Posted to the Gilder forum - August 12, 2000

IP browsing with Lambda transport?

How do you reconcile IP browsing with Lambda transport? This is a fascinating topological problem.

Carl Russo, Cisco's Chief of Optical Networking has this to say (thanks BG!):

Well, let's back up. The network will never be -- certainly not in my lifetime -- all optical, because you will have to examine the packets to understand where they need to go.

You really have two segments of the network: One is the core, which hooks up all of the big switching centers into a big mesh, and the other is the metropolitan network, which backhauls all the traffic to those switching centers. The backhaul network will be optical, the core network will be optical, but, in between, you'll have a very big electrical interconnect, which is basically a switching [connection]. You'll come to the interconnect optically, and you'll leave the interconnect optically, but in between it will be handled electrically.

GG says that the only solution to the latency of light is storage and this is the reason he like Mirror Image so well. But Mirror Image does not solve the whole problem because it operates on IP and that implies electronic switching.

GG also says that the core of the net will be based on lambda transport but he does not say how the photons will move from the IP periphery to the lambda core.

Here is an idea that uses storage to replace switching:

If you divide the Telecosm into a Lambda core and an IP periphery, then it could work as follows. The core would be a giant party line or Ethernet like network made up of servers and caches. As soon as a server has new data it broadcasts this data to all caches meaning that all caches are up to date all the time (a super redundant storage array, not like Mirror Image which has only part of the server data). Client browsers never contact the servers but talk to the local (city?) cache only. This takes care of the download. The upload, being considerably less bulky, could travel on a parallel return all IP network from cache to server. In

this system, caches never interrogate servers but use the parallel return IP network to acknowledge receipt of the broadband transmissions.

This would save a lot of photon switching! What do you think?

Denny
"Demand creates queues. Supply gets rid of them."

PS: This was my starting point and I no longer believe that severs will "push" content to the caches. Dec 9, 2000