Carrier Ethernet and the New OSI

Unconventional Wisdom
Tom Nolle, Founder and President, CIMI Corporation

Introduction

Back in the mid-1970s, the organization then known as the CCITT (now the ITU-T) formulated a rather revolutionary thing called the "Basic Reference Model for Open Systems Interconnect", better known as the OSI model. This model created a seven-layer architecture that has been used ever since to describe network protocols and networks themselves.

Despite the longevity of the references to the OSI model, the conception of the OSI model has changed over the years. Some "layers" have been added, and some don't seem to be getting used very much. Most recently, there is talk about concepts like "virtual networks" and "abstract topologies" that don't clearly relate to the old OSI concepts. To make matters worse, the Internet's evolution, based on TCP/IP, never strictly followed the old OSI model at all. A reasonable person might ask whether people who talk about "Layer One" or "Layer Three" aren't blowing kisses at an old friend instead of recognizing the relevance of the original OSI model.

Are they? What is the real value of layered architectures today, the real role of layered protocols, and the relationship between protocols, services, networks, and users? We're proposing to look at these questions here.

Underneath the OSI Covers

The OSI model as it was originally described consisted of seven layers:

- Level 7: Application
- Level 6: Presentation
- Level 5: Session
- Level 4: Transport
- Level 3: Network
- Level 2: Data Link
- Level 1: Media

These layers or levels aren't all independent. They fall, in fact, into three distinct groups:

- Level 1 and Level 2 were envisioned as "point-to-point" or intermodal connections.
- Level 3 was envisioned as the "connectivity" layer that bound everything together.
- Levels 4-7 were envisioned as "end-to-end" or user-to-user behaviors, outside the scope of the network itself.

From this, you can kind of taste the flavor of OSI. The model envisioned a **connection network** that provided information exchange between addressable endpoints, each of which might serve a variety of applications and users. This connection network was created by a mesh of nodes that were interconnected by data links.

Why seven layers? The reason was that the OSI people were operating under a single guiding principle; that of **independence of layer behavior**. Every layer "saw" the network underneath as a single structure, defined by its boundary interface. Functions were segregated by layering so that there was no need for a given layer to know **how** the lower layer had done something, only the way it provided services.

The OSI model was the Word in the 1970s, but even in the late 1970s there were gaps in the process. The Internet, which at the time was a government activity, didn't have the seven layers, and many at the time noted this. In truth, it was a kind of media issue. TCP/IP defines Level 4 and Level 3, respectively, of the model. It's true that the Internet didn't standardize a Level 2 or 1 or 5 thru 7 (though later on "sessions" were added), but there was nothing in TCP/IP that invalidated the model. The real gaps were created by Bob Metcalfe, with Ethernet.

LANs in general, in fact, put the first stress on the OSI model because they offered what was traditionally a Level 3 service (connectivity) below Level 3. True, LAN protocols ended up creating "sublayers" of Level 2 and thus made the old notion of seven protocol layers kind of arithmetically oxymoronic, but that wasn't as much a problem as the fact that some of the features of Level 3 were being subsumed into the next layer down.

The IEEE 802.x protocols essentially create a connection model, the role of Level 3 in the OSI model. The contradiction was viewed at the time as perhaps an indication that OSI had become jaded (which may be true, but is still simplistic), and we moved on. One reason that this rather key evolutionary guidepost was largely missed was that carrier Ethernet wasn't a factor. Nobody really cared much about the OSI model on premises anyway, and the CCITT/ITU was a carrier body.

The higher-level part of OSI, the end-to-end stuff, was also evolving. The Internet created a world where service experience was presumed to be created "over the network" rather than in it, which fits rather nicely into OSI at one level. The fact that OSI was a carrier-driven process makes people wonder about why such a disintermediating concept as "over the net" could have developed there, and the answer is equally instructive; its connectionlessness.

The OSI model, in the early days, was based on connection-oriented behavior. The first OSI-compliant network was the packet switching X.25 standard, which is connection-oriented at Level 3. Frame relay and ATM are also connection-oriented. This created stateful network behavior, of course, but that was a logical extension of the PSTN concepts that still ruled the world in terms of revenue, profit, and investment.

