

# Next-Gen Services: The Internet or Around It?



## *Unconventional Wisdom*

*Tom Nolle, Founder and President, CIMI Corporation*

### Introduction

Formal packet networking has its conceptual roots in work by the Rand Corporation forty years ago. This was a time when data was exchanged in hundreds of bits per second and when a T1 line with 1.544 Mbps of capacity was considered enormously capable. Today, the average broadband consumer has more capacity than corporate sites had then, and the nature of the network has changed forever.

The technology implications of this change have been discussed here through much of this evolution. The business implications have also been discussed. What we propose to discuss here is something that's perhaps a mixture of the two.

The statement by Google that video threatened the Internet was a bit of a seismic shock, and despite the statement was challenged by those who were said to have made it (as shocking statements usually are these days), the fact remains that the shift in the nature of what gets exchanged across the network is rather likely to impact the network's nature.

A network's goal is to support its applications, and here we would note that "its applications" would consist of the stuff that's economically feasible and/or legally mandated. When the mission of the public network was to support pairwise communications, we developed a technology, topology, and business model that offered that. As the mission changed to provide access to millions of websites by a billion or more consumers, the network changed again, to the Internet model we see.

The mission of the future is content, video content in particular. The question is how that mission might change the network. Here, we'll look not so much at technology or business issues (though we'll dabble in both areas) as at the topology of the provider relationships.

### Three Topology Models

If we look at networks as topological relationships between providers, we can see that there are two basic models of relationships possible, and a third that represents a complex hybrid of the other two.

The model of networking that has emerged for public switched services and for common carrier data services is the **cascade model** where providers are interconnected with each other so that information flows along a path through a series of providers. The cascade model was inevitable for the PSTN because nearly every industrial nation started telecom as a regulated national monopoly if not a formal part of the government process (European providers were originally part of the "Postal, Telephone, and Telegraph" department of the government itself, from which the term "PTT" emerged). If everybody has their own national network and you are going to do cross-border calling, you support a cascade model.

The cascade model also emerged for the Internet, and in this case we'd assert that the model was based not so much on logic as on the financial model that emerged with commercialization. ISPs were valued based on customer count, and so they sought control of customers. Because like the PSTN, the Internet was valuable to the extent that it could support a broad range of relationships, interconnection of the ISPs was inevitable, and that interconnection created a cascade model.

Another model of relationships also emerged from the PSTN, the notion of a **parallel layer set** that offered several alternative choices. International Record Carriers (IRCs) were once the primary trans-border mechanism for connection, and some Internet players were backbone providers but not providers of customer access. Interexchange carriers in the US are also an example of this model.

The thing that seems to differentiate the two models is the question of explicit or implicit customer control. Access providers are more likely to create cascade models because access is tremendously capital- and operations-cost-intensive, and thus there are relatively few players who will operate there. However, it could also be said that the parallel layer set model has seen some diminution over time because low barriers to market entry have driven down incentives to play there.

The hybrids of these models have emerged primarily through business and regulatory transformation. Access providers who are permitted to play in the long-distance market create the notion of a combination of cascaded and layered models where the access provider owns local customers but also provides a cascade connection to other access players. This is often done through a “layer” of long-haul service. Most access incumbents have created long-haul extensions and thus have created a hybrid model of relationships.

The Internet has so far had less of this hybridization than the PSTN. One reason may be that there is no settlement on the Internet (bill and keep) and so interior players can't make money except from peering agreements, most of which are still settlement-free. An interior Internet player is also in competition with the notion of customer-owner players simply peering with each other and cutting them out.

What we think is happening now, or at least threatening to happen, is a development that might create a layered delivery structure in the Internet, or with IP services at least.

Access layering, meaning the existence of multiple credible access providers for the same customer set, is a competitive reality even if one doesn't consider the proactive regulation to encourage the practice here in the US and elsewhere. All people who are candidate owners for a given customer set are competitors. As we've said, the barrier to access layering has been cost.

Transport layering in the Internet seems at least possible if one considers the transport network not as a mechanism to support communications but as a mechanism to host experiences for distribution by the access network. The higher value on experiences means that there could be enough money on the table to build out a distribution network to provide them to users, even users that aren't under control of the experience host. That's what we've been calling the “portalnet” concept, and it's what a “content distribution network” or CDN really is.

What we can learn from this is that there is only one motivation to build a network—profit. There are, however, two paths to profit. One path exploits the customer, and the other exploits the relationship the customer seeks.

