

The Inescapable Inevitability of Convergence (Unless You "Help") Converging Campus Technologies: Evolution or Intelligent Re-Design?

NWACC 2006 Annual Conference (Portland, Oregon, June 9th, 2006)
(Cut/paste by Wayne Caswell from: <http://www.uoregon.edu/~joe/convergence/nwacc2006.ppt>)

Joe St Sauver, Ph.D. (joe@uoregon.edu)
Director, User Services and Network Applications
University of Oregon Computing Center
<http://www.uoregon.edu/~joe/convergence/>

Introduction

Welcome to the last session for this year's NWACC conference. I'd like to thank Marty for the opportunity to present this session, and I hope you've all enjoyed the rest of this year's meeting as much as I have.

It's rare for me to have a talk theme mesh so well with the overall theme of an event, or so closely with the major news events of the day, but I think that may be fortunate since I'm all that's between you and lunch (or a few hours of exploring Portland before heading home). I'll try to make sure you get your money's worth for the time you're investing.

Format of This Session/Handout

This session will be a half hour introduction/ overview followed by up to an hour for discussion.

While I'll begin by presenting one perspective on convergence, mine,* I hope you'll feel free to share your perspective during the discussion period, particularly if you see things differently than I do.

A note about this handout: I tend to cover a lot of material, so to help me stay on track, to facilitate later review by folks not here with us today, and to accommodate attendees who may be hearing impaired, I've scripted these slides in some detail (think of them as "closed captioning"). Not a spontaneous feeling, but I hope you'll indulge me.

* Disclaimer: And in fact, all opinions in this talk are strictly my own.

So What Is "Convergence"? It can mean different things.

- In a computing context, convergence might be taken to mean the near-ubiquitous adoption of a particular technology or product, such as x86 Intel/AMD CPUs (even Apple's using 'em)
- In a peripheral context, convergence might be associated with the development of multifunction devices (e.g., printer/scanner/fax/digital senders).
- In the network context, convergence is often taken to mean the consolidation of separate networks into a single Internet Protocol (IP)-based network.

We'll focus on this last type of convergence today.

The Traditional Approach to Delivering Voice, Video, and Data...

- Voice goes over the copper phone infrastructure (or via dedicated cellular infrastructure)
- Video goes over over a dedicated coaxial or fiber cable TV infrastructure, dedicated ISDN lines, minidish satellite, or broadcast TV, and
- Data goes over a dedicated data network.

All three redundant networks often run side-by-side at low levels of utilization and at considerable (potentially avoidable) expense. Combining all three of those onto one converged network is often called a "triple play" strategy or running a "packet-based multi-service network."

If We Just Share Some of the Physical Infrastructure, Are We "Converged?"

Occasionally you may run into situations where common physical infrastructure serves multiple purposes. For example, you might see voice and DSL service over telco copper, or cable TV and cable modem service over Cable Company coax.

In my opinion, this is not a true "converged" network – legacy services are still being delivered via legacy analog channels.

In a true converged network, all the services are delivered as interleaved IP packet traffic, getting encoded at their origin and decoded at their destination as may be necessary.

Some Advantages of Running Converged:

- Simplify your infrastructure and reduce CapEx by eliminating redundant networks; save money.
- Reduce dedicated specialized staff requirements and ongoing operational expenses; save money.
- Simplify local provisioning (just pull Ethernet to a location, no need to also worry about dedicated copper for voice or coax for video); save money.
- Increase your service footprint (if wanted, every Ethernet jack could also have voice & video) and your flexibility (zero turn up time for new installs or for service moves); save money.
- More potential features for carriers to sell.