There is an IEEE 802 DLC process that is connection-oriented, but the Internet and IP are connectionless. As the Internet took hold as the clear winner in data-to-the-consumer, connectionless behavior dominated at Level 3 and the notion of "connections" turned into the notion of "sessions" at Level 4 (TCP) or at Level 5 (SIP). This is how OSI, a carrier-driven process, ended up supporting a carrier-exploiting service model. But more was to come.

Carrier Ethernet and the New OSI

As we noted earlier, carrier Ethernet wasn't around when the OSI model was invented, but it's sure around today. In fact, it is fair to say that **carrier Ethernet is the major driver of change in networking today, both philosophically and in terms of investment**. The scope of its impacts is far too broad to cover in a single piece, so we're going to focus on how it's altering the conception of the "OSI model".

Think of the OSI model not as a "standard" that has to be adhered to but as a model that describes the behavioral goals of application communication. Since service provider networks support application communications that goes off-site, the behavioral goals set by an evolution to the OSI model then set the infrastructure requirements to which provider must build.

There have been three major forces unleashed by carrier Ethernet:

1. Since it offers both point-to-point and multipoint network capability at Level 2, carrier Ethernet presents at least a partial service alternative to Level 3. In effect, carrier Ethernet makes it possible to make Level 3 protocols (like IP) end-to-end like those of the other higher layers already are. Would a user who obtained IP packets over an Ethernet connection know whether there was a router in place? Not from a service perspective.
2. The addition of PBT to the carrier Ethernet repertoire adds stateful behavior back to the network at Level 2, which means that statefully-produced services can be offered by providers. One can argue that MPLS and TMPLS both do this, but that only magnifies the trend.

3. The subsumption/subduction of network behavior down to the data link layer of OSI encourages speculation that maybe further subduction could be undertaken. The concept of "path agility" across network layers ranging from the optical (traditional Level 1) to Ethernet (traditionally Level 2) to IP/MPLS (Level 3) and pseudowires (Level 4?) starts to really blur the concept of the lower layers of the model.

This last point may be the most critical of all. The conception of connections as tunnels that can be created at any layer that's convenient and migrated from one to another under conditions of traffic growth, cost reduction, or pretty much anything else, largely eradicates distinctions between layers of protocol and in fact largely disconnects network protocols from service protocols. The growth of the Internet may have spawned the success of IP as the framework for service projection, but it may just as well have spawned its death as a universal network technology.

So what? Isn't all this layer stuff just pap for the masses anyway? Not really, because the more layers you smear services over the more equipment vendors can link their features directly to the money flow. If all services are at Level 3 then anyone below that is just plumbing. If service concepts are layer-agile, then any layer into which they can migrate creates a new set of players - a new class of competitors.

What is happening in effect is a mash-up of OSI layers, the creation of a new conception of the OSI model, and one that will have a major impact on spending and competition over the next decade. The new conception recognizes three layers that roughly correspond to the layer groups we talked about, but that's where the similarity ends.

The bottom of the "new OSI" is the Connection Layer. This layer is responsible for information delivery to the user access point, the DMARK of old, and it supports the three basic service models that the MEF defined and that exist today; E-LINE, E-LAN, and E-TREE. There is no specific protocol for this layer, no specific infrastructure. Its only requirement is that it delivers the stuff that's introduced to the other endpoints in conformance to the service model and SLA.

The middle layer is the Facilitation Layer. This layer provides services to users/applications to enable their use of the Connection Layer and to provide other facilitation for even higher-layer services. It can be viewed as the place where relationships are managed. We'd include in this stuff like directory services, address translation (DNS and ENUM), identity and trust, etc.

The top layer is the Application Layer, and in this layer all of what we would call "user" and "application" services is structured. This layer recognizes that most "experiences" sold will be framed over-the-net, but might draw on in-net services from the Facilitation Layer. VoIP and IPTV clearly fit here, as does the web and all the stuff we recognize. And service devices tend to be tunnel on/off-ramps in the new world, not network devices that switch or route.

Consumer services aren't open connections or meshes, they're stars built around experience hosts. Any network technology that can create pipes from a server farm to a central office can be used to deliver IP video. There doesn't have to be a router in the picture at all. Interestingly, this architecture is a validation of the old BRAS architecture that's still represented by Juniper's E-series and that was recently acquired by Ericsson in its deal to get Redback.

