the VSP:

a) End-user experience is mission-critical so outsourcing the end user experience to a transit provider or CDN puts the VSP at risk.
b) The VSP has visibility into what video are being released, which ones are likely to be hot and which ones don't require special infrastructure adjustments.
c) Internet Video distribution is so new that the VSP prefers control. This is a strategic focus for the VSP: ensuring reliability, scalability, through the constant monitoring and evolution of the infrastructure to ensure the end user experiences during these early phases of Internet Video Distribution.
d) The traditional CDN may be ill-suited to distribute very large video objects, therefore we have to do it yourself.

One reviewer points out a few additional motivations for some companies to build their own CDN:

e) Content producers (Movie Studios for example) may want to bypass distributors because they can, whereas they could not bypass distributors before the Internet.
f) Some profiles of web objects are more difficult for CDNs to handle well: large volume of accesses of large objects and small volumes of accesses of small objects don't fit the CDN model as they rarely are served out of CDN edge cache.
g) Empire Building – some DIY CDNs are built not out of necessity but by staff need to build their own importance within the organization, aka build and operate the infrastructure 'better'.

<u>Definition</u>: Peering is a business relationship whereby two entities reciprocally exchange access to each others networks.

Is Peering for Content? Many of the largest Content Providers in the world are peering now. There are three primary motivations for peering. Peering provides the lowest latency path between two points. Peering bypasses the transit provider providing cost savings for both parties. Peering also provides greater control over the end user experience. Yahoo!, Microsoft, and Google for example have built out substantial networks and are peering at exchange points around the world. Yahoo! currently has over 640 peering sessions and a multiple- OC-192 (10Gbps) global backbone to distribute its content itself to peering and transit relationships[19].

Peering is not Free. There is a recurring flat monthly fee associated with public peering (switch port fees, collocation expenses, etc.) and private peering (cross connect or circuit fees, collocation, etc.). To compare peering against transit, one compares the unit cost of transit (in Mbps) against the unit cost of peering (in Mbps). Since peering costs are fixed monthly, the unit cost of peering varies based on the amount of traffic peered during the month.

The Peering Break Even Point (where the unit cost of transit equals the unit cost of Peering) is calculated using a lot of assumptions but in 2006 approaches

[19] Brokaw Price (Yahoo!) presentation, Sydney Peering Forum, 2005.

300Mbps[20]. That is, the cost of peering is completely covered by the cost savings of peering at least 300Mbps of traffic that would otherwise traverse a metered transit service. Fortunately, as we will see, video traffic is so large that this breakeven point is met in even the light load model.

Con: Requires the Content Provider become a Network Operator. This means a content provider interested in peering needs a 24/7 NOC, networking expertise on staff, etc. a seasoned and talented Peering Coordinator is required to obtain peering. These requirements can be expensive and away from the core competence of a video service provider. On the other hand, some of these things may be already be in place, perhaps as required for some other operations activities.

We will make a number of simplifying assumptions for all 3 Transit/Peering models:

Three IX Deployments. Jeffrey Papen points out that three sites are sufficient to meet most geographic diversity peering prerequisites for the eyeball networks that would receive video traffic. More sites than that he argues would provide diminishing returns; the costs may exceed the incremental peers picked up.

No Backbone. These models assume that no backbone is used for the distributed VSP implementation based on Jim Gray (Microsoft) assertion that the most cost effective way of distributing large amounts of content is to UPS overnight disk drives[21]. So, the servers and routers are shipped pre- configured with videos. It further assumes that updates (new releases) are infrequent and are distributed over a tunnel using a commodity transit services.

Each deployment independently can handle load. For simplicity, we will use three copies of the same equipment, ignoring efficiencies of distributing the load across all IXes. (If we assumed a uniform traffic load distribution across all IXes, we could for example chose fewer servers and smaller routers.) So Model 1A, 1B and 1C will each be deployed to three IXes across the U.S. resulting in three times the deployment cost.

[20] For example, 300Mbps@$25/Mbps=a $7500 gigE peering port. Any peering beyond 300Mbps saves money.

[21] Discussion on YouTube and the "Google Test" http://www.venturebeat.com/contributors/2006/09/21/the-google-test/

Purchase transit, peer in each location with an open peering policy and expert Peering Coordinator(s) as part of the network operations group**.** Active expert peering evangelism is required given the following five hurdles:

1) BGP is not perceived as a content providers expertise, so ISPs may be skeptical of their network competence. Since they will be woken up in the middle of the night, the ISP Peering Coordinator may not be receptive to peering discussions with perceived novices.
2) Peering requests may be seen as a revenue opportunity by the ISP approached. Sales Revenue is preferred by the ISPs, particularly given the large amount of traffic the ISP is expected to receive from the VSP.
3) Personality Clashes have prevented peering from going forward where it otherwise is in both companies best interests. A skilled Peering Coordinator knows the venues, the personalities, the approaches that will work and already has a rolodex with the right contacts to initiate discussions.
4) The potential load is so large that the peer may need to rapidly upgrade peering infrastructure. Without revenues to support this activity, there may be reluctance for the ISP to peer.
5) The peer may see a business clash in peering if their company sells videos to their customers also. They may not want to encourage or enable a far away competitor that will use their bandwidth to compete with them (aka the Net Neutrality Debate).