The Worries

So why aren't all networks converged today? There are some potential worries, including:

- "Quality" (Jitter/dropped packets due to commodity network data traffic potentially interfering with sensitive voice/video traffic, leading to poor sound quality or video artifacts)
- "Reliability" (Ssay what you will, the traditional phone system has been engineered to be very, very reliable, including during emergencies; users may need to be shown that a converged network can be as reliable)
- "Security" (e.g., People may believe that POTS service is inherently more secure than VoIP. It's not.)
- "Ergonomics" (If all you've ever seen is VoIP that requires a headset to avoid echo-related issues, you would not be willing to give up your traditional telephone.)
- "Costs" (This line of worry runs along the lines of "there may be some material unanticipated cost that will pragmatically wipe our any savings associated with convergence," including the classic "I'll need to replace my entire network" worry.)
- "Risk Aversion" (Succinctly put, "What we have currently works, and while it isn't perfect, I'm not going to get fired for just continuing to do what everyone else is doing.")
- "Interoperability/Standards Status" ("I'll just wait a little longer until the standards, uh, solidify...")
- "Vendor Attempts at Product Differentiation" (If you make network hardware, you may be tempted to promote some feature your product (and only your product) supports, even if that means overemphasizing the magnitude and prevalence of some rarely seen problem, or hinders interoperability/standardization efforts.)
- "The Farrier Problem" ("All I know how to do is shoe horses, thus I fear the automobile because it has the potential to make me obsolete.")
- "Regulatory Compliance Issues"

The Converged Network Advocate's Rejoinder: "Try It, It Just Works"

Demonstration of successful convergence can be a powerful persuasive tool, but proof-by- example is only persuasive, not conclusive:

- How do you know that the success seen in a trial will replicate and scale ubiquitously?
- What if my pilot project works great, but my production roll out crashes and burns? What's my failover / remediation option then?
- Is there some sort of technical "insurance" I can buy that will keep the demons at bay?

There's an almost irresistible urge to doubt or complicate an elegantly simple solution – it just must be too good to be true.

Convergence in Managed/Enterprise Markets vs. Ad Hoc/Consumer Markets

One possibility is that convergence could be happening in just managed/enterprise environments, or just ad hoc/consumer marketplaces, but not both....

Regardless of the doubts or worries in some minds, convergence is a reality in both the ad hoc/consumer market and in the managed/ enterprise market.

There are some big names pushing hard in this area...

In the Consumer Market Convergence Is Happening

- 12.3% VoIP penetration (residential) in 2005 (http://blogs.pulver.com/jarnold/archives/2006/04/residential_voi.html)
- Vonage (hardware VoIP): 1.6 million customers as of April 1st, 2006, but "some" financial issues, including an accumulated deficit of \$455.1 million as of March 2006, presumably due in part to spending \$331.7 million on marketing during '05 and Q1/06. (<http://www.sec.gov/Archives/edgar/data/1272830/000104746906005887/a2169686zs-1a.htm>)
- Skype (software VoIP): over 100 million registered users, worth \$2.4 billion, at least to eBay. (<http://investor.ebay.com/ReleaseDetail.cfm?ReleaseID=195324&FYear=>)
- TiVo (digital TV recorder, includes ability to transfer content over the network to a laptop): 4.36 million subscribers (<http://biz.yahoo.com/e/060414/tivo10-k.html>)
- "Microsoft TV Overview" (http://www.microsoft.com/tv/MSTV_Overview.aspx)
- "IPTV prepares for prime time." 6/5/06 article discusses ATT's plans for rolling out IPTV in production this summer. (http://news.com.com/IPTV+prepares+for+prime+time/2100-1037_3-6079710.html)

Convergence Is Happening in the Enterprise Market, Too

- Cisco and Avaya lead the corporate VoIP market. (<http://www.forbes.com/markets/2006/05/22/avaya-cisco-0521markets02.html>)
- "[Cisco's CEO] said the company's enterprise-communications group, which specializes in the voice-over-Internet Protocol market, saw sales increase by 40% over the last year." (http://www.forbes.com/technology/2006/05/10/cisco-networking-voip_cx_df_0510cisco.html)
- Or consider Avaya's deal with the US Army: "Avaya has been named one of 10 companies selected to participate in a \$4 billion U.S. Army project that will overhaul voice and data communications infrastructures of U.S. Army bases worldwide." The Infrastructure Modernization (IMOD) contract will support the Army's Installation Information Infrastructure Modernization Program (I3MP) with a single integrated communications system to seamlessly integrate voice, data, inside/outside cable plant and transmission products and services. (<http://www.avaya.com/gcm/master-usa/en-us/corporate/pressroom/pressreleases/2006/pr-060524.htm>)

So Convergence is A "Done Deal," Right?

