Abundant Bandwidth

and how it affects us?

More Questions Than Answers

Tal Lavian tlavian@eecs.Berkeley.edu

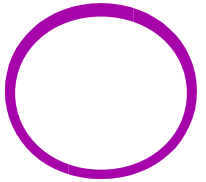

The Light at the end of the Tunnel

Our Networking Beliefs

- Let's challenge some of our networking beliefs
- Let's be a networking agnostic or skeptic for a moment
- Sorry…. I know it's provocative
- I could be wrong, but it's fun to challenge!

Agenda_

- Optical Internet & abundant bandwidth
- The economic factors (cheap bandwidth)
- Do we need protocol change?
- Do we need architectural change?
- Where are the bottlenecks?
- Summary

Abundant Bandwidth

Why does this change the playground?

- Optical core bandwidth is growing in an order of magnitude every 2 years, 4 orders of magnitude in 9 years
  - 1992 – 100Mbs (100FX, OC-3)
  - 2001 – 1.6Tbs (160 DWDM of OC-192)
  - OC-768 (40Gbs) on single ! is commercial (80Gbs in lab)
- 2-3 orders of magnitude bandwidth growth in many dimensions
  - Core – Optical bandwidth - (155mb/s ! 1Tb/s)
  - Core Metro – DWDM optical aggregation – (2.4Gb/s ! N*10Gb/s)
  - Metro – Access for businesses (T1 ! OC3, 100FX, 1-Gb/s)
  - Access – Cable, DSL, 3G – (28kb/s! 10mb/s, 1.5mb/s, 384kb/s)
  - LAN – (10mbp/s ! 10Gbp/s)

Why Does This Matter?

- How do these photonic breakthroughs affect us as researchers?
- This is a radical change to the current internet architecture
- The WAN is no longer the bottleneck
  - How congestion control/avoidance affected?
  - Why DiffServ if you can get all the bandwidth that you need?
  - Why do we need QoS?
  - Why do we need cache? (if we can have big pipes)
  - Where to put the data? (centralized, distributed)
  - What changes in network architecture needed?
  - What changes in system architecture needed?
  - Distributed computing, central computing, cluster computing
  - Any changes to the current routing?

