Thinking About Your Wide Area Network Connectivity

NWACC Networking for the 21st Century
Portland, Oregon, June 15th, 2001

Joe St Sauver, Ph.D.
(joe@oregon.uoregon.edu)
Computing Center
University of Oregon

I. Introduction

Why Bother Thinking About Your Wide Area Connectivity?

- If you are like many folks, it may be the area of networking you enjoy least
- Your existing connectivity may (or may not) fully meet your users' expectations
- Wide area connectivity is probably your single largest budget line item (except for staff salaries)
- You may have been doing the "same old thing" for some time, and may wonder if you should be looking at new options.

Disclaimer

If you’re happy with what you’re currently doing, by all means keep doing it, more power to you. Don’t mess with success.

Every situation is going to be subtly different, so before you change course based on anything I suggest, get a second opinion and be sure that YOU are completely comfortable with whatever approach you finally select.

Having Said That...

Let’s begin by outlining where I believe many of you may be now, and then we’ll talk about where you may want to go in the future (and why).

Parts of this talk may be old hat for some of you, but we’ll nonetheless start from ground zero and build from there; we apologize for any material that’s “old news” for you.

II. Understanding The Traditional "Smaller College" Internet Connectivity Model

T1 (or 2xT1)

For the last decade or so, many smaller liberal arts colleges have connected to the Internet via a flat rate T1 (1.5 Mbps), or perhaps via a pair of flat rate T1s (3.0 Mbps) serviced by a commodity Internet Service Provider (ISP).

This is the same model that many smaller local ISPs have themselves used.

T1 or 2xT1

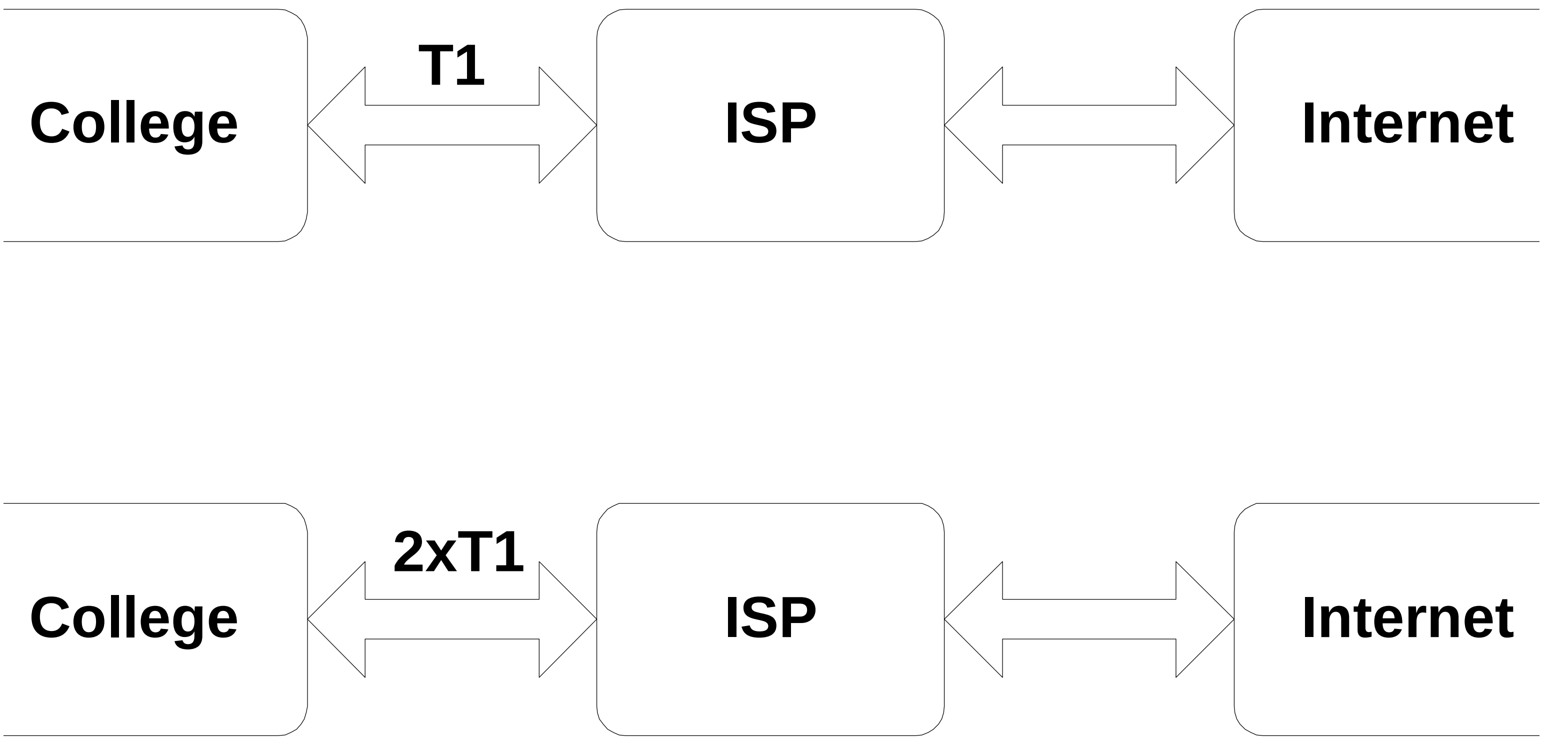

In General That Traditional Model Has Worked "Okay"...

It is straight forward/simple

For a typical smaller school, T1 (or 2xT1) connectivity is sufficient for faculty, student, and staff to have email access as well as casual access to the world wide web and it also usually sufficient to service the college's own web site.

Costs for this model are known/bounded

Total Costs for Traditional T1 (or 2xT1) Service Include:

- The ISP's port charge (the ISP's charge for agreeing to provide you with X Mbps worth of Internet drainage)
- Local loop costs (e.g., the cost of leasing a physical circuit from the phone company to the ISP's closest point of presence (POP))
- Network hardware costs (router, CSU/DSU)
- One time installation costs

Typical T1 Port Charges From National Backbone Providers

AT&T	$1,970/mo + $1K setup
UUNet	$1,795/mo + $3K setup
Sprint	$1,420/mo + $1K setup
Globix	$1,328/mo + $549 setup
CWIX	$1,300/mo + $300 setup
Qwest	$1,185/mo + $1K setup
Verio	$995/mo + $500 setup (1 yr/term)
ELI	$700/mo + $1K setup (1 yr/term)

Data from http://www.boardwatch.com/

Local Loop Charges

… usually have a fixed and a milage sensitive component

Local loops may be provisioned by the ILEC (incumbent local exchange carrier, e.g., Qwest) or by a CLEC (competitive local exchange carrier such as ATT, ATG, ELI, etc.)