Unfortunately not... convergence is still subject to a variety of technical, political and institutional threats. We'll now go over a variety of different ways to stand in front of the on-rushing convergence train.

"Underprovision Your Bandwidth"

The #1 way you can hinder convergence is by under-provisioning your bandwidth.

Operationalizing this for colleges and universities:

- Avoid 10Mbps Ethernet drops and half duplex hubs (likewise avoid relying on 802.11b wireless)
- Use gig (or 10gig) in the core, not 100Mbps
- Have sufficient wide area bandwidth (and no, just avoiding flat-topping the mrtg / rrdtool graphs isn't enough; you need some headroom); in most cases NxT1 or fractional DS3 will be too small

Measure your performance (latency and variation are key indicators)... even wide area networks can do fine on these sort of measurements.

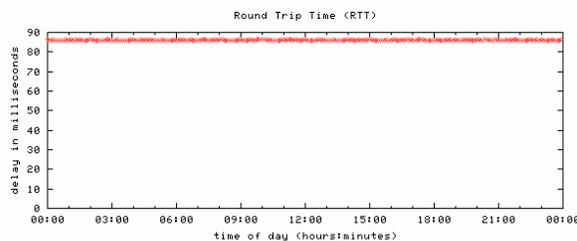
RTT And Loss Measurements

amp-bu from amp-uoregon

[\[NLNR\]](#) [\[AMP\]](#) [\[monitors\]](#) [\[weekly graphs\]](#) [\[reverse\]](#) [\[derivatives\]](#) [\[routes\]](#) [\[src info\]](#) [\[dst info\]](#) [\[2006.6.5\]](#) [\[2006.6.7\]](#)