Our Concept of the Internet

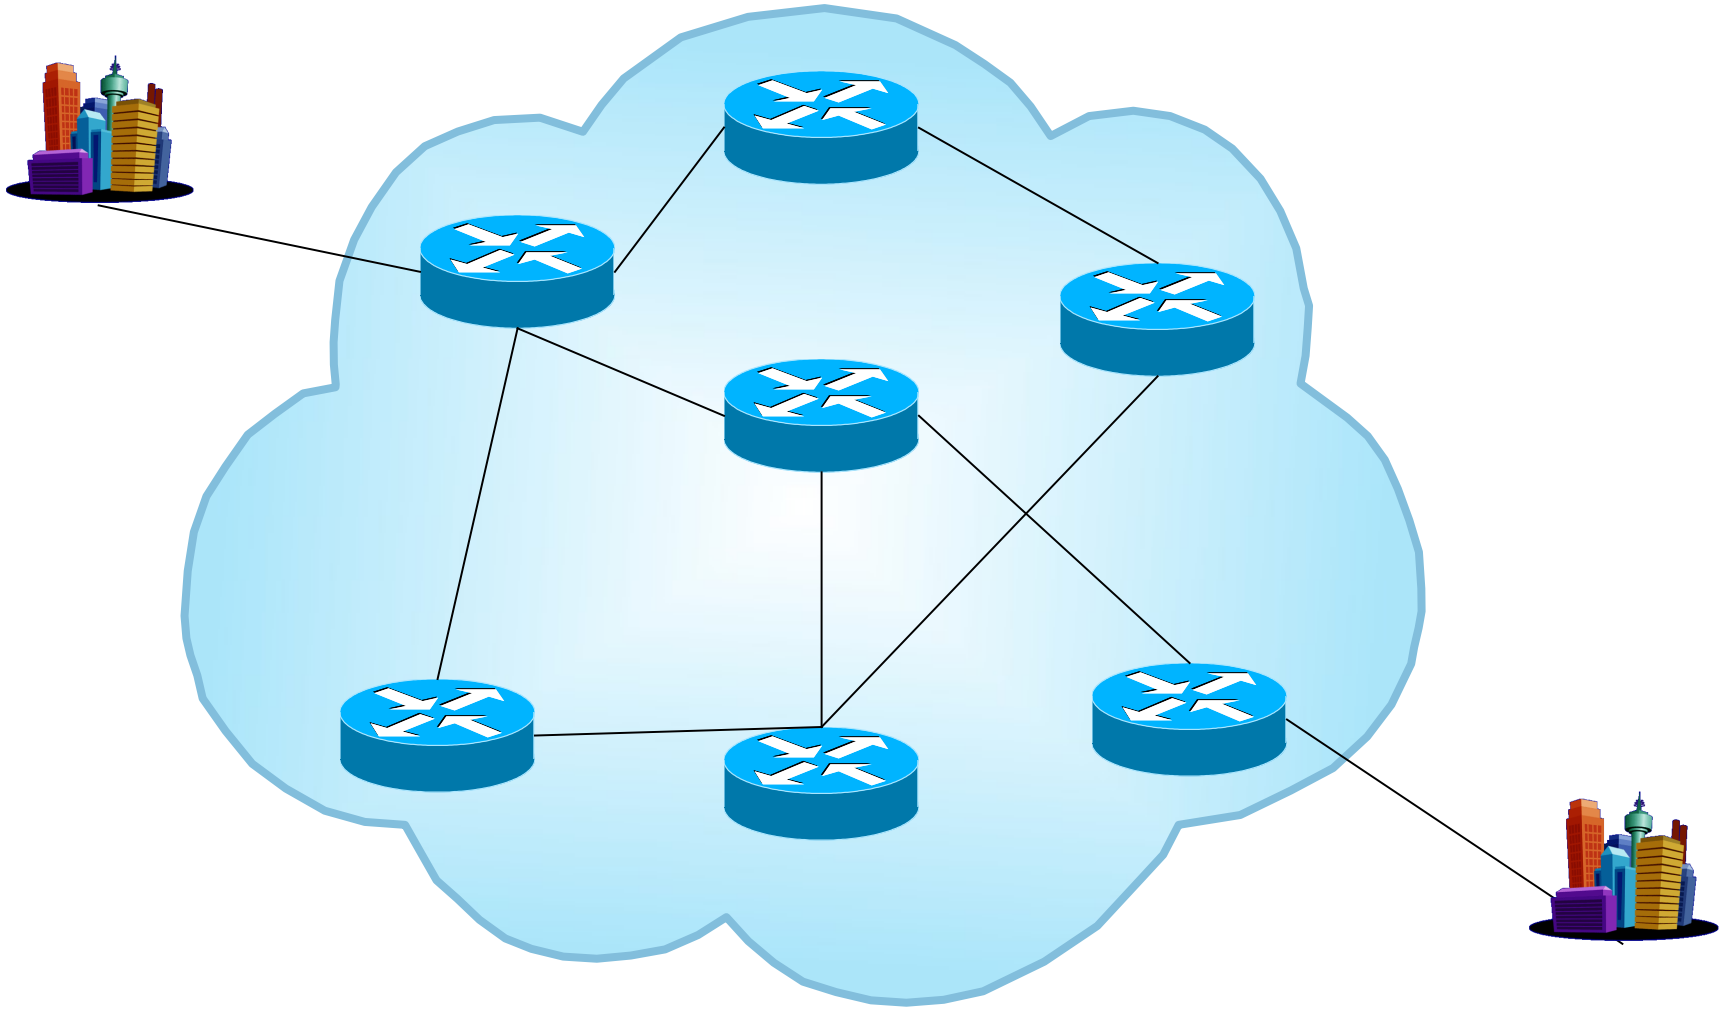

Internet Reality

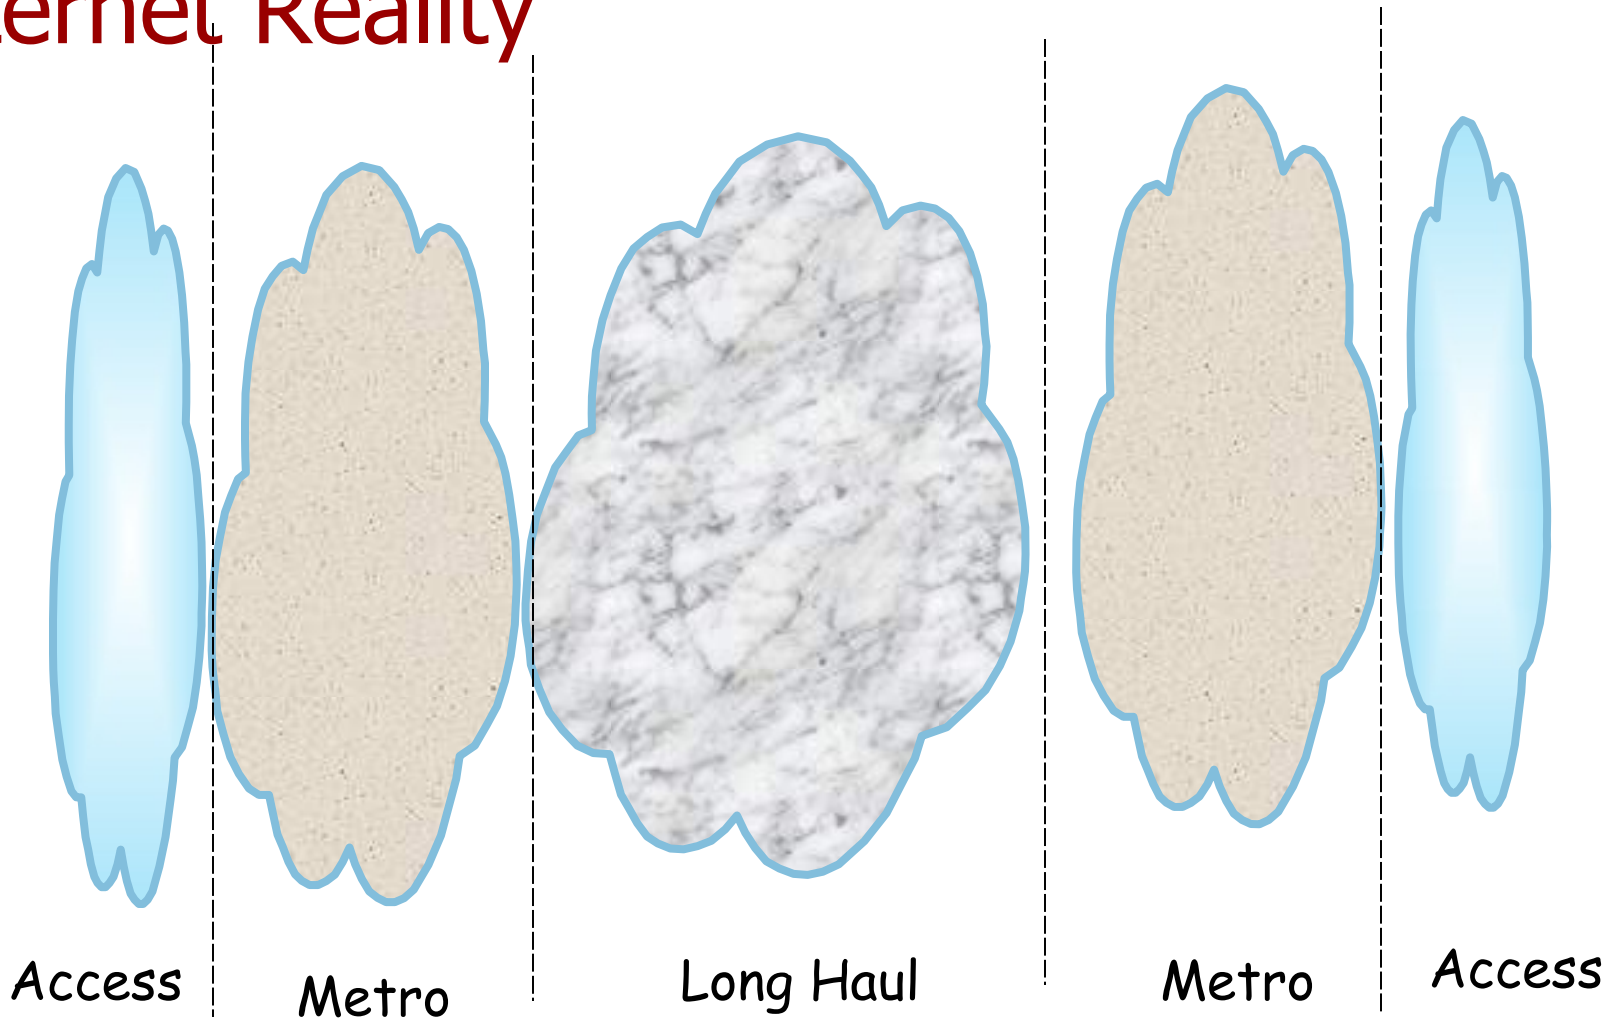

Internet Reality

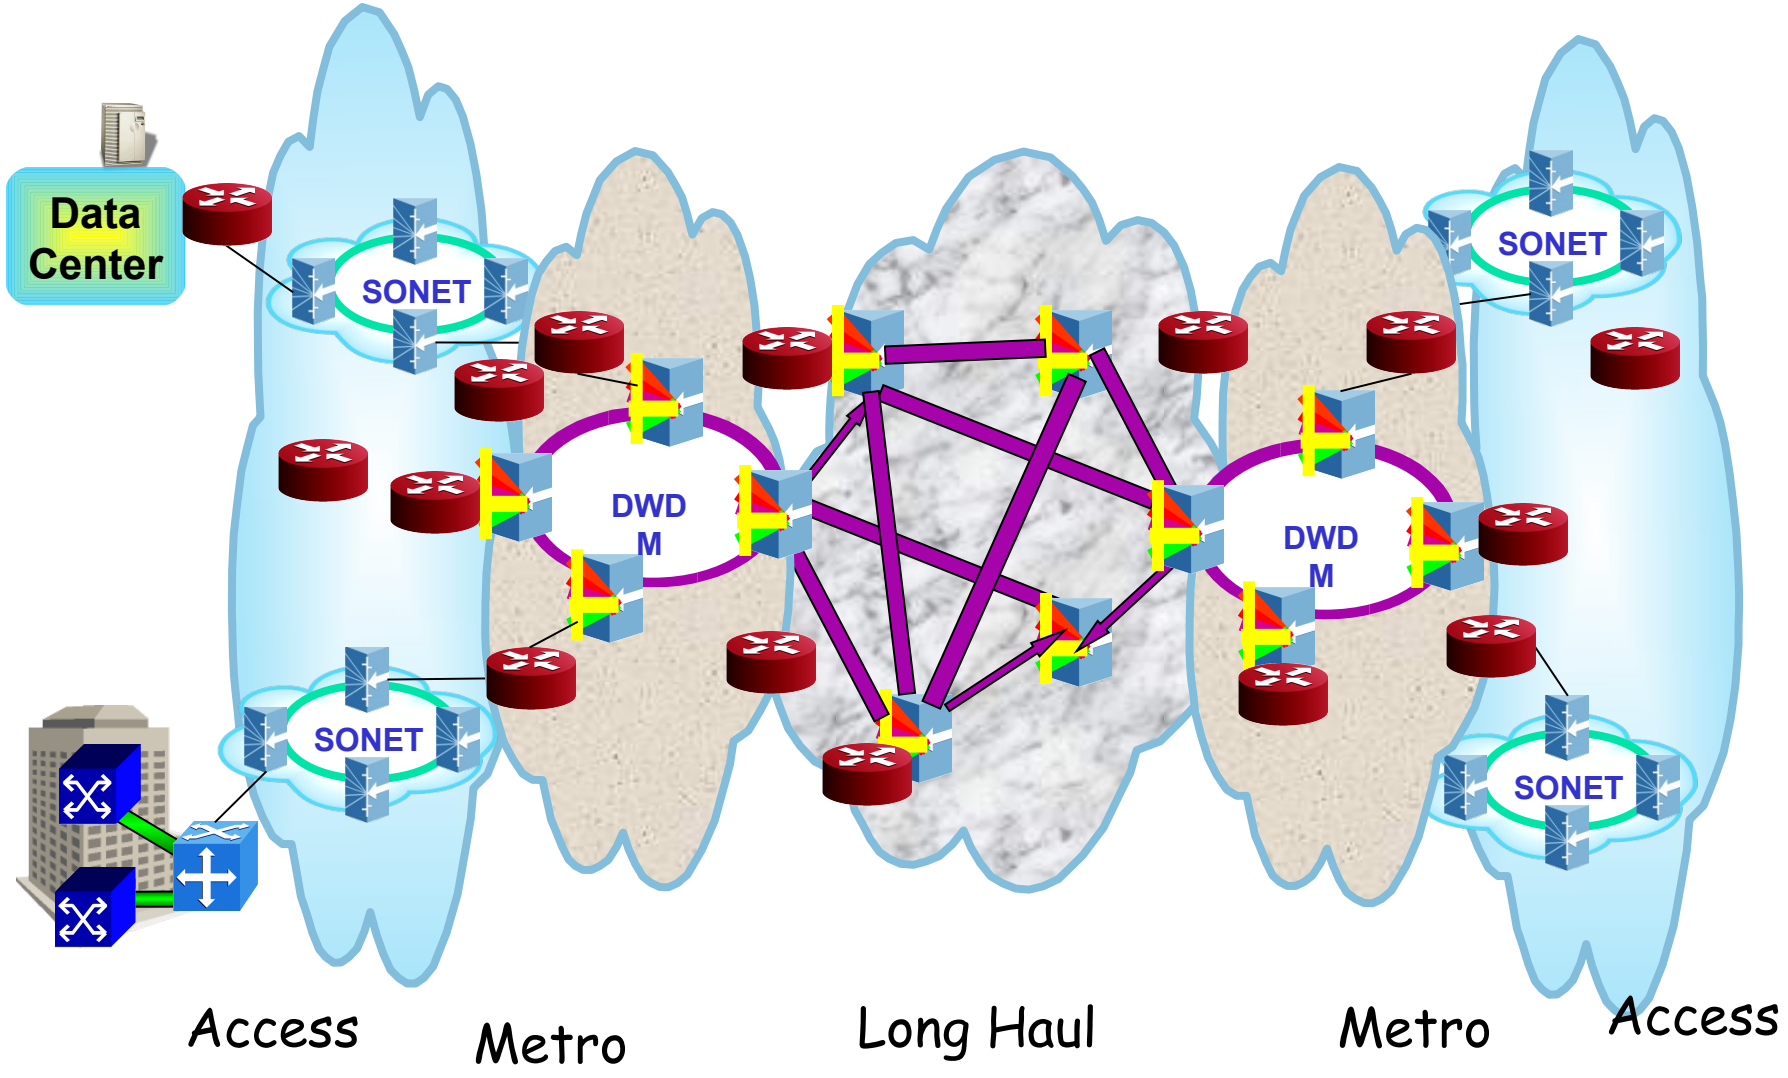

How Does this Affects our Lives?

- What are the new applications to use this abundant bandwidth?
- Distance learning?
- Telecommuting? (for the average person, not us)
- Broadcasting?(I want to see TV channel 48 from Japan)
- Video conference?
- What else? (this is a BIG question)
  - What are the new applications and services?