Model 3A: Light Load

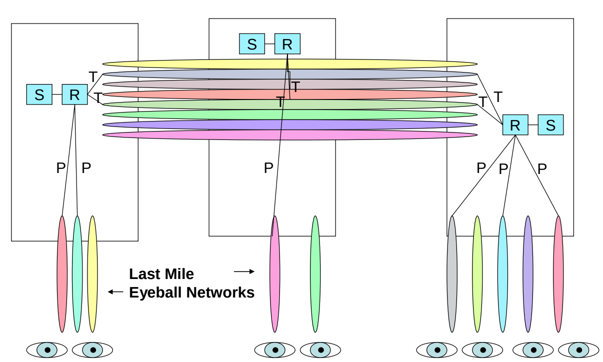

Here we will assume that the VSP is distributed to 3 IXes that are all reasonably well populated so the VSP can peer away 25% of their traffic. Jeffrey Papen (Peak Web Consulting) shared that his experience shows that 12%-18% of traffic can generally be peered in a single location. Barrett Lyon (BitGravity) believes the number to be closer to 40%, and with multiple sites connected by a backbone, Barrett claims a VSP should be able to peer away 60% of their traffic. We will assume that a conservative 25% of the VSP traffic is peered across the IXes resulting in the following financial analysis.

Average Load			400	Mbps	
95th Percentile Load		4	1,600	Mbps	
Peak Load		6.6	2,640	Mbps	
Model 3A - Blended Transit and Peering for Video distribution					
3 site	25%	peering			monthly
Transit Fee	1,200	mbps@	$25	perMbps	$30,000
Colo	3	rack@	$3,000	rack+port	$9,000
Network Equip	3	6509	$90,000	3yrAmort	$7,500
AggregationSwi	1		$10,000	3yrAmort	$278
Servers	3		$4,000	3yrAmort	$333
Maintenance				15%	$1,217
Staff	0.75		$180,000		$11,250
Total					$59,578
# videos downloaded per month					86,400
Cost per video downloaded					$0.69

Here we are assuming that the per-IX expense is $3000 per rack inclusive of a 1G peering ports. With this model, 25% of the traffic is peered across the IXes decreasing the transit expense but with a corresponding increase in staffing and infrastructure. The benefits of peering are more pronounced with larger scaled deployments discussed next

Model 3B: Medium Load

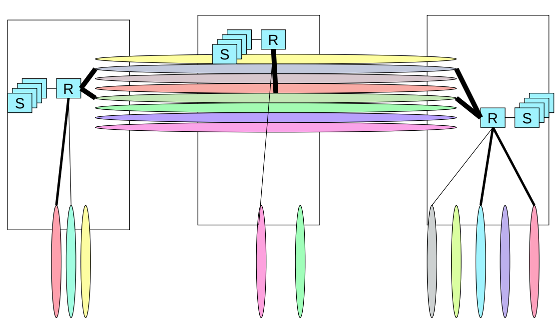

With this model we purchase the larger 10G peering port at greater expense, but at these higher loads, the monthly transit fees still dwarfs all other expenses. We have effectively however reduced the most expensive component by 25% resulting in an attractive cost per video.

Average Load			4,000	Mbps	
95th Percentile Load		4	16,000	Mbps	
Peak Load		6.6	26,400	Mbps	
Model 3B - Blended Transit and Peering for Video distribution					
3 site	25%	peering			monthly
Transit Fee	12,000	mbps@	$18	perMbps	$216,000
Colo	3	rack@	$10,000	rack+port	$30,000
Network Equip	3	6509	$90,000	3yrAmort	$7,500
AggregationSwi	2		$10,000	3yrAmort	$556
Servers	12		$4,000	3yrAmort	$1,333
Maintenance				15%	$1,408
Staff	1		$180,000		$15,000
Total					$271,797
# videos downloaded per month					864,000
Cost per video downloaded					$0.31

Model 3C: Large Load

The large load model presents significant hurdles from a political perspective.

First we are now offloading ten's of Gbps to peers who most likely will need to spend capital to upgrade their equipment to handle the additional peering.