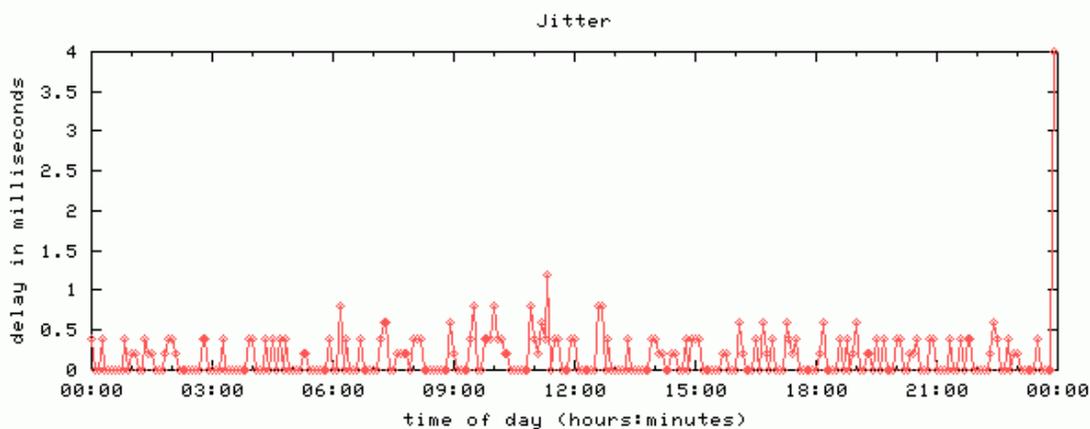
Tue Jun 6 2006

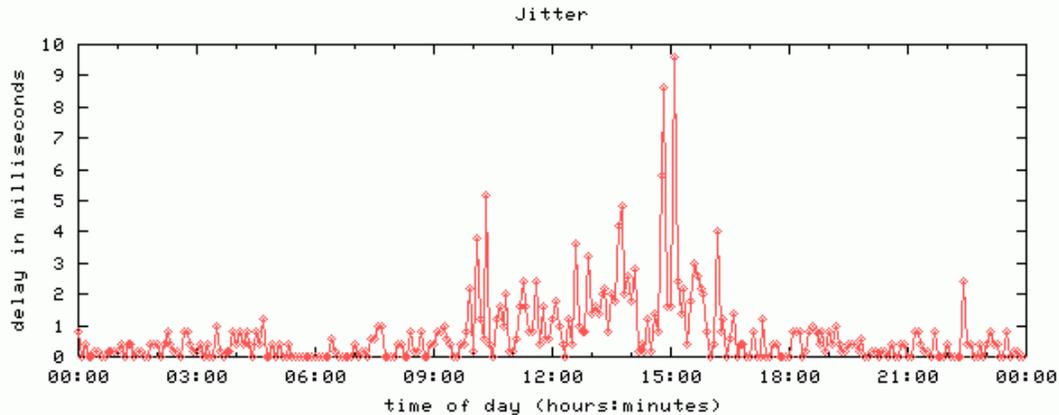
min: 86.00, mean: 86.09, max: 88.00, stddev: 0.28 (ms)



<http://watt.nlanr.net/active/cgi-bin/daily.cgi?amp-uoregon/HPC/data/amp-bu/106.6.6>

Guess Which Network Will Have More Of a Problem With Convergence?





The top graph is associated with the UO->BU graph shown on the previous page; the bottom graph (also from UO) will go unidentified. Note the different scales used.

"But We're Overrun With P2P Traffic!"

Not all P2P traffic is inherently bad (e.g., y'all know that Skype is P2P-based, for example...)

Define and architect the service you provide so that the institution (and the network) won't die if people actually use what you've built!

You can manage peer-to-peer traffic with a Packeteer or similar appliance (for now, but that may become impossible for technical or "net neutrality" reasons in the future)

Calibrate your service against what's available from the consumer marketplace, e.g., Comcast now offers 6Mbps down (the equivalent of 4xT1!) for \$57.95/mo and that's for non-cable customers!

"We Just Can't Afford Enough Bandwidth!"

Then the pricing/funding model you're using is wrong. Internet access is NOT a costless service. Just like water or electricity, it needs appropriate funding relative to its value/importance.

Consider the two scenarios (as extremes)...

- Small college (2,000 total users, 9 mo/year, \$10/month "value" (e.g., like el cheapo dialup)): \$180,000/year
- Large university (20,000 total users, 9 months/year, \$57.95/month "value" (like cable modem service): over \$10 million/year

You SHOULD be able to buy a lot of bandwidth for \$180,000 to \$10+ million per year.

"What If We Just Meter Usage?"

In a metered usage scheme, users pay by the byte for their network traffic. Variations may provide for some base traffic allocation before any charges actually accrue, and at some sites departments (rather than individual users) may actually end up being billed, etc., etc., but in any such metering scheme, "the meter's running" when traffic is flowing over the network.

Voice has relatively modest bandwidth requirements so metering will likely not preclude voice convergence. Metering will, however, render video convergence financially difficult. (e.g. $1.5\text{Mbps} / 8 * 60 * 60 * 24 * 30 = >486 \text{ GB/month}$, or 70 full length movies)

Other Metering Issues

Most Americans are used to local landline calls being unmetered (cell phones have tried to change that paradigm, but that paradigm's being rapidly eroded, e.g., with free incoming calls, free nights and weekend calling, etc.)

A metered environment needs a billing support system to handle revenue collection. That can be expensive! And once you start metering, people start looking for ways to "game" the system (open jacks, anyone?), and an adversarial model is created.

IMHO, metering is just really a bad idea.

Artificial/Unrealistic Demands

At the same time I oppose metering, you should also know that I oppose artificial/unrealistic "tests" or "challenges" of converged networks.