Fast Links, Slow Routers

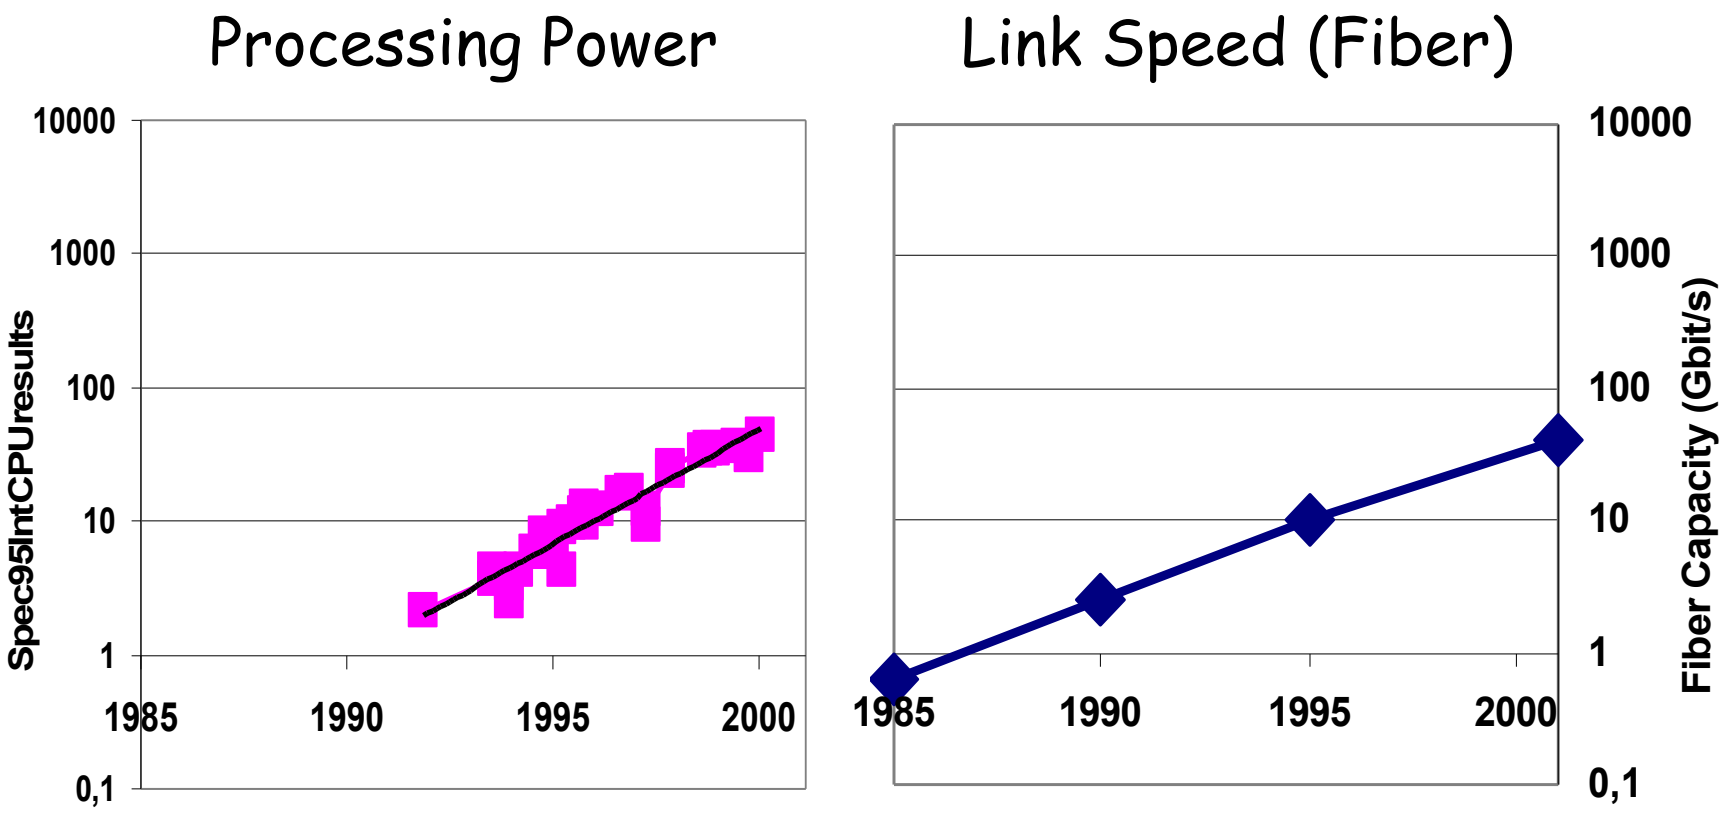

Source: Nike McKeown, Stanford

Fast Links, Slow Routers

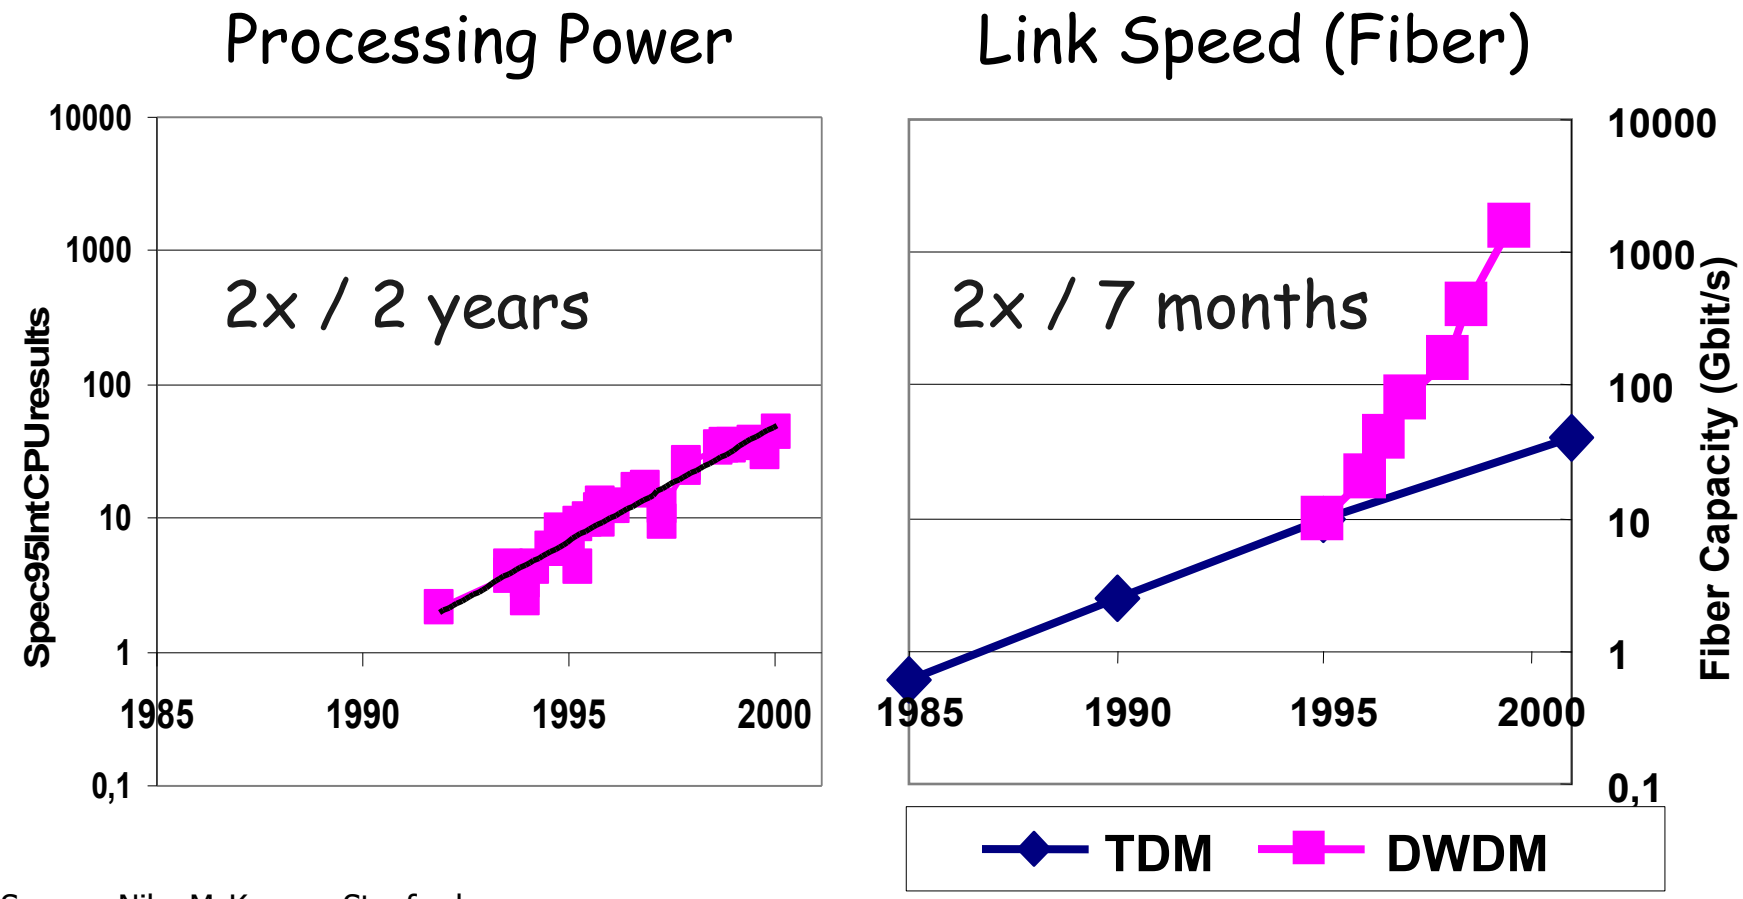

Source: Nike McKeown, Stanford

Agenda_

- Optical Internet & abundant bandwidth
- The economic factors (cheap bandwidth)
- Do we need protocol change?
- Do we need architectural change?
- Where are the bottlenecks?
- Summary