Sample (confusing) tariffs are available online from http://tariffs.uswest.com/

Figure roughly “n” hundred dollars/month

Hardware Costs

For a T1, you are typically looking at something like a Cisco 2620 or 2650 router, (which can be purchased with an integrated CSU/DSU you'll also need) for ~$2K--2.5K

Costs for hardware at the T1 level are quite modest, but required hardware costs can quickly become material at higher connection speeds

Installation Costs

Designed to cover time and committed equipment on the provider's end, PLUS (at least in some cases) the installation charge is intended to act as a deterrent to provider-hopping customers ("Gee, we just paid $3K for install… even if this other provider is a lot cheaper, we DID just pay $3K!)

Installation costs often start high so they ca be negotiated down as a "deal sweetener"

End Result...

The institution has Internet connectivity

The institution usually uses non-portable IP address space obtained from the ISP

The institution needs little if any wide area networking expertise on site

Relationship is a businesslike (paying) customer to (paid) service provider, a defined arrangement of known parameters

Known Parameters==SLAs?

Many ISPs now routinely offer "service level agreements" whereby the customer can request a refund of a portion of the amount paid if services fail to meet agreed upon specifications (for example, outages exceed some stipulated level).

In reality, SLAs are simply a distraction -- trivially small SLA refunds cannot begin to compensate you for poor quality service. Buy a good ISP **don't** shop for 'good' SLAs

III. Making the Traditional Model Cheaper #1: Frame Relay As A Replacement For Point-to-Point Circuits

Substituting Frame Relay T1s For Point-to-Point T1s

If the goal is to drive down the cost of provisioning T1s, one alternative to consider is substituting frame relay T1s for point-to-point T1s.

In the frame relay model, locations connecting via frame relay connect to a common frame relay “cloud” rather than connecting directly on a point-to-point basis.

Frame Relay (continued…)

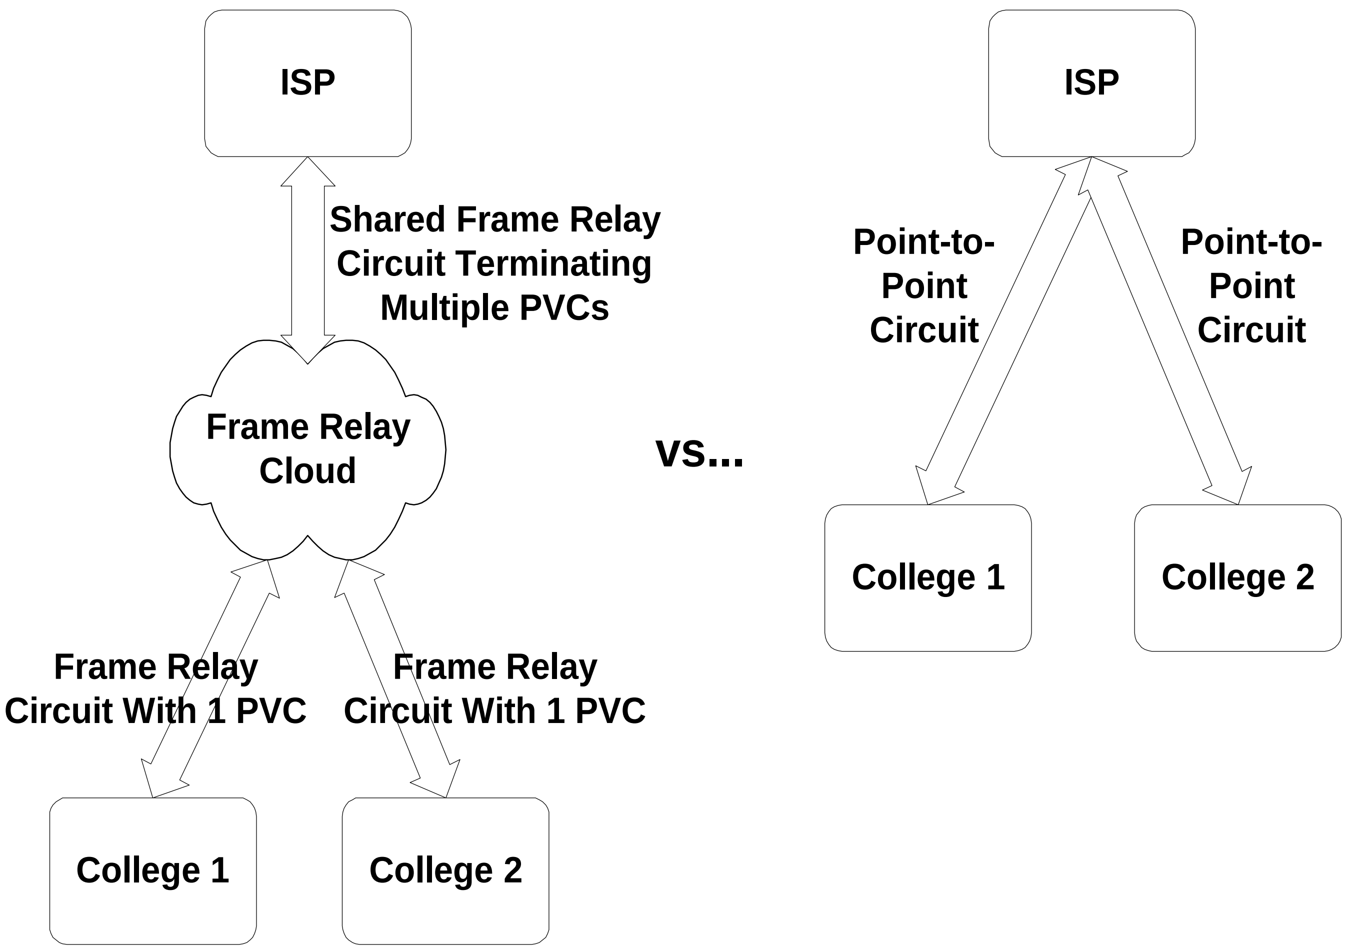

How Come A Frame Relay T1 Can Be Cheaper Than A Point-to-Point T1?

At least in Oregon under the State's Fast Packet contract, frame relay service is distance insensitive. This is hugely important if you're connecting a truly remote site.

Assuming the other end already connects to the frame relay cloud, you only end up paying for "one end" of the circuit plus (implicitly/indirectly) part of the other end

How Can Frame Relay T1's Be Cheaper (continued….)

The upstream (ISP) end does not need to ad equipment to terminate each new frame rela circuit on a 1:1 ratio (they can all be accommodated via a single shared device and shared frame relay connection)

Some ISPs are somewhat aggressive in how they underprovision/overcommit shared frame relay circuits

The CIR is seldom 100% of the nominal rat

"What's CIR?"

The Committed Information Rate (CIR) is the amount that you will be able to routinely send on a sustained basis. The CIR might be 256Kbps or 512Kbps or 768Kbps on a "T1"

Frame relay service is statistically designed to accommodate brief traffic bursts to the full nominal capacity of the circuit… but only for bursts, and then only when capacity happens to be available.

Example of When CIR Mattered...

We wanted to test Cisco's IP/TV MPEG1 IP multicast product with an Oregon Community College partner of ours. IP/TV would send multicast video traffic at 1.5Mbps on a more or less steady basis… this did NOT work well until we cranked the sending rate down to no more than the CIR (1Mbps in this case).