For example, a classic example of an unrealistic network demand for a converged network is **uncompressed high definition video over IP – that can run 1.2-1.5 gigabit per second**. At that rate, dedicated video networks make sense.

There's no problem handling MPEG1 video (at 1.5Mbps) however, or even reasonable amounts of MPEG2 video at 1.5 to 20Mbps (on a fast Ethernet connection going into a gig core).

Wayne Caswell – What would be unrealistic depends on aggregate network throughput. A 1.5 Mbps DSL connection can easily stream Internet radio and MP3 music but would be overwhelmed by HD video compressed with MPEG2, since that would require ~20 Mbps. A 100 Mbps FTTH connection can handle several HDTV streams, and with a terabit network, you could even stream uncompressed HDTV. But these examples assume that everyone isn't doing the same thing at once. The network operators' worry is that if broadcast TV is replaced by on-demand IPTV, a city of 3 million people might need aggregate performance of 20 Tbps, assuming that 1/3 of the population was watching TV at any one time, that they were each watching something different or at a different place in time, and that they weren't receiving streams each (for picture-in-picture or to record multiple shows on a DVR).

QoS – "Let's Make the Network Complicated"

Complexity is the 2nd biggest enemy of network convergence.

Anyone remember ATM (the network protocol, not the cash machines)? Classic example of a complicated network protocol with lots of knobs; a technology that could be counted on to often end up misconfigured with tragicomic results.

"Today's ATM" consists of complicated QoS schemes imposed on top of what would otherwise be a perfectly usable packet network.

"But, but, but, we NEED quality of service for converged networks... don't we?"

Do We Need QoS?

In a lightly loaded ("overprovisioned") network, a network with QoS and a network without QoS work effectively the same (QoS provides protection against packets being dropped, but so does extra headroom, and extra headroom is a far simpler and more robust solution).

For network engineers I highly recommend John Kristoff's Internet2 02/06/2006 Joint Tech's cleverly named talk, "Tripping on QoS", but let me just give you John's bottom line: "In a nutshell, I think you usually don't need [QoS] [...] but theology may trump science."

(<http://www.internet2.edu/presentations/jt2006feb/20060206-qos-kristoff.pdf>)

QoS: Neither On the LAN Nor the Internet

Let me also be explicit that when I say QoS isn't necessary or a good idea, I'm talking both about on the LAN and over the Internet as a whole.

On the LAN, it is cheap to provision fast connections with lots of headroom instead of deploying QoS.

When it comes to going to the Internet, connectivity is more expensive, but we simply have never figured out how to make wide area inter-domain premium QoS work.

But don't take my word for it...

Two Memorable QoS Quotes

#1: "After several years of experience attempting to deploy an inter-domain, [expedited forwarding] - based, virtual wire service in the Internet2 environment, the Internet2 QoS working group has concluded that any reservation-based form of QoS faces prohibitively difficult deployment obstacles, including:

- All-or-nothing network upgrades for providers (e.g. all access interfaces must police)
- Dramatic changes to network operations, peering arrangements, and business models
- Absence of suitable means to verify service (by users or providers)
- Moreover, within the Internet2 environment very few application performance problems can be traced to network congestion. Instead, end-to-end performance is often hampered by faults on or near end-systems including: broken TCP stacks (e.g. inadequate socket buffering), Ethernet duplex mismatch, and crummy cabling (e. g. CAT3, shared media, or physical damage)."

(<http://qbone.internet2.edu/>)

#2: Scott Bradner, Harvard: "The Internet is not reliably crappy enough."

(<http://www.educause.edu/ir/library/powerpoint/LIVE063.pps>)

QoS Is Not the Only Form of Network Complexity to Avoid

While QoS is the most commonly considered form of network complexity that potentially inhibits convergence, it is not the only type of network complexity you're potentially going to encounter. The two other major types are:

- extensively VLAN'd architectures
- architectures that employ "middle boxes" (such as network address translation (NAT) boxes or firewalls)

VLAN'd Architectures

Sometimes you'll see sites deploy a "converged" network that actually makes extensive use of VLANs to partition traffic. For example, an office might get a data VLAN'd network drop, a VoIP'd VLAN'd network drop, etc., with each drop using a different subnet.