Breakthrough...Bandwidth

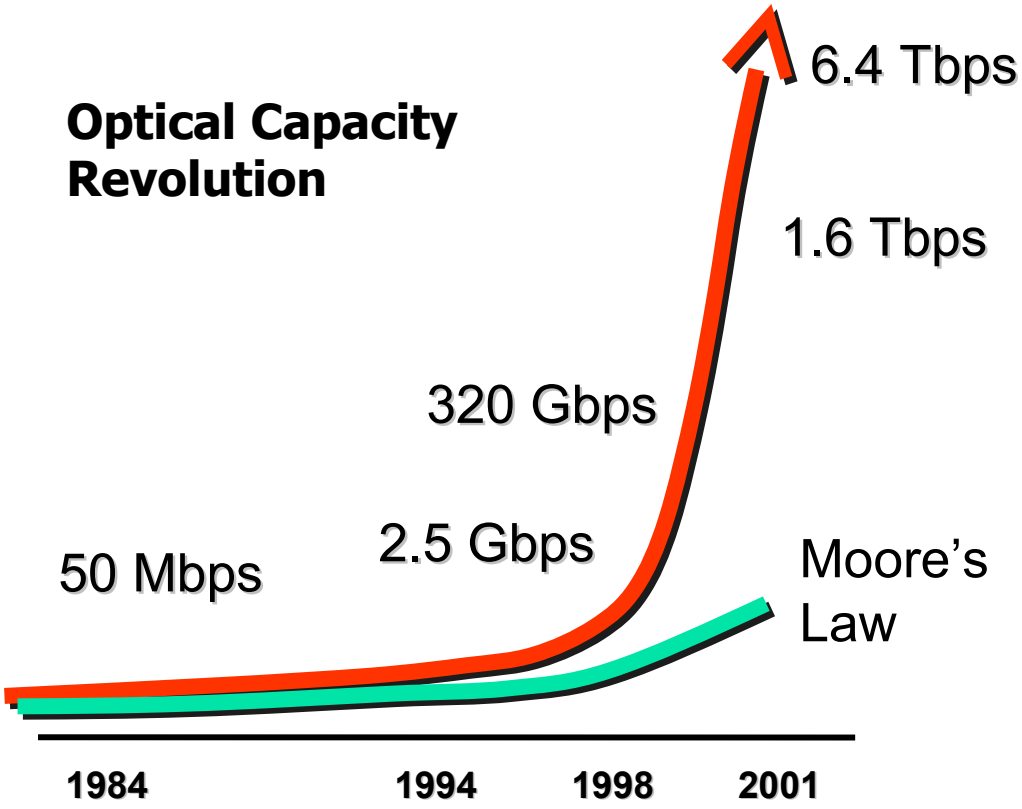

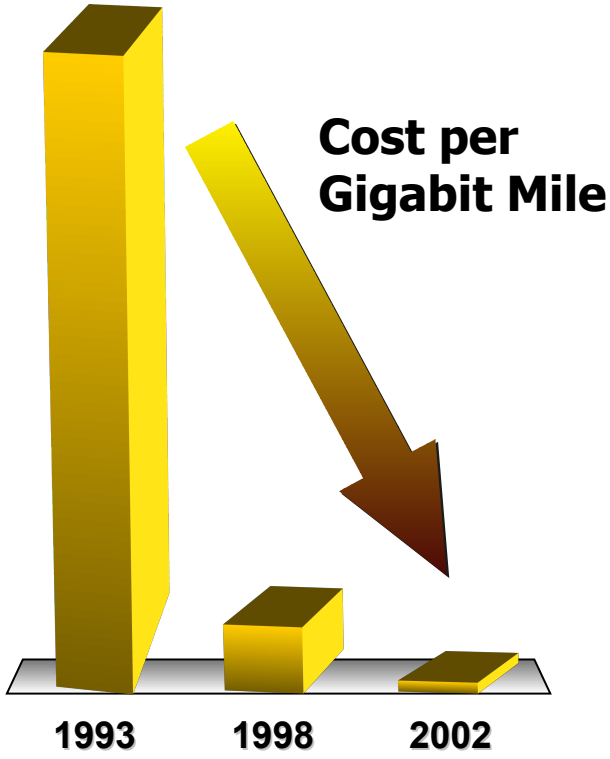

Wavelengths will become the communications circuits of the future...

Monthly Charges

- Current Connectivity
  - UUNET – OC12, $75- $140K
  - Sprint – OC12 - $78k
  - AOL – OC3 – $20k
  - XO – T1- $1500
- Current dedicated connection
  - OC3 SF-NY – $340k ($4M a year)
  - Only limited organizations could afford it
  - Optical bandwidth is changing dramatically

Bandwidth is Becoming Commodity

- Price per bit went down by 99% in the last 5 years on the optical side
  - This is one of the problems of the current telecom market
- Optical Metro – cheap high bandwidth access
  - $1000 a month for 100FX (in major cities)
  - This is less than the cost of T1 several years ago
- Optical Long-Haul and Metro access - change of the price point
  - Reasonable price drive more users (non residential)

Optical Ethernet

- New technologies are much cheaper
- Ethernet as the WAN access for businesses
- Will be at home if it is cheap enough
  - Charlottesville Virginia has become one of the first cities in the country to build its own Optical Ethernet network with 40,000 residents and 18,000 university students

If we had the bandwidth…

- What if we all had 100Mb/s at home?
  - Killer apps, other apps, services
  - Peer-to-peer video swapping
  - Is it TV, HDTV, something else?
- What if we had larger pipes at businesses?
  - 1Gbs home office, 10GE/DWDM large organizations
- How would the network architecture look, if we solve the last mile problem?

Agenda_

- Optical Internet & abundant bandwidth
- The economic factors (cheap bandwidth)
- Do we need protocol change?
- Do we need architectural change?
- Where are the bottlenecks?
- Summary

Possible changes

- Network architecture changes
  - Network computation on Edge devices
  - New services on Edge devices
  - Servers and servers farm location
  - Applications that interact with the network
  - Load balance switches, content switches, and server farms
  - Optical SAN connect directly to the networks with no servers
- Service model changes
  - New economic factors
  - Bandwidth and access is cheap
- Transport protocol changes
  - New protocol between hosts and edge devices
  - New protocol between the two sides of edge devices
  - End-to-End argument between edge devices and not end hosts

Assumption Changes

- Is TCP the right protocol?
  - BIG MAN & WAN pipes
  - No optical queues, no optical buffers
    - Like circuit switching (and not packet switching)
  - Extremely low bit lost (10^{-15})
  - Extremely low delays
  - 100Mb/s on every desk
  - Ratio change (file size/pipe size). No time to fill up the pipe
- Are we sure that in a new technology, losing packets means congestion? What if this is not true?
- TCP was designed for packet switching while optical is close in its characteristics to circuited switching

Do We Need Protocol Changes?

If there are no queues, how TCP "slow start" helps us?

- How this fits to the sliding windows?
  - Why don't we start dumping packets at our link speed?
    - Most HTTP files are relatively small (few K's)
    - For 100KB file, no time to fills up the pipe
    - The max Wind size is 16 bit=64kb
    - For 1Mbs wind we need about 20 RTT
      - If RTT is 10ms --> 200ms.
      - What if RTT 100ms ? That's 2000ms!!
      - What if RTT 500ms? (Australia on a bad day)? That's 10,000ms!!!
      - But just burst at 100Mb/s link speed - is 10ms
  - Assuming that we need daily backup of 100GB over a 10GE line – Do we need the same TCP assumptions?
    - Just dump - about 100 seconds (and correct at once in the end)
    - TCP with very high bit lose (say 10^{-9}) - might be much longer
      - 10^{12}Gb/ 10^{9} = 1 thousand restarts

Agenda_

- Optical Internet & abundant bandwidth
- The economic factors (cheap bandwidth)
- Do we need protocol change?
- Do we need architectural change?
- Where are the bottlenecks?
- Summary

