



Buying Metro Ethernet Services

Peter J. Welcher

Introduction

I'm writing this on 1/2/2004, so Happy New Year! Rather belated, by when you will see this in print. But it sounds better than Happy March!

The previous three articles were about Security topics, including some fresh ideas on what's changed recently in the security world. I do hope the links were useful to you. I have lots more security thoughts, but it seems like it's time for a change of topic for a while.

So this month I'm going to write about Metro Ethernet Services. The focus of this article will be on what you're buying, with a glimpse of the technology under the hood. I hope this demonstrates that I've actually learned something: when I wrote about MPLS, I first wrote about how it works, which is what I think is neat. But frankly, you don't need to know all that to be a consumer of MPLS IP connectivity or WAN services. What you do need to know is how to design for that situation. The result was a later article, *MPLS VPN's From the Customer Side*, which can be found at the URL <http://www.netcraftsmen.net/welcher/papers/mplspece.html>.

The same applies to optical networks. What Cisco is putting out is heavily Service Provider flavored. That's fine if you're a Service Provider and want to know about how the optical side of things is going to work. There are some very important engineering and provisioning details that need to get done right. Yet the data (bridge, switch, route) side of things is arguably more complex from the protocol side of things. There are certainly some important details that need to be paid attention to there. From my point of view, what the devices do (and don't) do on the data side is what's interesting, and there's some surprising vagueness there. More on that later. And by the way, no, I haven't seen any vendors do a better job of documenting what they do than Cisco. Sometimes far worse!

I started envisioning this article when I attended two good Networkers 2003 presentations. I found myself thinking that what I was hearing was some of what I'd been wanting to know for a while, but also thinking that folks need to know how to think through what they're getting, because most of the vendors seem to be staying rather vague about it. My concern would be that a newcomer to the area might come with certain expectations or assumptions that might not in fact be valid. You'd hate to go and buy some service and then discover it doesn't do what you thought it would do for you. So my mission in this article is to set you up to be an informed consumer of Metro Ethernet services.

Why Metro Ethernet?

Fast cheap bandwidth. Tunable, in the sense that your provider may be able to add more bandwidth for you quickly.

Also, you can buy this sort of thing now. It is real, if geographically spotty. We have at least one enterprise consulting customer who is buying it. The service is attractive, it's affordable. Your mileage may vary, since providers may or may not be doing it in your geographic area. And even if they are, the prices may be shocking (high or low).

We're also working with two governmental entities and a large enterprise, on using fiber plant they've bought or leased to implement Metro Ethernet services for government departments. If you're like these two government entities, with your own fiber, you're the Metro Ethernet Service Provider (MESP? -- my acronym). So you do need to deal with all the SP details. Still, this article may be relevant to thinking through what you're trying to do with all that fiber, and designing the data side of things.

What Metro Ethernet does is gets you a connection, generally within a metropolitan area (hence "Metro"), generally running

over optical links and hardware. The connection physically terminates in an Ethernet cable. We're talking Fast or Gig Ethernet connections here, at prices in the range of \$1000-2000/month. You may be rate-limited, e.g. to 200 Mbps. Did that get your attention?

Some ideas as to how you'd use it:

- 1 High speed link from Point A to Point B (usually the first one goes to backup data center or between data centers)
- 1 High speed link to Internet
- 1 High speed link to Internet Interconnect or Colocation Facility
- 1 Distributed campus (which can also be done without the optical gear under some conditions)

Let me explain the third option (with thanks to the guys who asked me what I thought of it -- you know who you are). I ran into this in Long Island. Broadwing was thinking creatively, and proposed to a customer to provide T1 services to remote sites, terminating in Channelized T3's on routers. This is arguably how you make modern leased lines competitive with Frame Relay in terms of port costs. They also take distance charges out of the mix, which also helps. It's all done with optical transport. The extra twist to this is that the termination would be on a managed Cisco router in Manhattan, with GigEthernet Metro Ethernet to deliver it to the customer premises out on Long Island. The cost of the access would be less than Verizon pricing on several T3's. Broadwing would manage this end to end. That's pretty attractive!

This customer needs to scale up some QWest Frame Relay, and likes diversity, so they then started talking to OpenAccess and QWest about doing the same for their existing FR services. And then asked me for comments. My comments boil down to "neat", with some caveats that I go into briefly below.

The consulting customer then realized that having a presence at an major Internet interconnect might make other high-speed services affordable, which also seems to be true to some extent. Some good thinking there too!

Kudos to LightPath and OpenAccess for being competitive and customer-oriented in terms of price and services, and Broadwing for thinking creatively.

I know of another good sample use of Metro Ethernet. Use EPL services (below) to connect a couple of large Board of Education HQ L3 switch/routers via Metro Ethernet to smaller switches in city or county school buildings. That keeps the gear and design complexity in each school simpler (central switch doing trunking to per-classroom switches, with each room in its own VLAN?).