Frame Relay Pricing Under The OR Fast Packet Contract

T1: $450/month/end with one permanent virtual circuit (plus $616.74 installation) See: http://telecom.das.state.or.us/data/billing/nonwan1.htm

This is JUST local access (e.g., you're only replacing the point to point T1 local loop with frame relay service), this is not inclusive of Internet transit (port) charges

Internet Service (Port Charge) on a Frame Relay Circuit

Frame relay T1s often tend to end up being serviced via a local/regional ISP rather than a national backbone, which have different cost structures than national backbones

See: http://thelist.internet.com/areacode/541 (or whatever area code's of interest) for lots of providers offering frame relay service in the 541 area code

So Should I Convert My Point to Point T1s to Frame Relay?

The answer really is "it depends on a case by case basis." FR <u>isn't</u> always cheaper.

You really need to look at the cost of all the various options in your case (including paying attention to CIR rates when making comparisons or assessing the value of a point-to-point T1 vs. a frame relay "T1")

Also be aware that frame relay's traditional niche is "under siege" from DSL services...

IV. Making the Traditional Model Cheaper #2: DSL

DSL Service

Depending on where you're located, another alternative <u>might</u> be high bit rate DSL service, e.g., here in Qwest territory:

-- 1.2Mbps down/1.1Mbps up ($80/mo)
-- 4.4Mbps down/1.1Mbps up ($150/mo)
-- 7.1Mbps down/1.1Mbps up ($250/mo)

… this is <u>just DSL loop</u>, not Internet service

Internet Service for that DSL Loop… Add...

If you buy Internet service from Qwest.Net (one of many potential ISPs to service that DSL circuit):

-- 1Mbps $115/mo (e.g., total of $195/mo)
-- 4Mbps $135/mo (total of $285/mo)
-- 7Mbps $155/mo (total of $405/mo)

see: http://www.qwest.net/nav4/solutions/internet/ow_pricing.html (be sure to also figure in the cost of a business phone line)

"Wow… It Looks Like DSL <u>Could</u> Be a Real Bargain…"

<u>Key point 1</u>: you WILL NOT be able to order DSL service <u>everywhere</u> (particularly in the case of high bit rate DSL service). DSL service (particularly high bit rate DSL service) will only be available for sites "close" (18,000' max) to a suitably equipped telco central office, and then only at sites lucky enough to have excellent quality copper with no bridge taps, etc.

DSL... Bargain? (cont)

Key Point 2: High Bit Rate DSL is asymmetric… while you may get 7Mbps down to your site from the Internet, your upload speed to the Internet will only be ~1Mbps. [On the other hand, that asymmetry may mesh rather well with the usage pattern typical of most smaller colleges (where downloads from the Internet dominate uploads to the Internet)]

DSL… Bargain? (cont.)

Key Point 3: At least some DSL-servicing ISPs (such as Qwest.Net) force web traffic thru the ISP’s web cache, thereby breaking Internet transparency. [You wondered how they could afford to offer those great prices -- they are taking advantage of the fact that some large fraction of your traffic will probably be http, and probably web pages serviceable from a local web cache at that.]

DSL… Bargain? (cont.)

Key Point 4: Your equipment options may be limited. Most DSL service providers have standardized on a particular brand of DSL modem (such as the ~$300 Cisco 675 or 678), and that is what you will need to use if you want DSL service, like it or not. [Why is this an issue? Well, for example, many DSL modems are clearly consumer-grade rather than carrier-grade gear…]

DSL… Bargain? (cont.)

Key Point 5: You will have a restricted set of possible ISPs to offer service on your DSL circuit. E.g., most national backbones and many regional ISPs will not be available, the result being that you need to pick either a local DSL-servicing ISP or the phone company's unregulated ISP affiliate: http://www.qwest.com/dsl/learn/isplist.html

So Is DSL The Right Replacement for Traditional Point-to-Point T1 Service?

Probably not yet, at least not for mission critical services (good though DSL's pricing may be). DSL is still rather young and is really positioned as a consumer access technology, rather than a large business/ institutional access technology. Outage response and repair time is one particularly critical but currently unbounded factor...

So Is DSL The Right Replacement (continued…)

On the other hand, DSL **is** cheap and your equipment investment would be low, and since there is no need to enter a long term contract you **could** easily experiment with DSL if you had a mind to do so.

Unfortunately, if you're like many, your location may not let you get the DSL service you might want to buy (at least not yet, anyhow), making this all rather moot.

V. Limitations of the Traditional Connectivity Model (In Any of Its Permutations)

The Zen-Like Paradox of the Traditional Connection Model

The basic T1 (or 2xT1) connection model is built on a very subtle paradox: it only works well when people:

(a) don’t use it (or don’t use it much) and

(b) don’t depend on it.

Why Do We Say That a T1 (or 2xT1) Only Works Well When People Don’t Use It?

A T1 (or 2xT1) is easy for even a single directly connected user to saturate… to say nothing of hundreds (or thousands) of simultaneous directly connected users.

1.5 to 3 Mbps TCP flows used to be a big deal for a end system to generate, but now projects like www.web100.org will push routine system throughput to 100Mbs...

Or Consider, As A Benchmark, the Capacity of a T1 in "CD's"

1.5Mbps ==>

$$\frac{1.5\text{Mbps}}{8\ \text{bits/byte}} * (60\ \text{min/hr})(60\ \text{sec/min}) ==>$$

675 Mbyte/hour (e.g., roughly one "CD equivalent" worth of "stuff" per hour) -- and many free software distributions now are multi CD affairs….

There Is No Financial Incentive <u>Not</u> to Use Capacity

Moreover, virtually no college (at least none that I'm aware of) implements usage sensitive pricing: use a lot, use a little, it doesn't (financially) matter to the end user.

We know what happens when "free" shared resources are made available: the "Tragedy of the Commons" results in increasing levels of use by economically rationally users until the resource is overused/ruined.

“OK, What Do You Mean About Not Depending On It?”

Consider electrical power or plain old telephone service -- those are mission critical resources which you “depend on” and which are available everywhere

For a long time and at many sites, however, Internet service has been viewed somewhat more casually, with access, reliability and redundancy a secondary consideration.

“Not Depending On It…” (continued)

For example:

-- consider the casualness with which we joke about “the network being down” (as long as it isn’t down TOO long) or

-- consider the fact that while traveling, we accept the fact that network connectivity may be poor (e.g., dialin), or non-existent at many locations

-- we don’t REALLY depend on it

“Not Depending On It…” (continued)

We have routinely seen multiple parallel “dedicated-purpose” T1s deployed to shelter one category of traffic from another (prime example: generic Internet traffic vs. H.323 video conferencing traffic). Why? Users are unwilling to take a “chance” that the network will be busy due to random traffic at the time a high priority apps (like a class videoconference) needs bandwidth...