Second, the largest cable company and telephone company peers offer competing video services, so may not be interested in making it inexpensive for VSPs to get to their eyeballs.

Finally, there is some evidence that these eyeball heavy networks operate network devices to throttle back heavy network users. The motivation for this deployment was initially the excessive peer-2-peer file sharing (copyright violating typically) user load. The point is that the scale by itself will attract attention and generate hurdles for the VSP deploying this much video across the commodity Internet.

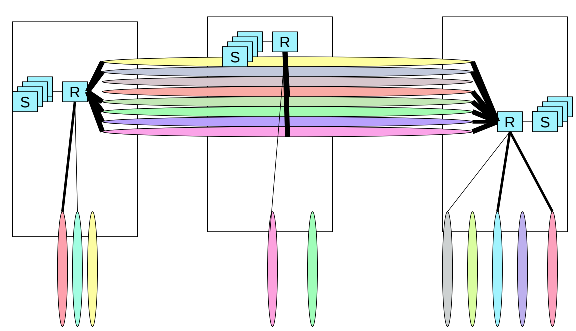

With this model we duplicate Model 1C across all IXes deploying a large number of servers. We assume the large number of servers will lead to a reduced average price per rack inclusive of peering ports. Again, the numbers in the model can be adjusted.

Average Load				40,000	Mbps	
95th Percentile Load			4	160,000	Mbps	
Peak Load			6.6	264,000	Mbps	
Model 3C - Blended Transit and Peering for Video distribution						
3 site	25%	peering				monthly
Transit Fee	120,000	mbps@		$10	perMbps	$1,200,000
Colo	42	rack@		$2,000	rack+port	$84,000
Network Equip	12		6509	$150,000	3yrAmort	$50,000
AggregationSw	42			$10,000	3yrAmort	$3,889
Servers	792			$4,000	3yrAmort	$88,000
Maintenance					15%	$21,283
Staff	3			$180,000		$45,000
Total						$1,492,172
# videos downloaded per month						8,640,000
Cost per video downloaded						$0.17

By far the most surprising side effect of this research was the discovery of the peer-to-peer economies for video distribution, discussed next.

Model 4: Peer to Peer Distribution of Video Content

Business Model Premise: The current Internet Service Providers and CDNs at the core can not effectively handle the video load across a single or even multiple locations:

a) Backbone, peering interconnects, and the hundreds of thousands of routers deployed can not handle the load or performance (latency, jitter, etc.) for video.

b) The leaf nodes (i.e. Grandma's PC) in aggregate have the cycles and network capacity, if shared, to distribute popular video content today. Jeffrey Payne (Grid Networks) claims that less than 1% of the aggregate last mile content is utilized today.

c) Popular content can be chopped up into small chunks such that many downloaders become sources, and topologically close downloaders will prefer the topologically close sources. This 'swarmcasting' requires only a source 'seed', and a lookup mechanism for the first downloaders to find the seed, and then to direct future downloaders to topologically closer sources. This has proven a viable technology even if the business models have been flawed in the past.

Peer-to-peer is compelling to the VSP since the load on the VSP infrastructure and therefore cost is minimal. With this low cost of entry, it should not be surprising that the author has discovered the most innovation occurring in the peer-2-peer space.

Peer-to-peer does however count on the good graces of the last mile operators allowing their customers to fill their pipes with this shared video content.

Cringeley suggests a Utopian peer-to-peer approach whereby 2.5% of the "Desperate Housewives" audience of 10,000,000 distributed randomly across the net could be given a feature like fast forward over commercials if they allow their machine and DSL line to be used in a peer-2-peer reflector mode. Then the rest of the 10,000,000 households would pull the pieces of the video files from these randomly distributed reflectors.

This spreading of the load is done with the expense of traffic distribution spread across the interested user base network infrastructure. The cost to the content provider here is only the cost of distributing the content to the 256,000 reflectors and perhaps some control plane overhead to manage this.

There will always be at least an initial spike in transit fees while the "seed" content is being distributed to "peers" in the peer-to-peer desktop sense. The transit load continues until the content is distributed to desktops topologically closer to the end-users requesting the file. Many p2p users today move completed downloads out of their shared directories making their sources unavailable to others. In response, many peer-2-peer networks use a give-to-get throttling tracker mechanism to encourage those people to keep their PCs seeding the content. As long as some desktops remain online and the source remains available, peer-2-peer remains a powerful and inexpensive alternative for distributing large scale video traffic.