The Access Bottleneck

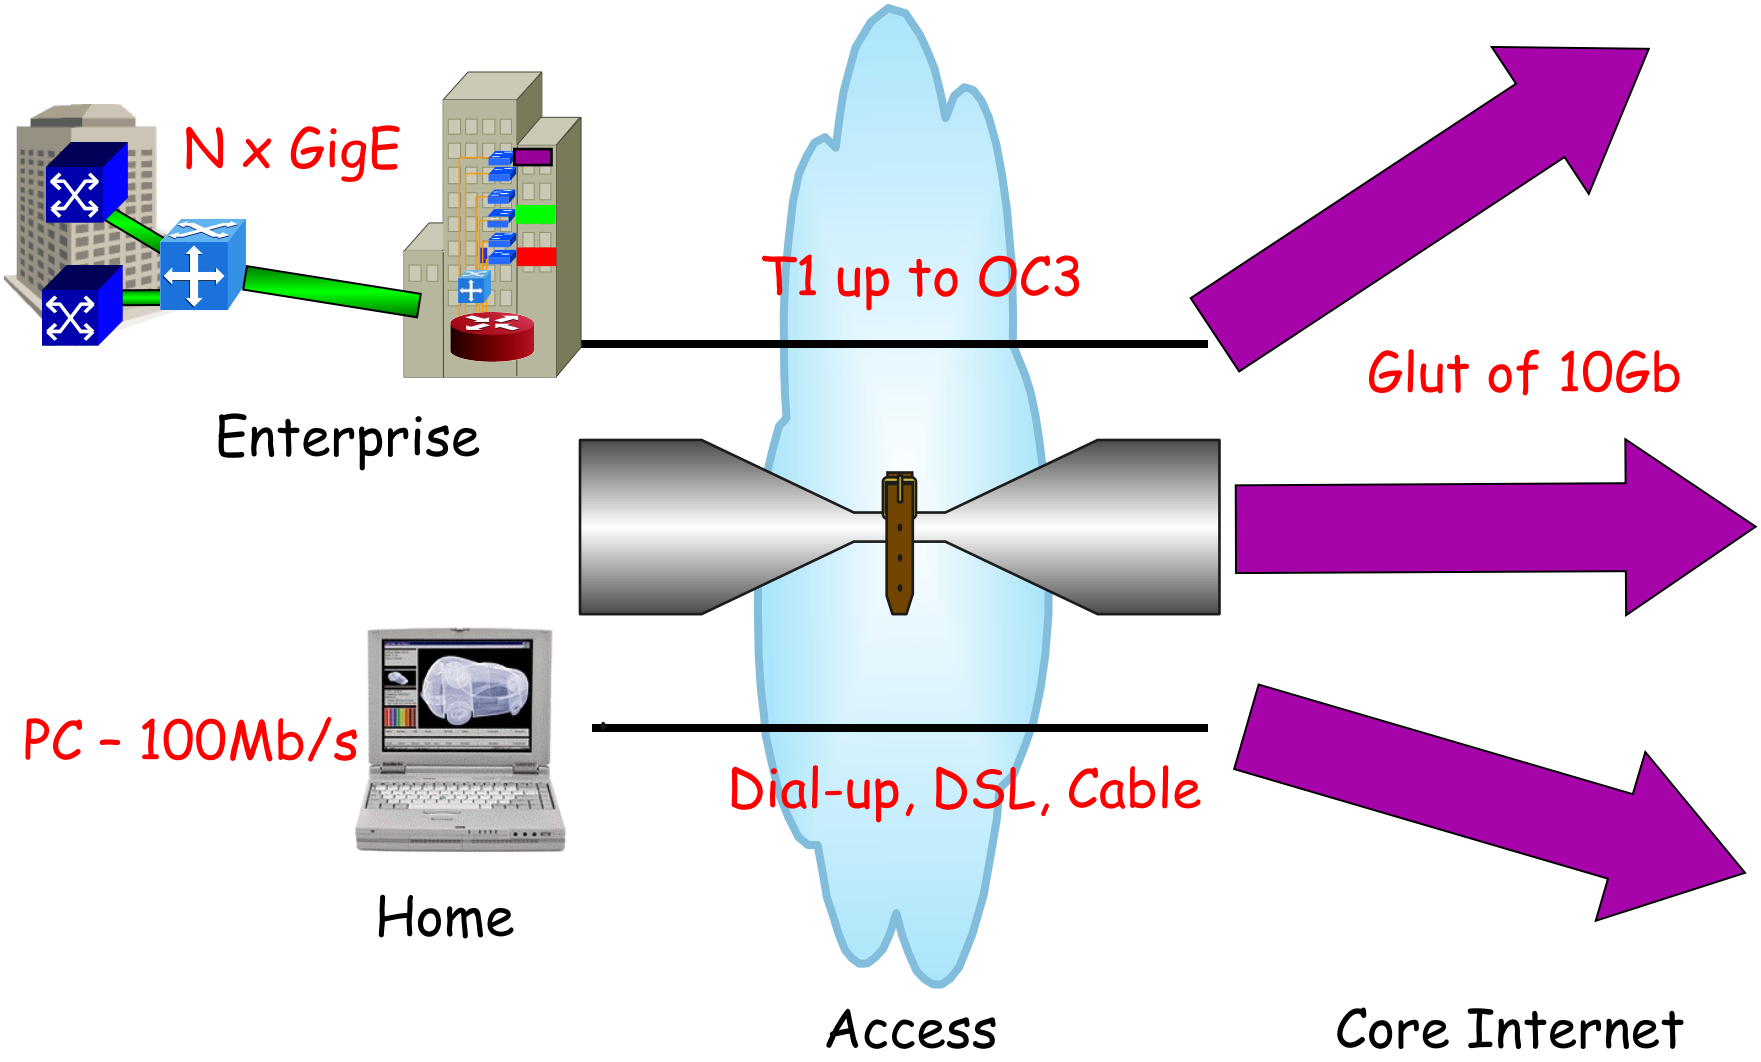

The Access

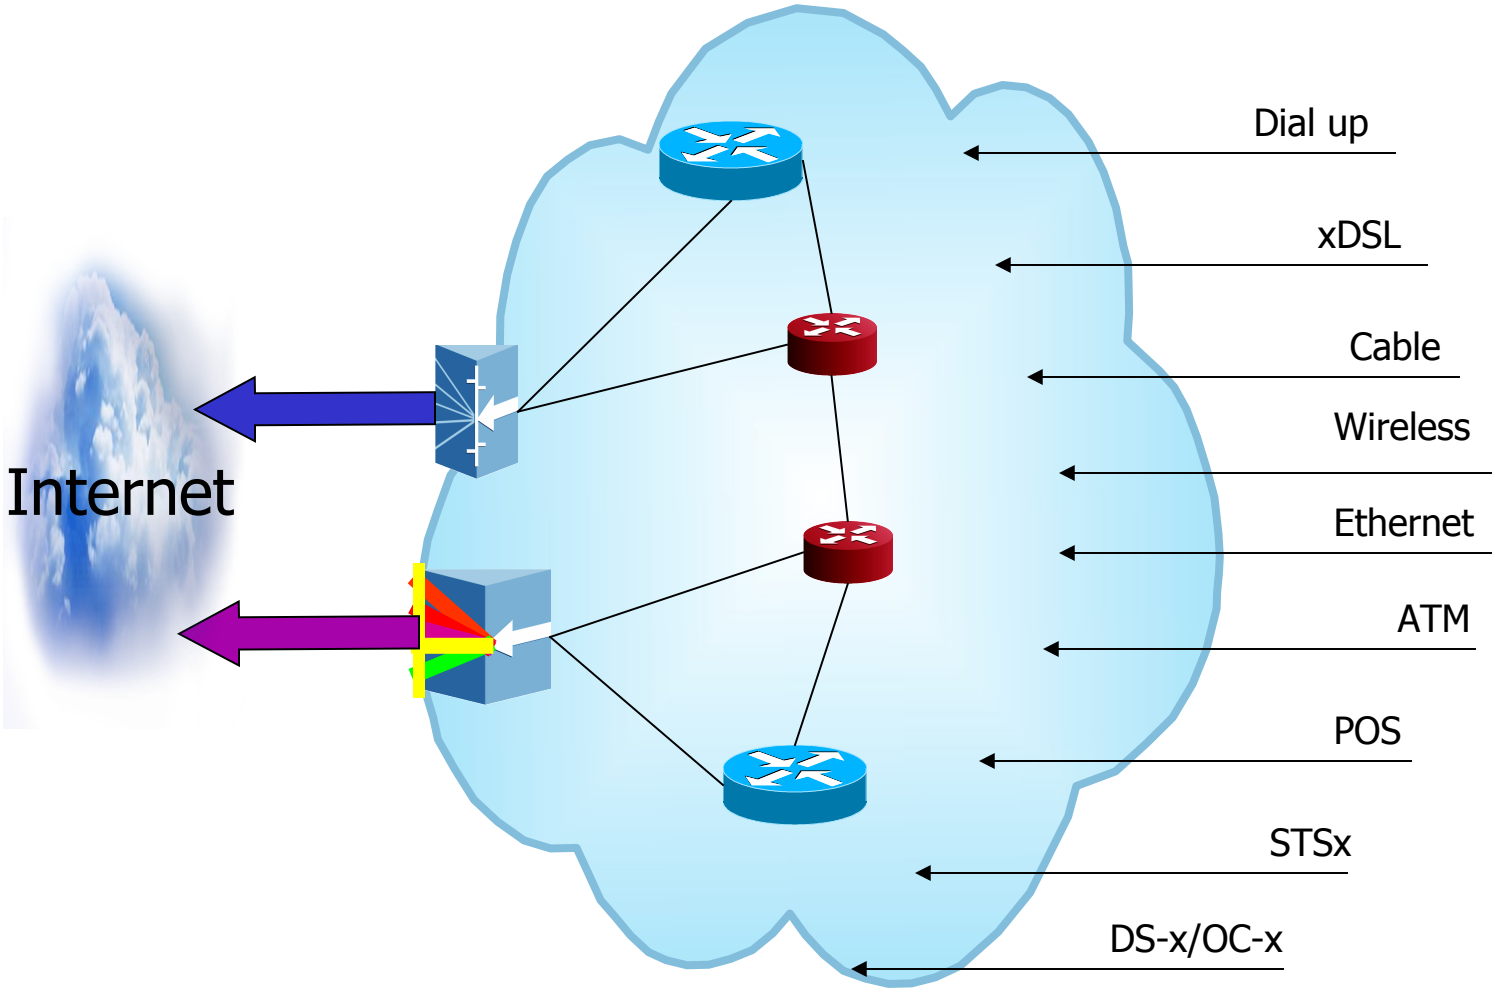
Internet
Dial up
xDSL
Cable
Wireless
Ethernet
ATM
POS
STSx
DS-x/OC-x

Access and Metro Networks?