I consider this to be "cheating" – yes, all the services are being delivered over IP, however at least at the edge, two, three or <N> separate networks are being presented to the user...

Yes, VLANs give you more control over your traffic, but at the cost of increased network complexity and loss of one-drop-for-everything.

Firewalls and NAT Boxes

Another way that you can simultaneously increase the complexity of your network and potentially thwart convergence is through the deployment of firewalls, NAT (network address translation) devices, and other "middle-boxes."

Yes, I know that firewall deployment is a matter of security dogma (particularly in some highly regulated environments, such as healthcare), however firewalls, NAT boxes and other middle-boxes greatly complicate deployment of converged services, particularly for incoming traffic.

Specifically, middle-boxes cause a loss of "Internet Transparency" and break the end-to-end model...

End-to-End Model & Internet Transparency

- "Architectural Principles of the Internet," Brian Carpenter, June 1996, <http://www.ietf.org/rfc/rfc1958.txt>
- "Internet Transparency," Brian Carpenter, February 2000, <http://www.ietf.org/rfc/rfc2775.txt> (Brian Carpenter is currently chair of the IETF.)
- "Some Internet Architectural Guidelines and Philosophy," Bush and Meyer, Dec 2002, <http://www.ietf.org/rfc/rfc3439.txt> (Dave Meyer is with UO, Cisco and the IAB.)

Regulatory and Policy-Related Issues, Such As Network Neutrality and CALEA

"We're From the Government, We're Here to Help"

The third way that convergence could stall is via regulation/policy, including things like the network neutrality debate and CALEA (Communication Assistance for Law Enforcement Act).

Regulatory issues are coming about now because convergence IS occurring...

- For example cable companies (and third parties) want to offer voice over IP; incumbent telcos (and third parties) want to offer IP video.
- Those third parties may economically threaten both the cable companies and the incumbent telcos.
- Law enforcement believes VoIP is "real" enough to potentially be used by criminals and terrorists.

How Did the Network Neutrality Issue Pop Up?

From facilities-Based Providers, Such As AT&T...

"They don't have any fiber out there. They don't have any wires. They don't have anything. They use my lines for free – and that's bull. For a Google or a Yahoo! or a Vonage or anybody to expect to use these pipes for free is nuts!"

AT&T Chairman Edward Whitacre Jr., "Rewired and Ready for Combat," Business Week, November 2005, http://www.businessweek.com/magazine/content/05_45/b3958089.htm.

And from Verizon...

"The network builders are spending a fortune constructing and maintaining the networks that Google intends to ride on with nothing but cheap servers. It is enjoying a free lunch that should, by any rational account, be the lunch of the facilities providers."

John Thorne, Verizon SVP and Deputy General Counsel, "Verizon's Executive Calls for End to Google's 'Free Lunch,'" Feb 7, 2006, <http://www.washingtonpost.com/wp-dyn/content/article/2006/02/06/AR2006020601624.html>

And BellSouth...

"A senior telecommunications executive said yesterday that Internet service providers should be allowed to strike deals to give certain Web sites or services priority in reaching computer users, a controversial system that would significantly change how the Internet operates."

"William L. Smith, chief technology officer for Atlanta-based BellSouth Corp., told reporters and analysts that an Internet service provider such as his firm should be able, for example, to charge Yahoo Inc. for the opportunity to have its search site load faster than that of Google Inc."

"Or, Smith said, his company should be allowed to charge a rival voice-over-Internet firm so that its service can operate with the same quality as BellSouth's offering."