Model 4A: Light Load

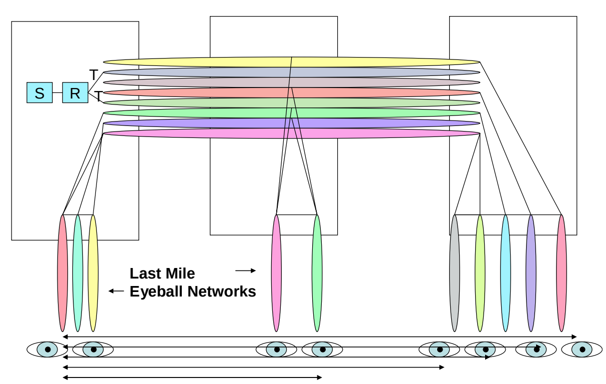

Average Load			400	Mbps	
95th Percentile Load		4	1,600	Mbps	
Peak Load		6.6	2,640	Mbps	
Model 4A - Peer-to-Peer Network for Video Distribution					
single-site stormcasting					monthly
Transit Fee	100	mbps@	$50	perMbps	$5,000
Colo	1	rack@	$1,500	perRack	$1,500
Network Equip	1	6503	$30,000	3yrAmort	$833
Servers	1		$4,000	3yrAmort	$111
Maintenance				15%	$367
Staff	0.5		$180,000		$7,500
Total					$15,311
# videos downloaded per month					86,400
Cost per video downloaded					$0.18

Note that the cost to the VSP is very low, leading to a very cost per video. Reviewers observed that the cost of distribution is being simply *shifted* to the last mile providers.

Model 4B: Medium Load

The cost effectiveness of peer-to-peer is even more pronounced as we evaluate the more popular video content. Once the seed gets out from the origin, there are no additional costs for popular (seeded) content.

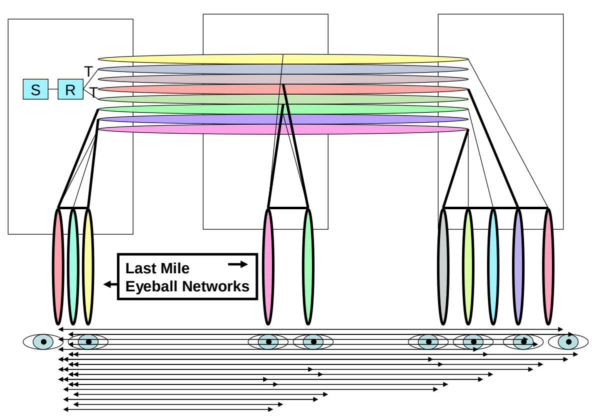

Average Load			4,000	Mbps	
95th Percentile Load		4	16,000	Mbps	
Peak Load		6.6	26,400	Mbps	
Model 4B - Peer-to-Peer Network for Video Distribution					
single-site stormcasting					monthly
Transit Fee	100	mbps@	$50	perMbps	$5,000
Colo	1	rack@	$1,500	perRack	$1,500
Network Equip	1	6503	$30,000	3yrAmort	$833
Servers	1		$4,000	3yrAmort	$111
Maintenance				15%	$367
Staff	0.5		$180,000		$7,500
Total					$15,311
# videos downloaded per month					864,000
Cost per video downloaded					$0.0177

These cost characteristics more closely mirror the broadcast economics: There is a one time cost of broadcasting, and the more popular the video content, the more viewers and the more cost effective the (seeded) distribution.

Model 4C: Large Load

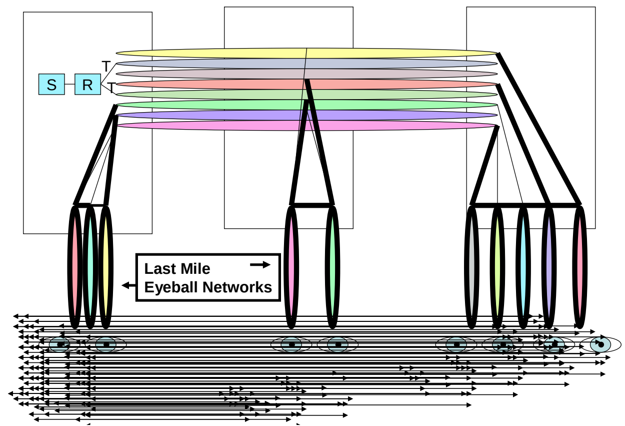

Average Load		Mbps		Mbps	
95th Percentile Load		4	160,000	Mbps	
Peak Load		6.6	264,000	Mbps	
Model 4C - Peer-to-Peer Network for Video Distribution					
single-site stormcasting					monthly
Transit Fee	100	mbps@	$50	perMbps	$5,000
Colo	1	rack@	$1,500	perRack	$1,500
Network Equip	1	6503	$30,000	3yrAmort	$833
Servers	1		$4,000	3yrAmort	$111
Maintenance				15%	$367
Staff	0.5		$180,000		$7,500
Total					$15,311
# videos downloaded per month					8,640,000
Cost per video downloaded					$0.0018

The cost of distribution is VERY LOW!

