

Posted to the Gilder forum - December 13, 2000

Four Islands? Thirty Mirrors?

I have a bit of a short circuit that I need help with.

In February, GG told us that we need 30 or so Mirror Image caches to make the Telecosm work at an adequate rate. Mirror Image would get rid of the problem created by the latency of the speed of light by bringing the data closer to the end user.

Now, in December, the problem is Connectivity and this is solved by thousands of lambdas traveling in each fiber optic cable which, with the help of four optical islands (switching palaces) will connect all user to each other. What happened to latency? Did GG just forget about it, a momentary lapse, or is this something that is no longer of interest?

A few days ago GG cleared up the question a bit:

Cao was depicting the center of the Fibersphere--the so called all optical core of the network--as it might evolve over the next decade toward the theoretical nirvana of a nearly switchless lambdanet governed by the speed of light. The faster Cao's visions come true the quicker Exodus and Mirror Image become the crucial keys to the web. When nearly all latency on the net is lightspeed, Exodus centers on the edge become the optimal hubs of the Internet, balanced between the need for proximity to users and the need for sorting and searching scale for storewidth. Competitors such as Akamai sacrifice searching and delivery scale to an unnecessary proximity and incur major penalties in coherency as data changes.

--GG

My take is somewhat different. There are two ways of connecting people: client-server and peer-to-peer. Mirror Image caches do client-server and Cao's multiple lambdas and switching palaces do peer-to-peer. But they are not mutually exclusive, rather, they are complementary. My feeling is that the bulk of the traffic will be client-server and selected communications will be peer-to-peer, the ones where the participants want to be in touch in real time. If these

suppositions are true, then what are the consequences for the Telecosmic topology?

If the bulk of the traffic is client-server then the core will be lambda transport and the edge will be IP switching over Ethernet transport. The core does not need "connectivity" but fast bit rate to transport a lot of data from the servers to the caches. The edge, on the other hand, might benefit from connectivity, it would be a way to reduce the traffic congestion on the Ethernet transport. If this is correct, then what will change radically is the topology of metro, and local access and not the core, after all, all you need to do in the core is to interconnect 30 Mirror Image CAPs. The users and the servers are the edge.

Later, a reply to Sandy

Sandy:

The users are all at the extreme edge (the core has no users at all) and it is these edge denizens that need to be connected.

If you were to connect all users with each other, each additional user would need to be connected to all existing users and the Telecosm would require as many new lambdas for each new user as it has current users. (Pascal's Triangle gives the number).

But most communications are NOT going to be peer-to-peer but client-server so what you need to do is connect all users to their respective Mirror Image CAP. Now the number of lambdas is a heck of a lot less. You can still create peer-to-peer connections but instead of being direct, they would go through 2 Mirror Image CAPs.

User-MICAP-MICAP-User

To create a switchless core having 30 Mirror Image CAPs all connected to each other, all you need is 435 connections. Since these would have very high traffic densities and individual circuits are not required, you might as well use high bitrate connections instead of multi lambda connections. The choice would be based on cost since both can transmit a lot of data.

This thinking leads me to believe that Avanex will be an edge solution and that is why Cisco is so interested. The bulk of lambdas will be at the edge, not at the core. At the core you will have high bitrate or lambdas, whichever is cheaper.

Denny

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

[Software Times](#)