<http://www.washingtonpost.com/wp-dyn/content/article/2005/11/30/AR2005113002109.html>

Some Incumbent Facilities-Based Providers Arguments Against Net Neutrality

- *Net Neutrality???* You're regulating the Internet!
- We're a market economy, and we should be free to use our assets in a economically rational way, as the market may bear; if you want us to make additional investments in infrastructure, you need to let us earn a return on those investments.
- We're all routinely used to other examples of differentiated services with differential pricing (FedEx vs. "US Mail; First Class vs. Coach; etc.); why should network capacity be any different?
- Problems with differential pricing are all theoretical so far; let's not worry about hypothetical 'problems.'
- W/O pricing flexibility, some P2P bandwidth hog customer will end up being unfairly subsidized by innocent minimal-usage customers (Grandma reading her grandkids' email messages)
- Differentiated services are technically needed to accommodate voice, video and other advanced services [remember, their argument, not mine]
- Unless we have discretion when it comes to whom we partner/peer with, network stability and reliability may be adversely impacted
- US regulation will make our country fall behind unregulated overseas competitors
- If net neutrality regulations pass, the now-regulated carriers might be stripped of their ability to deal with spammers, denial of service attacks, etc.
- It might even be impossible to do such things as passively cache some content (many providers currently save content to local web cache servers when it is first retrieved, and then serve subsequent requests for popular pages from the local copy, thereby reducing bandwidth usage and accelerating delivery of that content)...
The problem? When the cache is out of space, it flushes unpopular content to make room, thereby "playing favorites" w.r.t. the most popular content.

And here's a novel argument...

- *"Forcing net neutrality would arguably violate the incumbent carrier's First Amendment rights by forcing them to carry speech they might otherwise wish to editorially exclude."* Miami Herald Publishing Co. v. Tornillo, 418 U.S. 241 (1974), http://www.epic.org/free_speech/tornillo.html

Arguments In Favor of Network Neutrality

- ***Fairness: all that's wanted is a level playing field.***
- ***Net neutrality is the network's '1st Amendment.'***
- Without network neutrality, users won't be able to get to some destinations at all (or if they can, those destinations might be slower than they are now)
- Without network neutrality, content providers might have to pay to be assured of acceptable network throughput, and those costs will be passed on to consumers, increasing everyone's costs
- Providers are common carriers using public assets (such as right-of-way or licensed spectrum) on the public's behalf, and they should not be allowed to discriminate in their marketing of those resources.
- With recent FCC actions with respect to the unbundling of network elements, incumbent facilities-based providers have a virtual monopoly on wireline broadband access. Monopoly or near monopoly providers should be held to a stricter standard of conduct than when competition and market forces can set prices/standards.
- Facilities-based incumbents are both providers of network access AND competitors to services, which might be offered over that network access (e.g., POTS from the incumbent vs. VoIP from Vonage). It would be naïve to assume that the incumbent could fulfill both roles fairly w/o regulatory guidance.

- When the network ceases being a clear channel, the task of debugging network problems may become more difficult or impossible ("Well, something's happening to your traffic after it leaves your computer before it gets to its destination, but we're having a hard time telling *what* that is...")
- The Internet's biggest advantage is the way it allows innovation w/o the need for prior approval or consultation. That freedom would be lost w/o network neutrality, and innovation would be stifled.
- Some cellular carriers have already adopted broad restrictions on what can be downloaded to phones. (<http://www.vzwdevelopers.com/aims/public/wapContGuide.jsp>)
- The "without assurance that investments can earn a return, we won't invest" argument of the incumbent telcos is undercut by the reality of what allegedly didn't happen after the 1996 telecommunication deregulation experiment. (See, for example, Bruce Kushnick's "\$200 Billion Broadband Scandal," <http://www.newnetworks.com/broadbandscandals.htm>)

Higher Ed and Network Neutrality

Virtually all higher education-related network or IT-related organizations (such as Educause, Internet2, ACE, AAU, NASULGC, ALA, etc.) support network neutrality rather than the carrier-favored network non-regulation. (See: <http://www.educause.edu/netneutrality/>)

But we need to "thread the needle" carefully...

- Remember that higher ed runs its own uncongested/exclusive network, Internet2
- Higher ed also routinely relies on packet shaping to control recreational traffic on its own networks
- We all want carriers to deal with their abusers...