However, consider a few observations from the field:

1) You are counting on the good graces of the last mile providers for your business continuity. Specifically, there is no contractual arrangement requiring that peer-to-peer video be carried at all on their network.

2) Last mile bandwidth in the U.S. is asymmetric. Upstream bandwidth is typically a fraction of download bandwidth, so there may not be sufficient upstream bandwidth on any seed machines to handle all the spot event downloaders. Once others have downloaded the seed however, the swarming does distribute the video traffic pretty effectively, preferring the low latency higher bandwidth seeds over the slower ones.

3) Last mile providers have deployed traffic shaping, peer2peer mitigation devices to throttle back the load[22]. One has to consider the effects of these actions and future actions on the delivery of services.

On the other hand,

1) Peer-to-peer has a proven track record for distributing content over large distances for very little cost to the content providers, and

2) Peer-to-peer is already in the early stages of adoption by the movie studios[23].

3) PPLive service is demonstrating today the ability to stream even live events using a distributed p2p model[24].

Summary and Implications

We discussed video distribution methods with over one hundred companies and documented the current thinking on video distribution over the Internet. We developed a model for each of four distribution techniques across three load models and highlighted the cost per video for using each model.

Models	A:10 videos	B: 100	C: 1000
1: Transit	1A: $0.60	1B: $0.36	1C: $0.20
2: CDN	2A: $0.77	2B: $0.44	2C: $0.24
3: Hybrid	3A: $0.69	3B: $0.31	3C: $0.17
4: P2P	4A:$0.18	4B: $0.0177	4C: $0.0018

From the data and analysis of the different network distribution schemes for video, given all the assumptions listed, we observe:

1) Commodity transit for large scale video traffic delivery from a single location does not work in 2007. The cascading failures span the IX power and space availability, the ISP peering and backbone infrastructure, the lack of high bandwidth standards for 40Gbps and 100Gbps to support this load, and the last mile infrastructure.

2) On a unit basis, video distribution clearly demonstrates economies of scale,

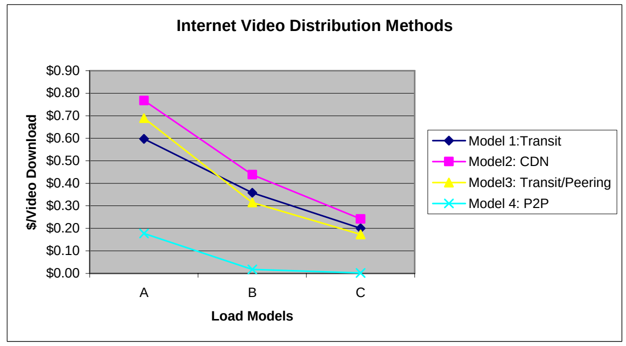

3) Transit expense is by far the largest component of the Internet Video distribution expense,

4) Transit with Peering approach requires expertise but scales well to deliver videos,

5) Peer2Peer distribution for Internet Videos is by far the least expensive to the Video Service Provider, but it accomplishes this by moving almost all of the distribution costs to the last mile providers.

Implications:

1) There appears to be a market failure that makes it unattractive for the largest ISPs in the U.S. to deploy infrastructure to service the pending video traffic.

2) We've focused on the demand side, but there is a serious supply side problem in the Internet today:

- The Tier 1 ISPs are turning away 10G transit service requests citing lack of 40G/100G backbone core and peering interface capacity, lack of 40G/100G local loops with service into the IXes, limited IX Power capacity, last mile infrastructure limitations and oversubscription, etc.). When these issues are overcome, each Tier 1 will need to do a forklift upgrade (challenging transition) to replace hundreds of large (Cisco 12000/Juniper T320 class) routers with equipment capable of handling bundles of 10G video traffic. Of course the peers also need to be able to accept this traffic so they will require upgrades as well. All of this takes time.

3) As a result of the above, the U.S. ecosystem

[22] Harald Willison (formerly a network engineer for Adelphia) shared that virtually all of the cable companies have adopted deep packet inspection devices to throttle back p2p, DOS, etc. traffic based on traffic signatures.

[23] "BitTorrent goes Hollywood", BusinessWeek http://www.businessweek.com/technology/content/may2006/tc20060508_693082.htm

[24] PPLive website listings: http://www.pplive.com/en/index.html

may experience a period of shortened transit supply and a resulting higher transit prices. We are already hearing about multi-10G transit orders being turned away.