It's clear that consumers would buy stuff at the top of the model, but we think that enterprises would, as they do today, consume directly at all layers. Thus, enterprises are sold connection service and facilitation services, and could in theory mix and match among providers to make all of this work for them. Enterprise services are more traffic-symmetric and so it may be that enterprise services will be a key element in deciding just how much IP-specific network technology (as opposed to service technology) would deploy.

Clearly this shift will impact equipment vendors, and the impact there could be very profound indeed. We think each of the following trends will evolve out of this new model:

- It will become increasingly difficult for someone to sell hardware at a single one of the old OSI layers. This will impact the providers of the "rich" layer, Level 3, first and so the router vendors will be faced with taking a broader view of networking than they'd like (since margins are higher at Level 3). Some of the

Level 2 (Ethernet switch) vendors, particularly the bigger ones like Extreme and Foundry, will be forced to create some sort of Level 3 accommodation.

- The network operators will focus on the Facilitation Layer as a means of boosting their profits (this is already visible in the BT API offerings) and this will force equipment vendors to add facilitating offerings to their product line or be disconnected from provider revenue focus. This may well be why Cisco is moving into the software area already.
- Application-layer players will see facilitation as the natural battleground with the network operators and will move to offer things there very quickly, which is why we believe Microsoft's CSF is being positioned the way it is and why Google is on an acquisition roll with things like security.

The big question here is what becomes of the principle of independence of layers. This, as we've noted, was one of the architectural guideposts for the original OSI model. In the new three-layer structure we're suggesting, the concept of an end-to-end service requires a multi-layer partnership doesn't it?

Maybe it sounds counter-intuitive, but the answer is "No." In fact, we think this new world is even more OSI-ish in terms of philosophy than the original model. The reason is that we are now accepting that service behavior is composited explicitly from a combination of transmission/connection and hosted features. OSI relegated hosted stuff to what was essentially outside the network, where the emerging model embraces it as part of the service.

Do We Care?

This is one of those discussions that seems painfully abstract. Do we care how many layers there are in a network model, or even if we have one? We do, for the same reasons we did in the mid-1970s. Communications depends on consensus. We can't network people or applications absent some procedural rules on how the process will work. Already in services like IP voice, we see individual providers and vendors building islands that are able to interconnect with each other only through the old PSTN. If VoIP requires the existence of the technology it's supposed to displace in order to provide universal service, we've left the path of wisdom here.

We also care if we're a network operator, an enterprise, or even an equipment vendor, because this is all shaping the techno-world in which we're making decisions. The number of service-formulating options is increasing in the new model. Even though the old OSI had a goal of layer independence, the fact is that all of the layers exerted a local gravity on their partners that tended to drag them into alignment around a pretty limited service concept. With the new concept, the marriage of all of the transport/connection layers into a single model opens up a new world of technology independence. Generally that's good for those who consume technology and less so for those that produce it.

In this case, the latter problem just might not be true. Over the last two decades we've been struggling to conceptualize a new architecture for a new public network. CIMI Corporation has been in the networking business for 25 years, and in that time we've seen ISDN, ATM, and IP convergence all rise—and fall. It may be that the problem we've had isn't the lack of value for these technologies per se, but the lack of a complete conception of future services and a link between those services and a fulfillment architecture.

Just saying "seven-into-three" and waving a magic wand doesn't create either one, of course, which is why the current situation may be a vendor opportunity. We may have spent the last 25 years taking baby steps in the direction of a new network model, and we may be on the cusp of being able to define it as clearly as the OSI model of old was defined. If this new model is more responsive to market conditions, which having originated out of them it's likely to be, then it could be the jumping-off point for a new set of network products and a new set of network successes.

Even this good news can be bad news, though. Equipment vendors are notorious for taking root and becoming trees. IBM might be the only example of a vendor who managed to reinvent itself despite the fact that the reinvention seemed to be threatening to its incumbency in every step taken along the path. But then IBM is the old-line vendor, the one who existed for decades. Network vendors are relative infants. Maybe there's still time to learn how to walk, and what direction to walk in.

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 Telemanagement 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 or Steven Taylor at taylor@webtorials.com.

Webtorials Briefs
Vol 2, Number 10

Published by Webtorials Editorial/Analyst Division
www.Webtorials.com

Division Cofounders:
Jim Metzler
jim@webtorials.com
Steven Taylor
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.