## **New Topology, New Internet?**

Today, the Internet is primarily a series of “access networks” that offer individuals and companies a portal onto a cascade model infrastructure. We've already noted that from the perspective of the users, this access structure can be viewed as a multi-layered model since a given user will typically have several options for Internet access. Since content distribution can be viewed similarly to web page distribution, one could argue that the content explosion would not necessarily change the model, but we think that it must.

The issue with content is that it is far more “concentrated” than web pages. While you could argue in abstract that things like YouTube create a framework where everyone is a content producer, the fact is that if you're talking about content people will pay for, the producer options are pretty much TV networks and movie studios. There are perhaps 900 major examples of these groups worldwide. In the early days of content networking, at least, it seems obvious that this small number of players (or an even smaller number of “content aggregators”, the group Google clearly plans to lead) will generate nearly all the profitable content distribution.

As the number of experience sources falls from tens of millions to potentially just “tens”, the nature of the network required to support the relationships changes. Content networks are logical stars built around the content resource, not meshes built to support any-to-any connections. As we’ve said before, these CDNs differ from a piece of the Internet in that they don’t carry any traffic other than that of the owner. To reach a user, they peer with the access players only, and there they attach not as an ISP (who owns customers, assigns their addresses and carries their traffic) but rather as a router peer, a guy who contributes his own “customers”—himself.

The effect of this is to begin to create a new structure, something that’s in parallel with the current “Internet”. All it takes to make this work is a content broker with enough opportunity to want to have a premium channel to the user. Google fits, as do others. This new structure doesn’t support n-way connection as the Internet does, only what is in effect a pure distribution or one-to-many connection. The growth of this second model is inevitable, it’s not some kind of sinister plot to undermine the Internet, it’s just a sinister plot to make money.

Two factors facilitate the growth of this new CDN model. First, there are only a small number of profitable content resources to distribute, as we’ve said, and thus it is very feasible (inevitable, in fact) to cache them locally in all the major metro areas. This means CDN scope will be limited, making it less expensive to build one. Second, the broadband revolution is reducing the number of access players with whom a CDN would have to peer. To support a thousand ISPs in peering connections takes a thousand nodes connections, but in theory one could cover about 85% of the US population with about 200 peering connections to the major RBOCs and MSOs.

As seen by the end user, all of this doesn’t necessarily change “the Internet” at all, but under the covers it brings about a major shift, and one that will ultimately make the whole net neutrality issue moot...or worse.

## The Business Impact of the Models

Content is clearly the major source of traffic growth for the Internet even without considering any “financial” issues. Even more clearly it’s the only credible source of what might be called “revenue traffic”, meaning stuff people would pay incrementally to obtain or for which QoS fees might be assessed. If this traffic flees the Internet, there is no real source of incremental funding to drive Internet expansion. The CDNs become the target networks for investment because they are driven by the applications with the highest ROI.

Google is right; the Internet can’t support video. Google certainly realizes that the layered model of CDNs will end up being the video delivery strategy of the future, and that the problem won’t be solved on “the Internet” as we know it, but on a parallel structure that will look to users like a part of the Internet, but that will instead be the thing that remakes the Internet as a business process.

Anyone with money can build a CDN, as long as there’s “C” to “D” (content to deliver, for those who can’t follow decomposition of acronyms) that can generate reasonable ROIs. We either have to believe those revenue opportunities, or we have to believe there’s no revenue future in public data services. One way or the other, the future is going to be radically different.

Google, with its plans and comments, is signaling that they realize that the search engine paradigm of revenue growth is very quickly going to dry up. They are not planning to deliver search results with their expanded network, they’re planning to deliver video content and their comments this month should demonstrate that to even the most skeptical. If you host the content you can insert ads—it’s as simple as that. Obviously, everyone else who expects to make money on ads will have to follow, and it’s not at all unlikely that movie studios and/or networks will build their own CDNs, or “rent” them.

Level 3 is clearly planning to be the guy from whom the renting is done. What the layered CDN requires is a high capacity network peered to a lot of key local access players, and that’s exactly what Level 3 has, and is expanding on. We believe that Akamai may also be looking at something along these lines, something that’s less Internet-cache intensive and more distribution via fat pipes. This illustrates how fundamental this new model of layered CDNs changes things; if Akamai doesn’t make a move it’s bypassed and sunk in a couple of years.

The Level 3 situation shows the equipment side of the story. Simple truth guys; routers are not only not the best way to build a CDN, they're not even a particularly good way. If you want a great CDN strategy you want a strategy that's all servers and Ethernet/PBT and optics. Does that sound something like what Level 3 has? It's also the rumored strategy of the search engine players who are looking at CDNs right now.