4) Some reviewers make the case that with 80-90% of all Internet traffic being video over the next couple years, and given the large peak characteristics described in this paper, that a new pricing model may be emerge; ISPs may no longer absorb the top 5% for free as in the current 95th percentile system. Others believe that the load on the peering and transit interconnects will be the limiting factor and therefore wholesale pricing quotes will more closely mimic the telephone system, with a traffic profile pricing scheme, so an ISP can determine the impact of the customer video traffic on their interconnects and network infrastructure before accepting the business.

Further Research /Discussion Topics

1. We talk about the waves of streaming video content but then discuss the methods of distribution using models for static 1.5GB movies. Apply the same math to streaming services. What is different?
2. Is there a clear break even point between model 2 and model 3 when it makes sense to build your own CDN?
3. Last Mile broadband Internet Access oversubscription ratios, some estimate to be 1000:1. In India you purchase Internet access and are told the oversubscription rate and charged accordingly.
4. What happens to the Internet charging schemes when 90% of all Internet traffic has 6.6:1 peak-to-mean ratio when the ISP community has priced and modeled using the 2:1 traditional Internet short HTML traffic flows?
5. Who wins and who loses with Internet Video traffic? Cable companies? CDNs? Traditional Movie and Television studios? New entrants into these markets?
6. This white paper talks about very expensive equipment required to service the "peak" demand and the expense of servicing a high "95th" percentile. What are the dimensions around the problem if you assume time-shifting can be accomplished; distributing the load to off-hours somehow, perhaps loading over a long period of time some clients or caches, or a set top box. Wouldn't that make the 95th percentile closer to the mean?
7. This white paper ignores the fact that some large percentage of the traffic is outbound (from video source to the Internet). Is there something that can be done with the inbound capacity?
8. We ignore the motivations for the content owners. We also ignore the security aspects and Digital Rights Management (DRM) issues. Might movie studios have a preference over the distribution method based on past experience with these methods?
9. At what point does the P2P model reach capacity in the last mile? If we assume no packet mitigation is occurring, will the uplink or downlink saturate first?

Acknowledgement

The white papers are constructed upon many hours of discussions with key members in the community and therefore have been made freely available to the community. I would like to thank the following people specifically who provided good ideas, insights, suggestions, or in some cases just allowed me to walk through the white paper with them allowing me to flush out the logic: Vish Yelsangikar (NetFlix), Peter Harrison (NetFlix), Aaron Weintraub (Cogent), Jon Nistor (TorIX), Barrett Lyon (BitGravity), Dave Knight (ISC), Aaron Hughes (Caridien), David Filo (Yahoo!), Jim Goetz (Sequoia Capital), Jason Holloway (DoveTail), Matt Peterson, Richard Steenbergen (nLayer), Lane Patterson (Equinix), Eric Schwartz (Equinix), Pete Ferris (Equinix), David Cheriton (Sun), Andy Bechtolsheim (Sun), Jeffrey Papen (Peak Web Consulting), KC Broberg (Rackable), Henk Goosen (Sun), Geoffrey Noer (Rackable), Jeff Turner (InterStream/nuMetra), Vab Goel (NorWest Venture Partners), Ken Hubbard (Quad), Matt Wood (Equinix), Gary Baldus (LimeLight Networks), Sylvie Laperriere (VSNL), Ted Seely (Sprint), Bryan Sutterfield (Cox), Patrick Gilmore (Akamai), Phil Thomas (Quad), Doug Wilson (Microsoft), Christian Nielsen (Microsoft), Harald Willison (Covad), Davin Overland, Douglas Galbi (FCC), Ren Provo (AT&T), Susan Martens (AT&T), Richard Clarke (AT&T), the Google Network Engineering and Network Architecture team, James Blessing (Entanet), Sean Donelan (Akamai).

About the Author

Mr. Norton's title is Co-Founder and Chief Technical Liaison for Equinix. In his current role, Mr. Norton focuses on research on large-scale interconnection and peering research, and in particular scaling Internet operations using optical networking. He has published and presented his research white papers ("Interconnections Strategies for ISPs", "Internet Service Providers and Peering", "A Business Case for Peering", "The Art of Peering: The Peering Playbook", "The Peering Simulation Game", "Do ATM-based Internet Exchange Points Make Sense Anymore?") in a variety of international operations and research forums.

From October 1987 to September 1998, Mr. Norton served in a variety of staff and managerial roles at Merit Network, Inc., including directing national and international network research and operations activities, and chairing the North American Network Operators Group (NANOG) Internet industry forum. Mr. Norton received a B.A. in computer Science and an M.B.A. from the Michigan Business School, and has been an active member of the Internet Engineering Task Force for the past 15 years.