“Not Depending On It” (continued…)

A single T1, or two T1s from the same provider, also means that you have no protection against provider-related difficulties -- if your one-and-only provider goes down, you have no redundancy.

You can obtain redundancy and a degree of survivability by multihoming (buying service from more than one ISP).

VI. Multihoming

Multihoming

Multihoming for redundancy and improved reliability is a relatively common practice today, although not necessarily at the T1 (or 2xT1 level).

By providing two paths to the Internet, failure of a single ISP should be insufficient to take you off line.

By multihoming, you also become less a "captive customer" of a single ISP.

Multihoming

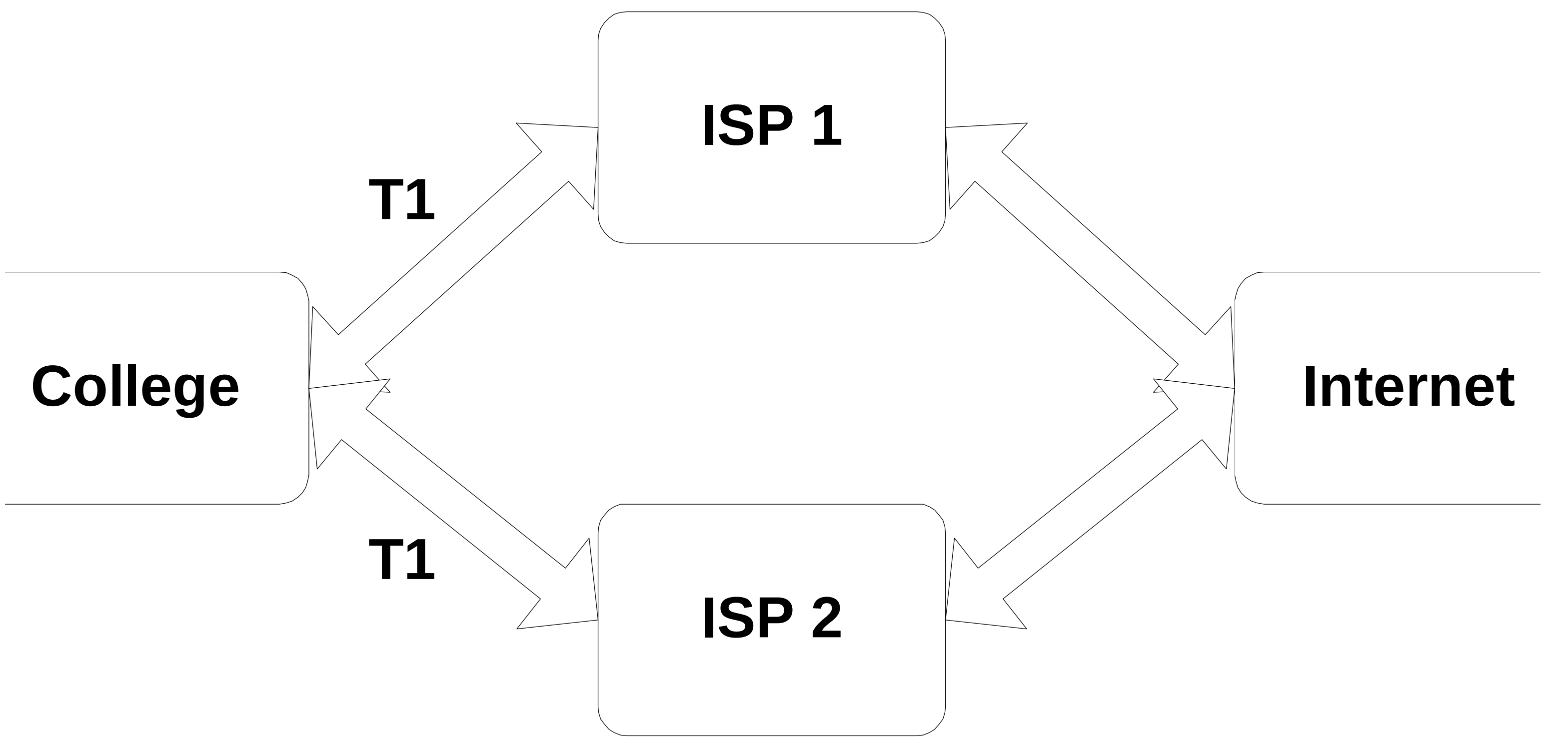

Some Limits of Multihoming

Of course, if you really want to protect your traffic, you need to have enough capacity to handle ALL your load via either single ISP, in case either of your ISPs has a failure)

Thus, if you are multihomed, multiply virtually all costs by two... and as you write those extra checks, remember you're buying 'insurance' network capacity which you may never actually need to use.

But You Won't Be Able to Resi Eating the Emergency Rations.

Of course, it takes tremendous discipline to provision redundant network capacity and then not routinely use it. You WILL be tempted to eat the emergency rations (e.g., to routinely utilize all your circuits to their full capacity), even though you multihomed for improved reliability, not to add extra capacity.

“Why Is It Bad to Routinely Use All Available Capacity In A Multihomed Scenario?”

If you are multihomed and routinely run both your circuits to full capacity, you will be horribly congested if/when one or the other of those circuits fails and ALL your load has to be serviced by the remaining circuit (which will have far less than the capacity it needs to do the job).

And Recognize That Your Load Will Not Automatically Balance

With two providers, you might assume that your load will automatically split nicely between the two in some magic fashion. This is NOT the case, particularly if your two providers are not of roughly equal influence Internet wide (e.g., assume one provider is a national backbone, and the other provider is a regional provider). Configuring to use both 50-50 will be tricky

And Recognize That Multihoming May Make Your Traffic Asymmetric

When we talk about traffic being "asymmetric," we mean packets may go out to an Internet destination via one path, but come back another. This *can* have some subtle but important implications; see, for example: http://www.internet-2.org.il/i2-asymmetry/index.htm

More Implications of Multihoming

The “simple” act of adding a second ISP also raises the technical bar for your wide area connectivity substantially, adding:

-- the need for provider-independent
   independently routable address space,
-- an autonomous system (AS) number,
-- larger/faster/more expensive routers
-- use of BGP4 for your external routing

And If You're Multihoming for Improved Reliability...

Be sure to also investigate how your local loops are being provisioned (and note that simply buying one local loop from your ILEC and another from a CLEC does NOT guarantee that the loops will share no common points of failure -- the "CLEC's" loop may end up being provisioned for the CLEC by the ILEC in the same conduit as your original loop)

So <u>Should I</u> Multihome?

If you are serious about improving your reliability, and if you have the resources (financial, technical, and managerial), yes.

In many cases, however, smaller schools may not have those sort of resources, and may have to forgo the advantages of multihoming.

VII. Peering

Internet Transit vs. Internet Peerage

When you buy Internet transit from an ISP, the ISP agrees (for a fee) to carry traffic for you to/from ANY Internet destination (up to the agreed upon capacity purchased).