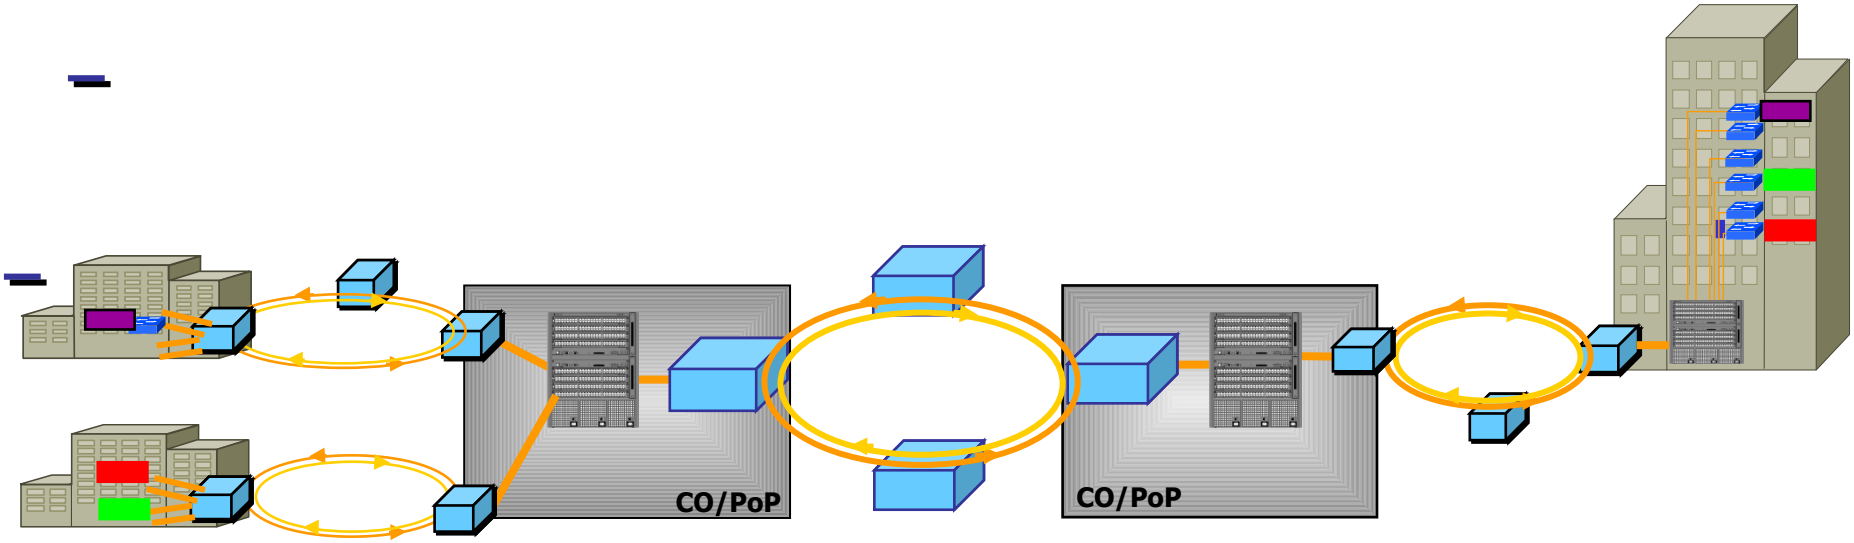

OC-3/OC-12 Access rings

OC-48/OC-192/Metro Backbone Rings

OC-3/OC-12 Access rings

Architecture Change

- End-to-end argument by the Edge instead of end hosts.
- Get some server functionality
- Services platform on the edge
- Overlay Networks
- Peer-to-Peer gateways
- Content Distribution Networks
- Load balance switch
- Bandwidth Auction – Weidong work

Services Platform on the Edge

- Can't do computation on the optical core
- Need to add the intelligence and the computation on the edge
- This might be a better place to add network services
- Services platform on the edge