About the White Paper Series - Network Operations Documents (NODs)

The Network Operations Documents (NODs) identify a critical but undocumented area of Internet Operations. We research that area with the Operations Community to document the area definitions, motivations, strategies, etc. The initial drafts are reviewed in "walk throughs", where Internet Operators provide their views, their data points, their criticisms, and their experience. These are credited in the Acknowledgements section and footnoted where appropriate for the next walk throughs. After enough walk throughs, the responses tend to migrate from constructive feedback to nods of acceptance, at which time a draft to is made available to the broader Internet Operations community. The papers are never "done" but rather are considered living documents, evolving with input from the community, hopefully reflecting the current practices in the previously undocumented area. Here are the NODs available from the author:

1. **<u>Interconnection Strategies for ISPs</u>** documents two dominant methods ISPs use to interconnect their networks. Over 200 ISPs helped create this white paper to identify when Internet Exchange Points make sense and the Direct Circuit interconnect method makes sense. Financial Models included in the paper quantify the tradeoffs between these two methods. All relevant data points are footnoted as to source.

2. **<u>Internet Service Providers and Peering</u>** answers the questions: "What is Peering and Transit? What are the motivations for Peering? What is the ISP Peering Coordinators Process for obtaining peering? What are criteria for IX selection?"

3. **<u>A Business Case for Peering</u>** builds upon the previous white papers but focuses on the questions important to the Chief Financial Officer: "When does Peering make sense from a financial standpoint? When do all the costs of Peering get completely offset by the cost savings?"

4. **<u>The Art of Peering: The Peering Playbook</u>** builds on the previous white papers by asking the Peering Coordinators to share the "Tricks of the Trade", methods of getting peering where otherwise they might not be able to get peering. These 20 tactics range from the straight forward to the obscure, from the clever to the borderline unethical. Nonetheless, Peering Coordinators might be interested in field-proven effective ways of obtaining peering in this highly controversial white paper.

5. **<u>The Peering Simulation Game</u>** finishes up my half day Peering Tutorial by engaging the audience in the role of the Peering Coordinator. Each ISP in turn rolls the dice, expands their network, collects revenue for each square of customer traffic, and pays transit fees to their upstream ISP. They quickly learn that if they peer with each other, the costs of traffic exchange are much less, but they need to negotiate how to cover the costs of the interconnect. ISP Peering coordinators have commented on how close the peering simulation game is to reality in terms of the dialog that takes place.

6. **<u>Do ATM-based Internet Exchange Points Make Sense Anymore?</u>** Applies the "Business Case for Peering" financial models to ATM and Ethernet-based IXes using current market prices for transit, transport, and IX Peering Costs.

7. **The Evolution of the U.S. Peering Ecosystem**, introduces and focuses on several fundamental changes in the Peering Ecosystem spurred by several events following the telecom collapse of 1999/2000.

8. **The Art of Peering: The IX Playbook** follows the same tact as The Peering Playbook; we first introduce the framework theory of how and why IXes are valuable from an economic perspective. We then enumerate about a dozen tactics IXes use to get over the "Start Up Hump", to build a strong critical mass of participants, and finally, defense tactics to maintain that population. (To be released at a future date.)

9. **The Asia Pacific Peering Guidebook** follows the "Evolution of the U.S. Peering Ecosystem" by exploring the Asia Internet environment from a peering perspective. What did Peering Coordinators find as counter-intuitive? What are the challenges peering in Tokyo, Hong Kong, Sydney and Singapore? This paper provides insights into these and related questions.

10. **The Great (Public vs. Private) Peering Debate** models the tradeoffs of using a large public peering port (10Gigabit Ethernet) versus lots of private cross connects. Both models are widely deployed and defended religiously.

11. **The Folly of Peering Ratios as Peering Discriminators** documents the discussions surrounding the peering debate at NANOG surrounding the peering ratio requirement many ISPs have in their peering policy. It lists and discusses in brief the strongest arguments on both sides of the debate.

Appendix A - References

The following resources were used to research this white paper.

Economics of Video and the Internet. http://www.videotechnology.com/economics_of_video.htm

http://www.ourmedia.org/

Open Media Network: 4M hours of TV and video traffic http://www.omn.org/

This week's top 50 videos: http://www.putfile.com/weekvideos

Barbie Girl: http://video.google.com/videoplay?docid=-5426933764767977068

Numa Numa Video: http://video.google.com/videoplay?docid=-6377855743675143177&q=original+numa+numa

Martial Arts Monkey http://youtube.com/watch?v=LZdGb8sW9us

WhereTheHellIsMatt: http://youtube.com/watch?v=qZSTfN3BJ_Y&search=man%20dancing

The Internet is for Porn http://video.google.com/videoplay?docid=5430343841227974645