Internet peerage is different. When you peer with someone, you agree to exchange direct customer traffic, and ONLY direct customer traffic, usually without financial settlements

Peering Does Not Eliminate the Need for Transit Connectivity

I want to stress that peering does NOT eliminate your need for transit connectivity. It should <u>REDUCE</u> the amount of transit connectivity you need to buy, but you will <u>never</u> peer with everyone, so you still need to have some guaranteed path (e.g., transit connectivity) to/from your site for all those providers with whom you <u>don't</u> peer.

The Mechanics of Circuit Based Peering: Bad Scalability

Peering could occur via a dedicated circuit established directly between each two interested parties. However, for more than two parties, that approach scales poorly. If four sites (A, B, C, and D) all wanted to directly peer amongst themselves, that would require six circuits (A-B, A-C, A-D, B-C, B-D, C-D); if 6 sites wanted to peer, that would require 15 circuits, etc. Ugh!

Voila! The Exchange Point

The preferred approach, therefore, is for all parties interested in peering to agree to meet at an exchange point, where customer traffic can be exchanged between multiple peers without requiring an ISP to provision a new circuit for each peer of interest.

The Oregon Internet Exchange is one example of a local exchange point (see: www.oregon-ix.net). See also www.ep.net

Direct Circuit-Based Peering vs. Use of An Exchange Point

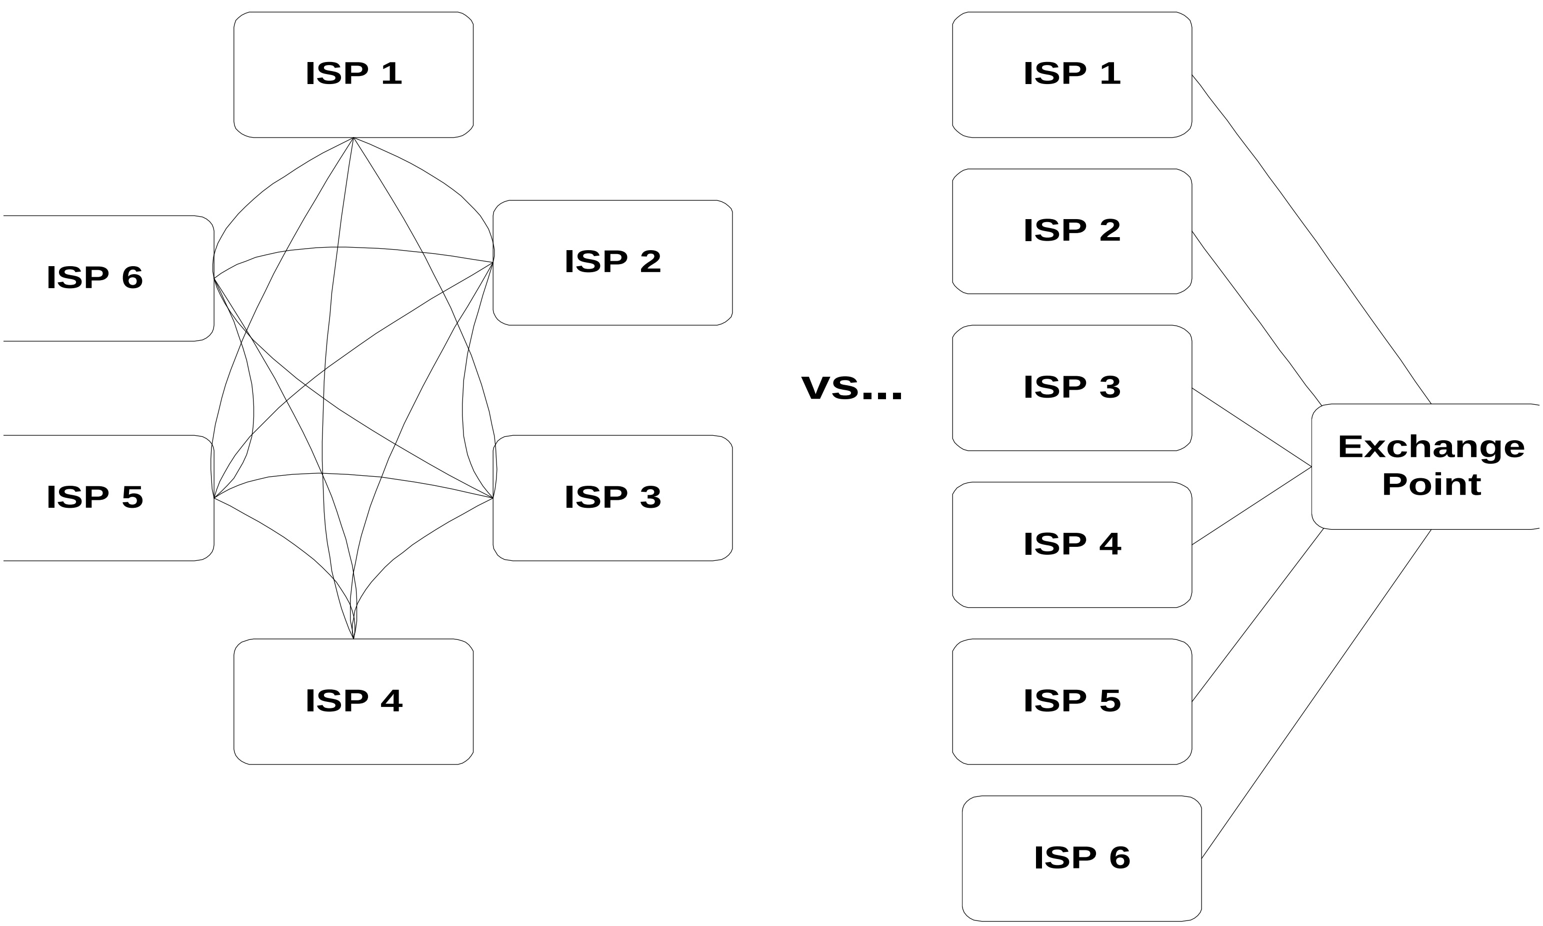

Multilateral vs. Bilateral Peering at Exchange Points

Arranging to exchange traffic among peers may be done via one of two models: multilateral peering or bilateral peering.

In multilateral peering, you agree to peer with anyone else present at that exchange point. Advantage: simple/straightforward.

In bilateral peering, peering is arranged separately between each pair of peers. Advantage: greater control and selectivity.

“Peering Sounds Great -- Why Doesn’t Everyone Do It?”

Participating at an exchange point raises the same technical challenges as multihoming; small sites may find the benefits not worth the hassle.

Participating at an exchange point will also have some direct costs (e.g., a circuit from your site to the exchange point, equipment for use at the exchange point, maybe exchange point participation fees, etc.)

‘Why Doesn’t Everyone Peer?’ (continued….)

Also, providers will usually only agree to peer when they are roughly the same size. [If someone’s small & a potential customer, peering with that potential customer may reduce the likelihood of that potential customer purchasing commodity transit.]

Peering only makes sense if you’ve got a material amount of customer traffic to exchange with the peers at an exchange.