Which brings us to Avici, or more properly their Soapstone Networks business unit. Announced on February 15<sup>th</sup>, Soapstone is a separate network control plane that can be laid on top of virtually any network technology and provides support for traffic management, path creation and maintenance, and application-to-network coupling. One of Avici's key missions with Soapstone, according to their release, is support for PBT. Is this all a coincidence? Not hardly. Avici is the router vendor for AT&T, where it runs what carrier statistics show is the most reliable converged IP network on the planet. Might these guys know something about AT&T's plans, and those of other providers? It would hardly be surprising.

We did a carrier sweep in the fall, and of the ten major providers we talked with, six had active PBT plans. Every single US provider is at least considering it. Why? Because PBT is a great way to support a CDN. Are you starting to see a pattern here?

The issue goes beyond CDNs, however. Enterprise services like leased line, frame relay and ATM, and even VPNs are different from "the Internet" in that promiscuous any-to-any connectivity is a vice for them and not a virtue. Suppose you stick these services on "CDNs" too? OK, it's true that remote access to VPNs for road warriors isn't an application for CDNs (you need Internet ubiquity), but nobody pays incrementally for that anyway. Branch office service VPNs are surely CDN-like applications; of the 1.5 million satellite sites of multi-site businesses in the US, 1 million are located in only 65 major metro areas. A CDN-like network with 65 access peering points could fulfill a heck of a lot of demand.

Adding in converged enterprise services also adds in service management as an issue. The Ethernet/optical behavior of CDNs needs to be linked to service management systems and high-level applications, including IMS. That, of course, is what Avici proposes to do with Soapstone, but it's evocative of yet another truth, which is that the network of the future is squeezing the "service of connection" that the Internet and routing represent between higher-level applications and server-based features and lower-level CDN-like transport.

Routing is getting marginalized. OK, that's not a popular perspective, but we remind our readers that we have been talking about this trend toward "layering" of CDNs for a full decade, and it wasn't popular for most of that period either. Truth isn't always palatable, and in fact many times it's quite unpalatable to a lot of the players. We think that's the case here.

## About the Author

Tom Nolle is the founder and president of [CIMI Corporation](#), a telecommunications, media, and technology consulting and research firm since 1982. Tom is a software architect and developer by background, but currently focuses on strategic market, product, financial, and regulatory issues. He's also the Chief Strategist for [ExperiaSphere](#), an open source initiative that focuses on development of Java-based service logic and service management frameworks for next-generation networking. Tom is a contributor to Network World, the No Jitter CMP publication, Telecommunications Magazine, Tech Target Search Telecom, and Internet Evolution. He's also the author and publisher of [Netwatcher](#), the oldest continuously published networking publication in the industry. He's a 25-year member of the IEEE Communications Society, a member of the Telemangement Forum, and a speaker at many industry events.



## About the Webtorials® Editorial/Analyst Division

The Webtorials® Editorial/Analyst Division, a joint venture of industry veterans Steven Taylor and Jim Metzler, is devoted to performing in-depth analysis and research in focused areas such as Metro Ethernet and MPLS, as well as in areas that cross the traditional functional boundaries of IT, such as Unified Communications and Application Delivery. The Editorial/Analyst Division's focus is on providing actionable insight through custom research with a forward looking viewpoint. Through reports that examine industry dynamics from both a demand and a supply perspective, the firm educates the marketplace both on emerging trends and the role that IT products, services and processes play in responding to those trends.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler at [jim@webtorials.com](mailto:jim@webtorials.com) or Steven Taylor at [taylor@webtorials.com](mailto:taylor@webtorials.com).

### Webtorials Briefs

Vol 2, Number 6

### Published by Webtorials Editorial/Analyst Division

[www.Webtorials.com](http://www.Webtorials.com)

### Division Cofounders:

Jim Metzler

[jim@webtorials.com](mailto:jim@webtorials.com)

Steven Taylor

[taylor@webtorials.com](mailto:taylor@webtorials.com)

### Professional Opinions Disclaimer

All information presented and opinions expressed in this publication represent the current opinions of the author(s) based on professional judgment and best available information at the time of the presentation. Consequently, the information is subject to change, and no liability for advice presented is assumed. Ultimate responsibility for choice of appropriate solutions remains with the reader.

### Copyright © 2008, Webtorials

For editorial and sponsorship information, contact Jim Metzler or Steven Taylor. The Webtorials Editorial/Analyst Division is an analyst and consulting joint venture of Steven Taylor and Jim Metzler.