Protocol and Services on Edge Devices

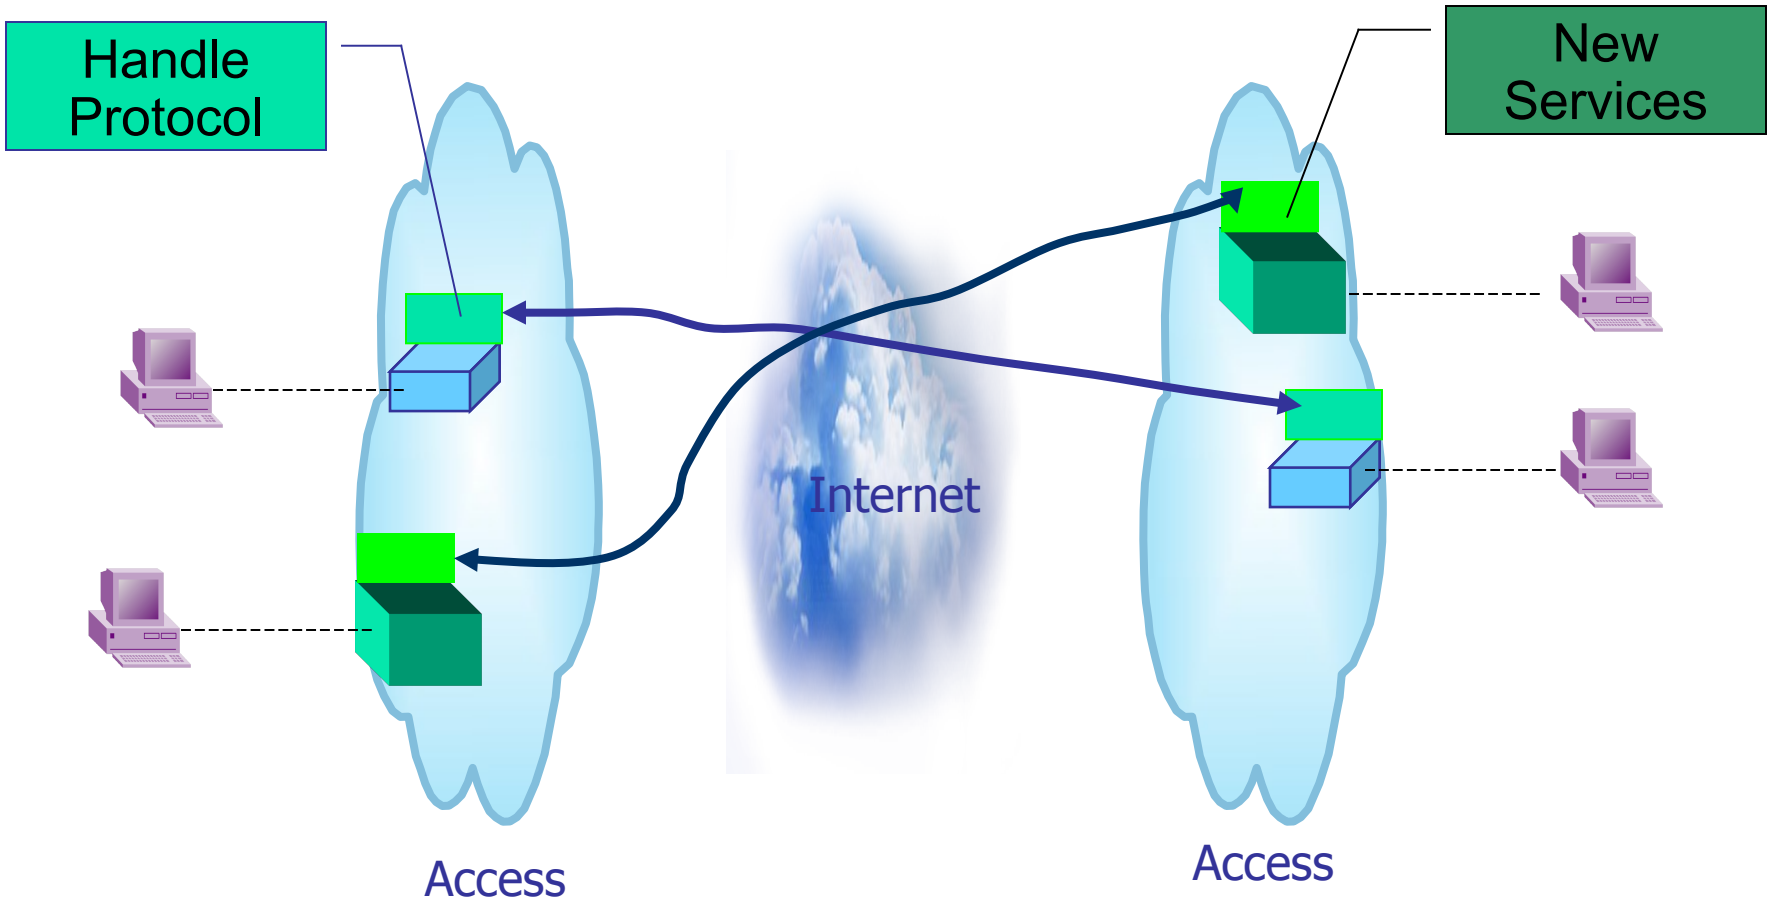

Bandwidth Trading

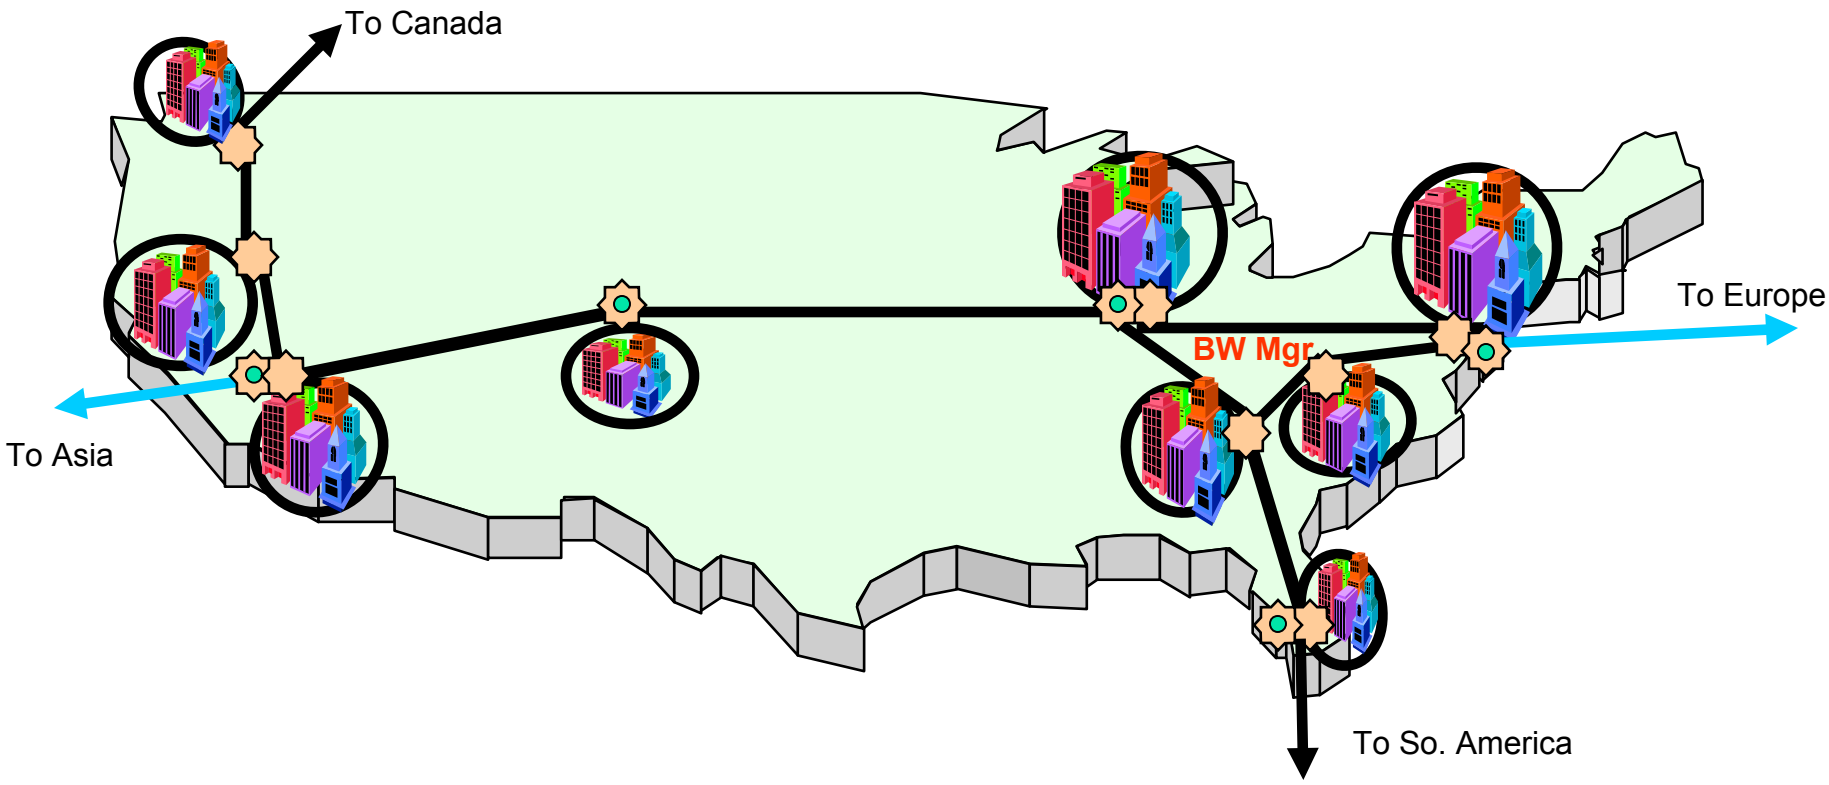

Agenda_

- Optical Internet & abundant bandwidth
- The economic factors (cheap bandwidth)
- Do we need protocol change?
- Do we need architectural change?
- Where are the bottlenecks?
- Summary

Where are the New Bottlenecks?

- Last mile? (for me it is the first)
- Aggregation routers?
- Between service providers?
- Between Metro and Long-Haul?
- Data centers? Clusters?
- Servers and CPU power?