Participating At An Exchange Point Does Mean More Than Reducing Transit Costs...

It is also important to point out that when you participate at an local exchange point, you get more than reduced transit costs -- local traffic stays local (rather than bouncing up to Seattle or down to the Bay Area or over to Denver to be switched from one transit provider to another), which means performance will usually improve.

It's a Bummer Being Small...

By now, you may be noting a recurring theme: there are lots of interesting things (like multihoming or participating at an exchange point) which are hard to do if you're a small college or other small entity. Are small colleges just out of luck?

No. Small colleges should strongly consider participating in network consortia, which can aggregate their demand.

VIII. Network Consortia

There Is Strength In Numbers

When multiple smaller sites band together, a lot of things which might be impractical for each of the individual smaller sites to do on their own suddenly become (at least potentially or theoretically) possible.

Some Examples of Local Network Consortia

A fine example of a network consortia is Oregon's OPEN statewide K12 network, interconnecting virtually all public K12 schools in the state of Oregon. See: http://www.open.k12.or.us/

Another example is OWEN/NERO, connecting virtually all public universities in the state of Oregon. See: http://www.nero.net/

Example Network Consortia Connectivity Model

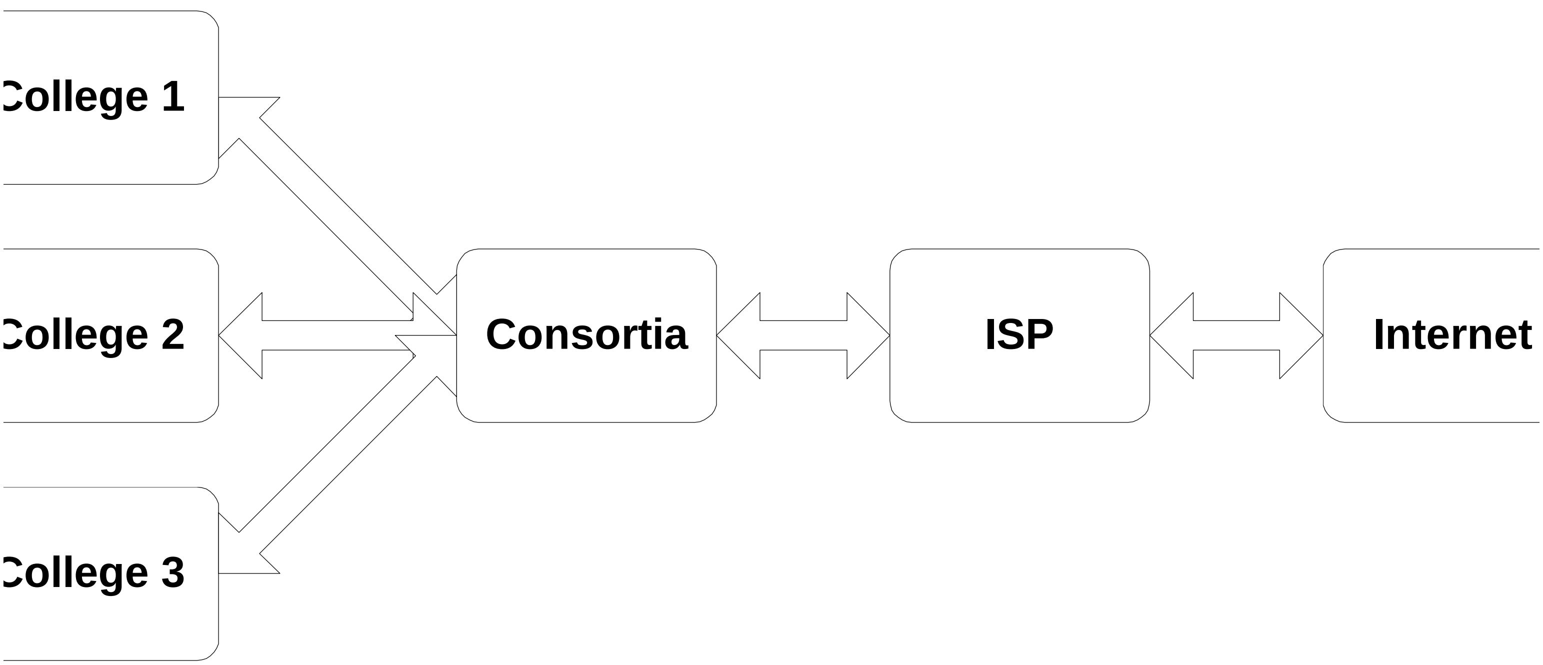

Network Consortia Model

In the network consortia model, the consortia acts as an intermediary between the end site and the ISP, buying wholesale network bandwidth from one or more internet service providers and then making that aggregate bandwidth available to the members of the consortia.

Implications of Participating in a Network Consortia

Participating in a network consortia has a number of interesting implications:

-- greater statistical multiplexing
-- improved reserves against load spikes
-- bifurcation of traffic into local (no cost) and non-local (regular Internet) traffic
-- increased opportunities for resource sharing & experimentation

Statistical Multiplexing of Demand

By aggregating a large number of users, brief peaks and valleys in demand end up “averaging out,” resulting in more efficient circuit utilization.

Compare the two graphs on the following slide...

Statistical Multiplexing

Highly aggregated traffic showing smoothing/statistical multiplexing

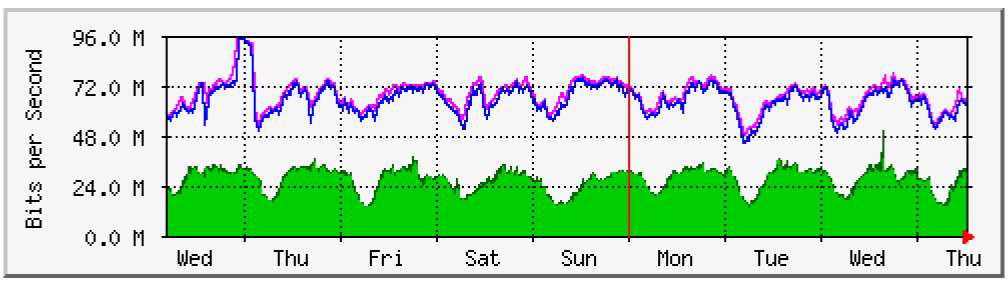

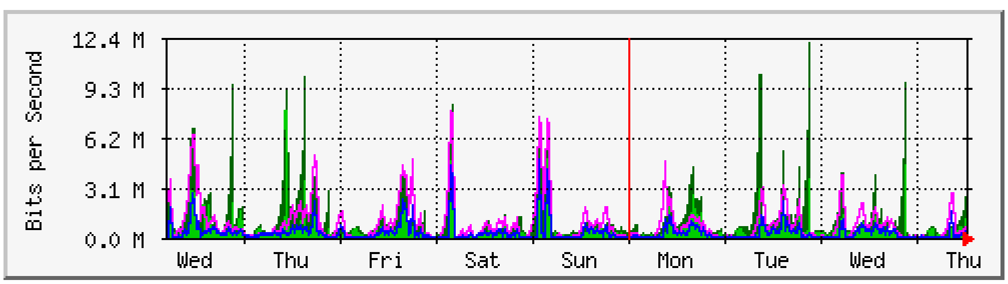

Versus traffic from a single subnet, showing much wider variation/more “peakyness”

Inreased Reserve Capacity to Handle Load Spikes

Aggregation also allows for greater reserve capacity to handle load spikes than individual sites could provision. Consider: (a) ten sites each buying T1s individually (and each never having more than a T1's worth of bandwidth) vs. (b) ten sites sharing 15Mbps, where at least some fraction of that capacity (let's say 5Mbps) is available for use to buffer brief spikes in demand...