Net Neutrality-Related Bills On The Hill...

HOUSE:

- **HR5252**, Barton's "Communication Opportunity, Promotion and Enhancement Act of 2006" (Telco-favored network non-regulation bill). Passed on House floor vote.
- **HR5273**, Markey's "Network Neutrality Act of 2006"
- **HR5417**, Sensenbrenner's "Internet Freedom and Nondiscrimination Act of 2006"

SENATE:

- **S1504**, Ensign's "Broadband Investment and Consumer Choice Act"
- **S2360**, Wyden's "Internet Non-Discrimination Act of 2006"
- **S2686**, Steven's "Communications, Consumer's Choice, and Broadband Deployment Act of 2006" Floor vote expected in August.
- **S2917**, Snowe's "Internet Freedom Preservation Act"

So What Do The Bills Provide?

Not going to attempt to summarize them here; too complex, and too subject to amendments in "real time" right now. Check <http://thomas.loc.gov/> for the current bill text and updated status.

Some partisan alignment has occurred (Democrats generally in favor network neutrality, Republicans generally in favor of telco non-regulation), but the issue puts Senators and Representatives in a difficult position because they are facing pressure from both well-heeled telco lobbyists and from large numbers of public organizations and members of the public (including some "odd bedfellows").

Do We Even Need New Laws?

We already have an agency that's supposed to be handling communication regulation, the FCC.

We have another agency, the FTC, that's supposed to be insuring businesses conduct their business in a fair and competitive manner.

Is it possible that these agencies could manage broadband access in a way that meets the desires of both the anti-regulation and the net neutrality camps?

The FCC has offered **four principles** that will guide its policy making...

- (1) Consumers are entitled to access the lawful Internet content of their choice;
- (2) Consumers are entitled to run applications and services of their choice, subject to the needs of law enforcement;
- (3) Consumers are entitled to connect their choice of legal devices that do not harm the network; and
- (4) Consumers are entitled to competition among network providers, application and service providers, and content providers.

Would those four principles be enough???

CALEA

CALEA (the Communications Assistance for Law Enforcement Act) mandates that facilities based network service providers configure their networks so as to enable them to respond to lawful intercept requests from law enforcement.

Recent FCC rulings have indicated that CALEA's scope includes higher education. See Educause's excellent CALEA resource center at http://www.educause.edu/Browse/645?PARENT_ID=698.

Depending on the outcome of pending litigation, you may need to architect your campus network to support CALEA-related requests.

How Might CALEA Affect Convergence?

If the litigation fails and CALEA compliance is required, **costs increase with network speed** (it is trickier/more expensive to handle lawful intercepts on a ten gig network rather than a one gig network).

Some **sites may postpone or roll back upgrades** (and remember, bandwidth is the universal solvent to convergence problems).

There may be a perception that a converged network is more likely to be the subject of CALEA-related requests than a non-converged network (but of course ad hoc VoIP will almost always be possible, and CALEA isn't limited to just voice communications anyhow).

More on CALEA and Convergence

The costs of complying with CALEA may have a displacing effect on convergence-related projects; limited funds and limited staff may be devoted to CALEA compliance rather than moving toward a converged network architecture.

Users may lose trust in the network, and shun it.

The antidote to CALEA's potential intrusiveness is generally considered to be strong encryption, but strong encryption can pose unique challenges when applied to jitter sensitive real-time applications. Interesting projects are beginning to emerge, including Phil Zimmerman's Zfone:

<http://www.philzimmermann.com/EN/zfone/index-start.html>

Conclusions

There are some technical obstacles that you can run into, but the good news is that **a clean overprovisioned network will generally be all that's needed to support convergence**. STRIVE to deploy fast, simple networks.

There are policy obstacles that you may run into, including most notably the current net neutrality /anti-net regulation debate, and CALEA. The outcome and impact of those unfolding regulatory areas is still unclear.

Convergence of voice, video and data networks is occurring, and IMHO **is** an inescapable inevitability unless people try to "help" too much.