Example of a new Bottleneck

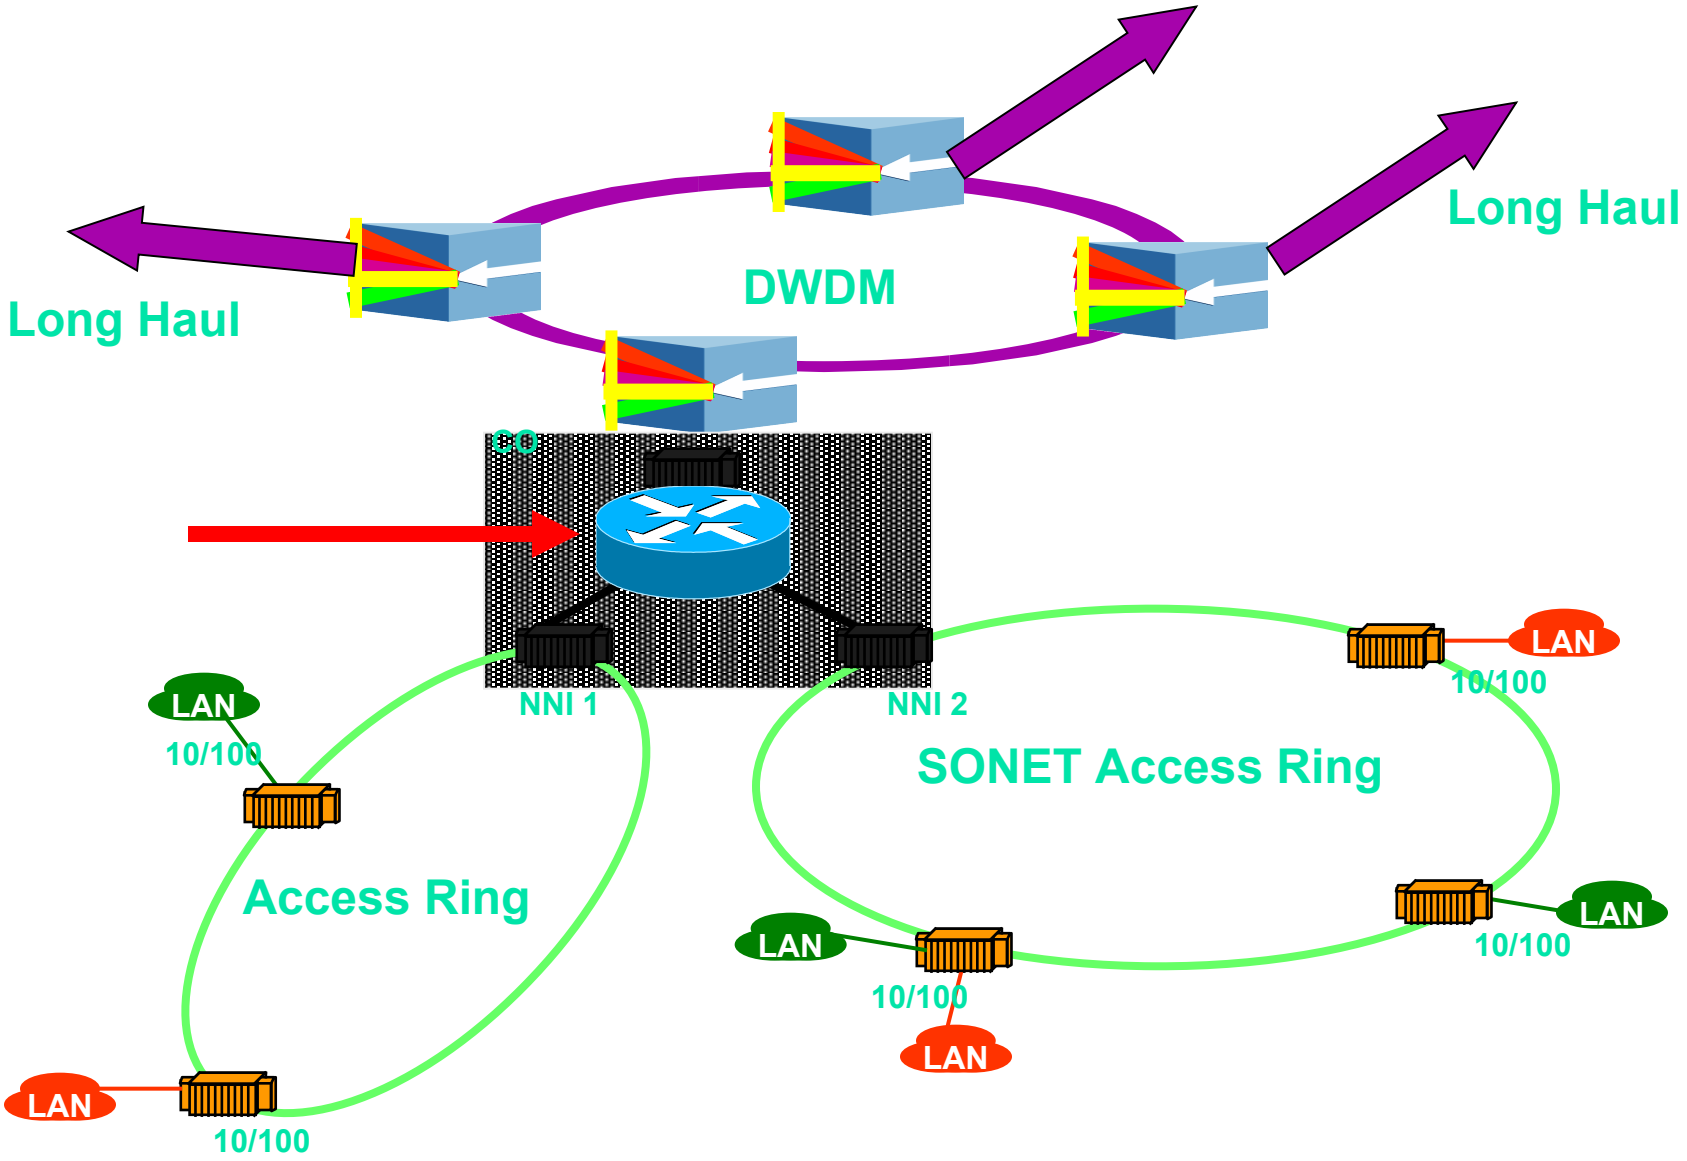

Example of a Bottleneck

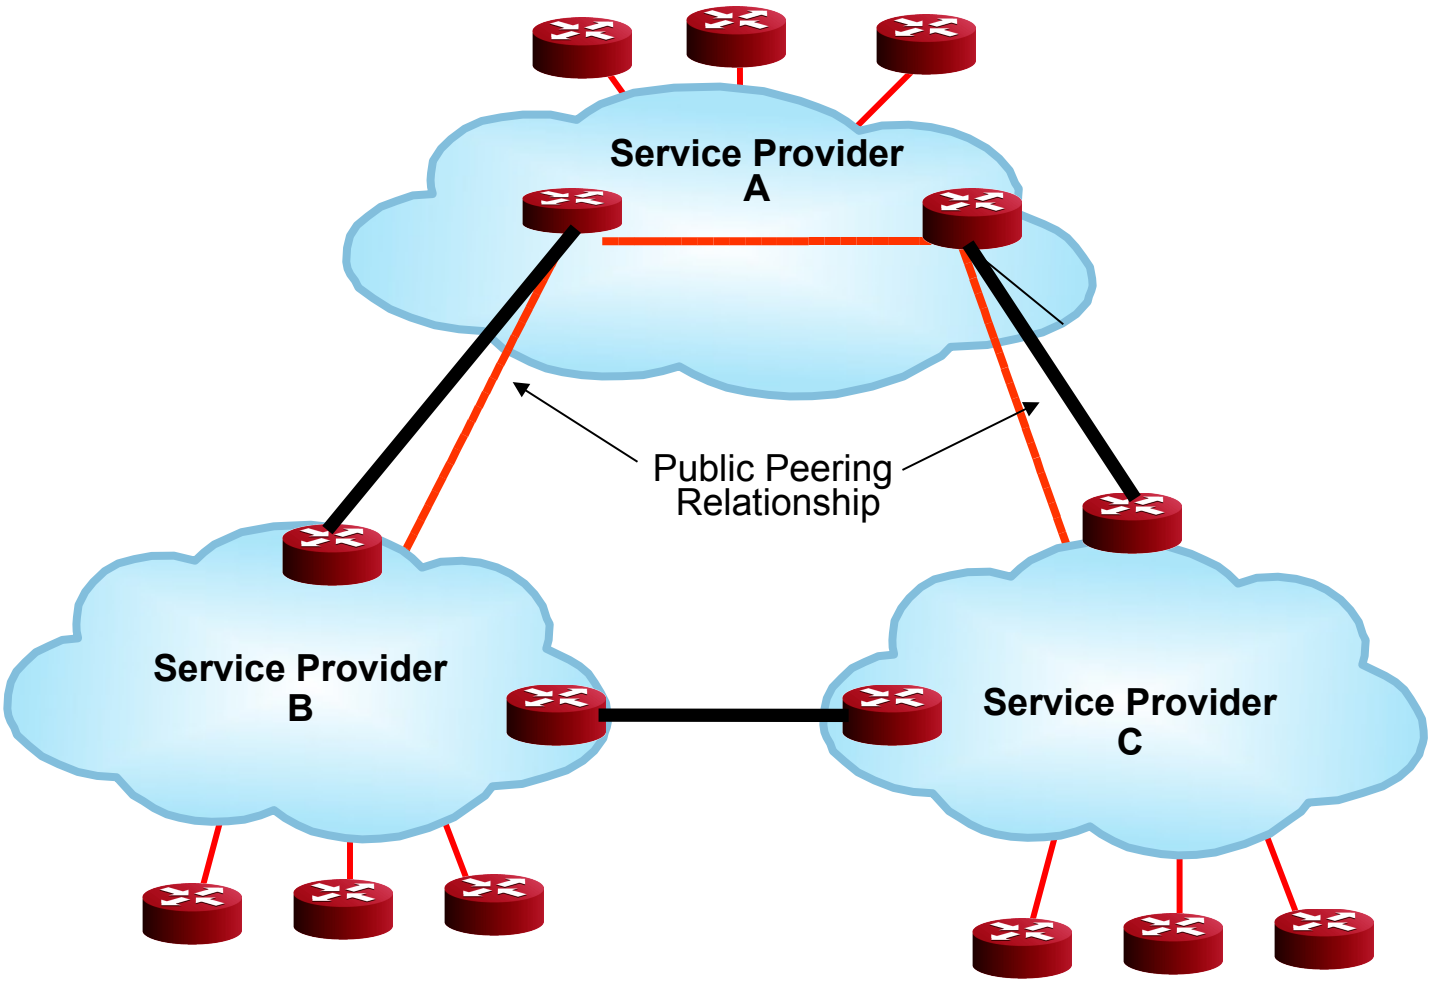

Open a Bottleneck

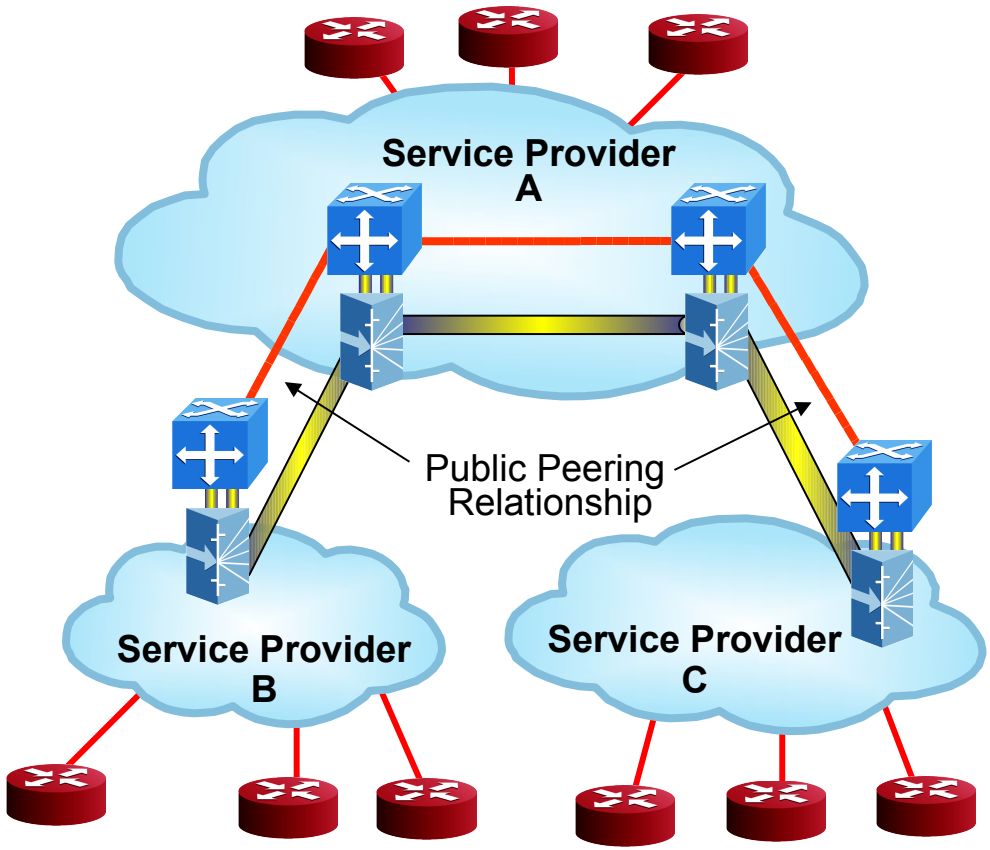

Open the Bottlenecks

- New products coming offer dramatic performance and capacity improvements that open some of the bottlenecks
  - Terabit Routers
  - Aggregation routers with optical output
  - Multipurpose boxes
    - Optical switch + IP router
    - SONET node + DWDM switch
    - SONET DCS + IP router
    - Long-Haul + Metro switch
  - Session switching vs. packet switching

Agenda_

- Optical Internet & abundant bandwidth
- The economic factors (cheap bandwidth)
- Do we need protocol change?
- Do we need architectural change?
- Where are the bottlenecks?
- Summary

Summary

- Disruptive technologies
- Optical Internet creates abundant bandwidth
- Dramatic changes in the cost per bit (99% in 5 years)
- Access is becoming cheap
- Opens several bottlenecks
- Need to rethink on architecture and protocol
- Our mission is to identify and build the services on top
- For most of the questions I simply don't know the answers

"Blindsided by Technology"

- When a base technology leaps ahead in a dramatic fashion relative to other technologies, it always reshapes what is possible
- It drives the basic fabric of how distributed systems will be built

Source – unidentified marketing

There is Light at the end of the Tunnel

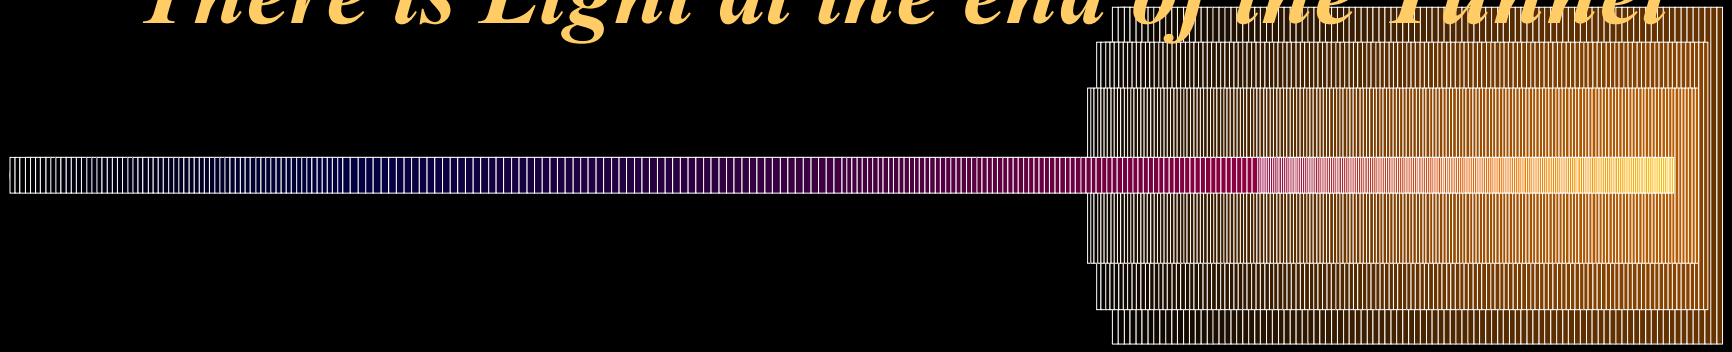

The Future is Bright

- **Imagine the next 5 years.**
- **There are more questions than answers.**