Reserve Peaking Load Capacity

Limited (1.5Mbps) peaking load capacity:

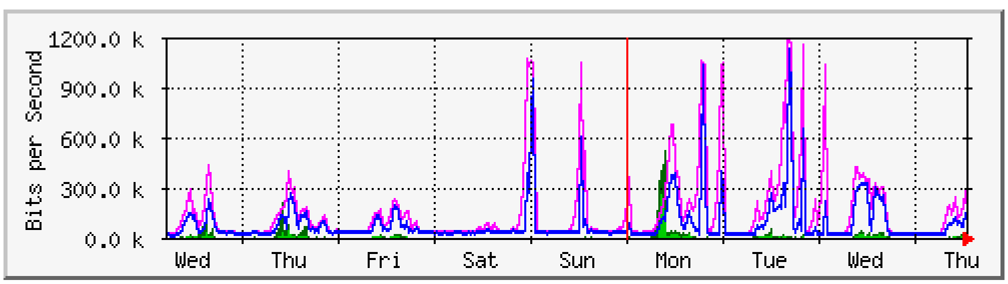

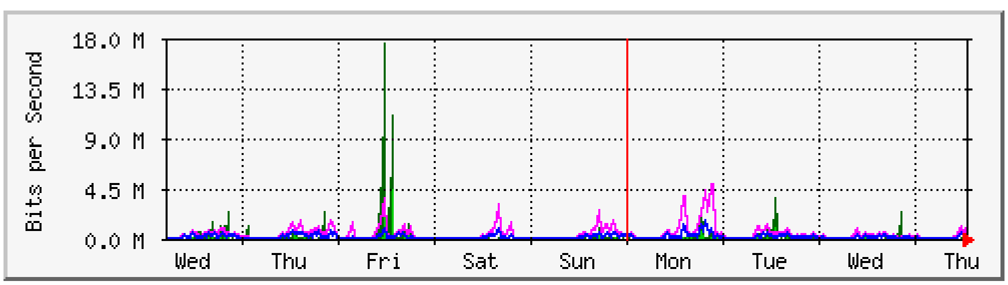

vs **extensive** excess load capacity available for peaking use (note scale on vertical axis)

Local vs. Non-Local Traffic

When you buy connectivity directly from an ISP, all your wide area traffic flows via that ISP's connection (and all that traffic counts against your purchased Internet bandwidth)

When you buy connectivity from a local consortia, traffic destined for another member of the consortia never leaves the local consortia's circuits, and hence doesn't use any expensive wide area connectivity

Local vs. Non-Local Traffic (continued…)

Thus, the utility of network consortia improves with the size of the consortia -- the more members participating in a consortia the better (just like peering points)

On the other hand, the bigger the consortia grows the harder it becomes to <u>informally</u> handle allocation of shared resources such as the consortia's wide area bandwidth

Formally Handling The Cost of Network Consortia Wide Area Bandwidth

- If you let bandwidth use run unchecked and just bill partnrs for actual usage, one or more partners may end up using all or most provisioned bandwidth (to the irritation of the other consortia members, and to the financial horror of the pigish partner)
- Trying to provision consortia capacity to meet that varying load is also pretty tough

An Alternative Approach

Let the consortia partners specify how much Internet capacity they want ("12 Mbps")

Charge them for that amount, whether they use it or not ("please pay us $9,000/month")

Reserve that amount of capacity for them

Drop traffic in excess of the contracted rate

Allow the partner to ratchet traffic upward (at the partner's discretion) ("we want 15 Mbps now and will pay you $11,250/mo")

Exempt local traffic from charges

In Fact You Can Also Probably Exempt Still More Traffic...

If your consortia members are like most sites, they will tend to consistently use more inbound bandwidth than outbound bandwidth. Since one cannot provision inbound bandwidth separately from outbound bandwidth, inbound demand ultimately determines the bandwidth which must be provisioned, and thus, as long as outbound traffic doesn’t get *too* hot, it can also be left unmetered

You Can Also Exempt Traffic to Motivate Desired Behaviors

You can also do interesting things with consortia bandwidth pricing to incent desired behaviors. For example, if you assume that some large percentage of all network traffic is http traffic, you might exempt from charges all http traffic that comes through a consortia partner's web cache, particularly if that web cache is part of a consortia-wide cache hierarchy.

Exempting Traffic to Motivate Desired Behaviors (cont….)

Web caching will be particularly important if consortia partners do not ratchet up their bandwidth to meet observed growth in inbound bandwidth demand. Exempted web cache traffic would thus run free and unconstrained, while users who elect to access web pages directly would get slower and slower page load times due to consortia traffic caps applicable to that partner

A Consortia Can Support Resources That Indvidual Sites May Not Be Able to Afford

For example, a consortia may run a central Usenet newsfeed box, taking inbound articles on a central system and then fanning them out so there aren’t multiple parallel/identical Usenet feed streams running over the consortia’s constrained inbound links.

Consortia and Content Delivery Networks

Another example of a shared resource valuable to consortia members, and one probably not directly obtainable by non-aggregated sites, is local access to content delivery network servers such as Akamai's. Oh, even if you're an independent, Akamai will still send pages to you, it is just that they will come in over your (paid) Internet transit rather than from a local (free) box.

Internet2 Access, Another Example of a Shared Resource

Another example of a consortia shared resource is access to high performance research and education networks such as Abilene/Internet2 (www.internet2.edu). While individual colleges may find it difficult to underwrite Abilene connectivity (the smallest Abilene pipe starts at OC3 speed, 155Mbps), a consortia of multiple colleges can share those costs.

I2 Sponsored Educational Group Participants

Participation in a consortia/statewide education network is effectively required for educational groups which would like to be connected to Internet2 as a Sponsored Educational Group Participant (SEGP), as Oregon's OPEN is… (For more information about SEGPs, see: http://www.internet2.edu/abilene/html/faq-sponsored.html)

So Should You Join A Network Consortia of Some Sort?

YES (no ifs, ands or buts about this one)

IX. What’s Still Ahead

A Confession

I might as well “come clean.” At the start of this talk, I said, “If you’re happy with what you’re currently doing, by all means keep doing it, more power to you. Don’t mess with success.” That was actually a lie.