Great Soccer Skills http://video.google.com/videoplay?docid=-6247272780250291395

Shakira: http://youtube.com/watch?v=8gUtHxtkfqQ&search=shakira

Shakira Spook: http://video.google.com/videoplay?docid=8718728501056290731

Video Shows

www.google.com/Top/Computers/Internet/Broadcasting/Video_Shows/

Internet Video Magazine

http://www.internetvideomag.com/

http://www.internetvideomag.com/BestVideosoftheWeek.htm

http://www.internetvideomag.com/BestVideoWebSitesoftheWeek.htm

Moms online and Off the Record: http://www.mommetv.com/

Food Portion Control: http://home.about.com/z/cg/vp.htm?ch=health&l=health/v/5&ap=1','18/1FQ/U8'

Porcupine Tree live in Germany: http://video.google.com/videoplay?docid=138242861598716315&q=concert

Chainsaw Artist: http://videos.caught-on-video.com/category/Amazing+Skilled/0/a040ca5d-5dea-4016-834c-980c0002db42.htm

Caught on Video in this medical video footage which shows from conception of sperm and egg to the first four weeks of new life. : http://videos.caught-on-video.com/hottestvideos/0/28c3899b-2875-4976-898f-980f00135c24.htm

Content Delivery Networks (CDNs): How They Work and How They're Used? http://www.campus-technology.com/techtalks/trans/010920content.asp

Content Networks by Christophe Deleuze http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_7-2/content_networks.html

Curvey of Content Delivery Networks (CDNs) http://cgi.di.uoa.gr/~grad0377/**cdns**urvey.pdf

CDNs are not just for content anymore Network World 01/14/2002 http://www.networkworld.com/news/2002/0114specialfocus.html

Content Delivery and Distribution Services: http://www.web-caching.com/cdns.html

Growing beyond a single CDN: Peering content networks Nov 14 2001 http://infocus.telephonyonline.com/ar/telecom_growing_beyond_single/index.htm

A Practical Guide to Streaming Media http://tnc2001.terena.nl/proceedings/Slides**MatthewListe**.pdf

First Report on Future NGI Control Architectures and NGI Services http://eurongi.enst.fr/archive/127/JRA141.pdf

http://www.campus-technology.com/techtalks/events/010920content.asp

Peer-to-Peer Content Distribution

http://www.cs.cmu.edu/~kunwadee/research/p2p/

"Mapping the Gnutella Network: Macroscopic Properties of Large Scale Peer-to-Peer Systems" by Matei Ripeanu, Ian Foster

Load Sharing using open source SW on cheap hardware: http://www.inlab.de/balanceng/example3.html

Back to the future of video compression: MPEG workshop on future directions in video compression. 20 April, 2005, Busan, Korea http://m7itb.nist.gov/**video**WS1/Ebrahimi.pdf

Notes on video compression and compression standards http://www.kan-ed.org/marratech/video-basics.html

http://www.astaskywire.com/news/200503/20050301.asp

Taking Video Service to the next level: www.bigbandnet.com/documents/Service_Reliability__Assurance.pdf

"Network Planning for Rich media Content and Application Providers" by Jeffrey Papen, Peak Web Consulting. Jeffrey@peakwebconsulting.com

<it would be good to break this down to $/video download, and add in CDN into the graph>

Amazon Elastic Compute Cloud (Amazon EC2) - *Limited Beta*

http://www.amazon.com/gp/browse.html?node=201590011

Really good academic writeup of the detailed options surrounding vide delivery servers. Streaming Video over the Internet: Approaches and Directions www.wu.ece.ufl.edu/mypapers/streamingVideo_camera.pdf

The AMS-IX Video Working Group meeting minutes in English describes the solutions for handling the exceptionally large video traffic, predicts a "collapse by May 2007" if they don't offload somehow and deal with the video traffic. https://www.ams-ix.net/video-wg/Minutes_StreamingWorkshop_170506.pdf https://www.ams-ix.net/video-wg/Minutes_300806EN.pdf

From Jeff Turner: you can cite estimates for YouTube bandwidth costs from some of the following links: http://www.datacenterknowledge.com/archives/2006/Sep/12/youtube_gets_bandwidth_boost_from_level_3.html – non-authoritatively references Forbes article (below) http://digg.com/tech_news/YouTube_Gets_Bandwidth_Boost_from_Level_3 - for the $1M per month figure in transit for YouTube... There is another blog entry at: http://willy.boerland.com/myblog/youtube_bandwidth_usage_25_petabytes_per_month - also stating a higher number. http://www.forbes.com/home/intelligentinfrastructure/2006/04/27/video-youtube-myspace_cx_df_0428video.html - the original business article which pointed this out... http://blog.forret.com/2006/05/youtube-bandwidth-terabytes-per-day/ - Terabytes per day estimates...