Irresistible technological and financial forces are gathering, and like it or not, you will almost certainly have to make some changes to your wide area networking.

The Coming Theme: <u>Fast and Cheap (F&C)</u>

F&C powerful commodity PC hardware

F&C local area hundred megabit and gigabit networking hardware

F&C high speed remote access (DSL & cable modems)

F&C specialized wide area connectivity (I2)

F&C commodity connectivity (e.g., Cogent Communications, Yipes, Telseon, et. al.)

F&C metropolitan fiber networks

...Clashing With Some Crucial Choke Points That Aren’t Changing [Fast Enough]

- Essentially invariant commodity T1 pricing
- A lack of readily available wide area network engineering talent
- Old approaches to new P2P apps
- Many sites are connecting to Internet2 and/or deploying private fiber or doing other advanced networking initiatives … but many more sites are not.

Fast And Cheap Commodity PC Hardware

Virtually any current commodity PC can easily source tens of Mbps worth of traffic on a sustained basis (e.g., a student's five hundred dollar hand-me-down PC can saturate your wide area T1 or 2xT1).

With relatively trivial tuning, that same PC can easily saturate a DS3 (45Mbps).

Local High Capacity Disk Is Becoming Dirt Cheap

40GB IDE drives are available for less than $100, which implies that students (for example) can build substantial disk arrays to provide “content” for their PCs to serve. Do not be surprised when you bump into students with half-terabyte or terabyte disk arrays within a year.

Network Cards Are Getting Faster and Cheaper

Fast ethernet (100Mbps) ethernet cards are now down to less than thirteen bucks each, quantity one (e.g., Netgear FA311TX from http://www.warehouse.com after $5 mail in rebate)

Gigabit over copper (1000Mbps) ethernet cards are now less than $120.00 each (e.g., Addtron AEG-320T)

And Even Ethernet Switches Have Become Absurdly Cheap

HP 4000M's with 40 10/100Mbps ports have been routinely available at less than $1300 after rebates (e.g., $32.50/fast ethernet port)!

Given the price of NICs and switch ports, from a LAN point of view, there's really no reason NOT to engineer your local network to go fast… except that you then slam that LAN firehose into a WAN drinking straw.

DSL Service and Cable Modems Have Displaced Traditional Dialup Modems

Students and faculty with 1/2 to 1/3 of a T1's worth of DSL service at home may not have their expectations me sharing a single T1 at work with hundreds or thousands of other users

Cable modem service (up to 8Mbps on an asymmetric basis) for ~$40/month may REALLY spoil users for T1 or 2xT1 service

I2 Is Connecting More Users & Is Getting Ever faster, But There Are Many Non-I2'ers

The I2 SEGP program is greatly increasing the number of users connected to Internet2

A growing number of foreign research and education networks peer with Internet2

I2 is no longer selling any connections slower than OC12 (622 Mbps) [although old OC3's have been grandfathered in]

==> A new "digital divide" will hit independent colleges particularly hard...

<u>Absurdly</u> Cheap Wide Area Bandwidth Is Now Available … At Least In <u>Some</u> Cities

In the top 20 markets (basically, the so-called "NFL franchise cities") you could purchase 100Mbps worth of ethernet provisioned commodity transit from Cogent Communications for $3K/month, or 1Gbps worth of commodity transit for $10K/month (see: http://www.cogentco.com/)

See also Yipes, Telseon, etc.

Yet The Price Little Guys Pay For T1s Isn't Dropping...

As you look at everything that's getting faster and cheaper, it is ironic that one thing that's not getting faster and cheaper is the T1 (or pair of T1s) that you probably rely on for wide area connectivity.

Beyond T1s, DS3 and fractional DS3s have gotten cheaper, but they still aren't cheap, and greater than T1 speed local loops aren't available everywhere in Oregon anyhow.

Maybe The Hardest Choke Point to Overcome: The Shortage of Network Engineers

You may find that you face an insurmountable obstacle: competent wide area network engineers are hard to find and hire at <u>any</u> price… yet without good network geeks, you're going to be hard pressed to execute any sort of advanced network strategy… (beware hired guns who build it, turn it on, get paid, and leave)

And Meanwhile, Back At the Ranch: P2P Apps

Your users continue to discover new bandwidth intensive peer-to-peer file sharing applications like Audio Galaxy (www.audiogalaxy.com) or Kazaa (www.kazaa.com) or EDonkey2000 ("harness the power of 2000 electronic donkeys!"), and as P2P applications start to port hop and encrypt their traffic, you will find it harder and harder to know what's up

And No, I Don’t Think Network Policing Boxes Are the Answer to P2P Apps

Any automated network policing box which you can install to detect and limit P2P apps will eventually be overcome by smarter and smarter P2P apps (so don’t even go down that P2P “arms race” road).

Automatic app-neutral traffic shapers don’t fix the problem you have, either -- they just hide the symptoms and make it harder to detect and treat the underlying problem

And All The While, Users Are Probably Telling You… It's Too Slow!

And you know what? They're right. The world has changed and is still changing. If you're doing the same thing now that you were doing ten years ago, you're in trouble.

If money is no object and you live in the right place, you MAY be able to buy your way out by just throwing money at the problem (e.g., replace your T1s with DS3s). But in most cases, you'll need to work hard.

X. So What Should You Do?

Our Top 10 Recommendations (In No Particular Order)

If you aren't part of a statewide network consortia yet, join one.

Track what's going on. You need to be monitoring your network usage at least with MRTG or a comparable "big picture" macroscopic network monitoring tool.

Plan to spend more on wide area networking. [As a rule of thumb, plan to spend as much for wide area networking as you do for telephones or electric power.]

Our Top 10 Recommendations (continued…)

Plan strategically. For example, you know you will need fiber within a few years, and since it takes time to find/obtain access to an asset like that, begin working on it now.

Hire a competent wide area network engineer. Yes they are hard to find and yes they are expensive, but they are worth their weight in gold. You need one. Hire one.

Our Top 10 Recommendations (continued…)

Make yourself use and **depend** on the network. When something breaks, find out why and make sure it gets fixed.

Act like one of your users. If you can get DSL or cable modem service at home, do it -- they all are. How does your service compare? Try new peer to peer applications yourself -- that's the only way you'll understand their allure for your users.

Our Top 10 Recommendations (continued…)

- Figure out one new thing to offer every term. Web email. 802.11 wireless. PalmOS support. IP multicast. IPv6. Whatever… conciously force yourself to evolve your network service offerings. Don't stagnate.
- Make sure that whatever you do, your approach scales well. Don't be the victim of your own success. Assume users will love whatever you give them, and want more.

Our Top 10 Recommendations (continued…)

- Do it yourself, locally. If you start down th "lets outsource" road, or if you settle for folks "doing for you" elsewhere, you're on the slippery slope to hell. Develop and cherish local expertise and your local infrastructure. It will matter in the long run Zen koan for future meditation: reconcile "doing for yourself, locally," with being "p consortia." (This one **does** have an answer

Thanks for the chance to talk today!

Questions?