What you might not expect is that the link may not behave exactly like a plain Ethernet cable. One simple example that goes with the first example above: one of the MESP's has Cisco boxes in between Customer Points A and B, with CDP turned off. The two Cisco devices at the end of the GigE "virtual link" therefore do NOT see each other via CDP. They also don't see the MESP boxes. So it isn't Plain Old Ethernet. And that's really the point of this article.

Types of Service

By now I hope you can see where this might or might not be relevant to your organization.

Let's talk about what you might be buying. There are several types of service, and different vendors do use the terms differently. I'm going to try to follow the Cisco terminology, and I'll describe each variant so you can check if the other vendors are using the words in the same way.

First we need some terms describing what the service does. Here are some of the characteristics of the service you might buy:

- 1 Types of connection: point-to-point (P2P) or multipoint (MP).
- 1 Customer equipment (CE): router or bridge (switch).
- 1 Architecture: VPWS, VPLS, EoOT
- 1 Service Characteristics: Service Multiplexing, VLAN Transparency, Bundling, L2 PDU Transparency

Explanations:

VPWS = Virtual Private Wire Services

VPLS = Virtual Private LAN Service

EoOT (my acronym) = Ethernet over Optical Transport, which might be SONET/SDH, or EoxWDM, Ethernet over Coarse or Dense Wave Division Multiplexing (CWDM/DWDM).

Service multiplexing means you can use one link to reach multiple sites, using VLAN's to control target site.

VLAN transparency means your VLAN's get trunked through the link.

L2 PDU transparency refers to BPDU's and CDP packets: transport of L2 management frames. If you don't have L2 transparency, your switches will not be able to detect Spanning Tree (ST) loops!

Some services defined by Cisco:

Service	Type	Architecture	CE	Characteristics
Ethernet Private Line	P2P	EoOT	Router or Switch	VLAN Transparency, Bundling, L2 Transparency
Ethernet Relay Service	P2P	VPWS	Router	Service Multiplexing
Ethernet Wire Service	P2P	VPWS	Router or Switch	VLAN Transparency, Bundling, L2 Transparency
Ethernet Multipoint Service	MP	VPLS	Router or Switch	VLAN Transparency, Bundling, L2 Transparency
Ethernet Relay Multipoint Service	MP	VPLS	Router	Service Multiplexing
MPLS L3 VPN	MP	L3 MPLS VPN	Router or Switch	Routed Ethernet connection, with MESP providing routing services for you (see the prior article , mentioned above)

Ethernet Private Line (EPL) is a point-to-point service, interconnecting two Ethernet ports. There is no multiplexing of services, such as say Frame Relay or ATM do. The EPL acts like a long point-to-point Ethernet cable, transparent to whatever your gear sends down the wire. This would be built using EoOT optical technology. It provides Bundling in the sense that you can make the link a trunk, and your VLAN tags will be transported, possibly using QinQ (see below).

The OpenAccess service mentioned earlier is a variation of this, intended for interconnect of two routers. It provides P2P Ethernet but no transparency.

Ethernet Relay Service (ERS) is VLAN-based. I think of this as being like Frame Relay, with VLAN acting in place of the DLCI. That is, your site may trunk to the MESP. Each VLAN then gets transported to a different remote site. BPDU's and CDP do not pass through transparently. This could be built using L2 MPLS VPN pseudo-wires, with one pseudo-wire per VLAN. The first cut at Cisco AToM EoMPLS services could provide this sort of connectivity.

Ethernet Wire Service (EWS) is a point-to-point service, transparent to BPDU's. It is probably based on L2 MPLS VPN services by the MESP. The distinction is that the entire trunk is tied to one pseudo-wire to another location, so the service provides a P2P trunk between sites, in effect. If your Cisco IOS code doesn't support full EWS yet, it can be built using QinQ (see below) plus the first release of EoMPLS pseudo-wire (one VLAN at one site connected to the 2nd site).

Ethernet Multipoint Service (EMS) makes all CE devices peers. There is no service multiplexing, in that all VLAN's appear at all sites. BPDU's pass through transparently. **Transparent LAN Service (TLS)** is another name for this. It might be built using QinQ (see below) on top of a multipoint VLAN service. Or it might be provided more directly by the transport device.

Ethernet Relay Multipoint Service (ERMS) allows P2P and MP to co-exist on one Ethernet link ("UNI"). Some VLAN's might be P2P and others might be MP. This hybrid service is (probably) opaque to customer BPDU's.

For the situations above which do **not** provide L2 transparency, the CE device should probably be a router for this. Otherwise you and the MESP run some solid risk of being down due to some really large scale Spanning Tree issues. I've had to deal with some lately, and diagnosing large ST loops is not easy nor fun. The real issue here is, due to the lack of transparency, ST loops are highly likely to happen as soon as some customer creates some "backdoor" L2 link between

sites.

The same might also apply to transparent links, if you're cautious. If the MESP allows switches on one or both ends, my next question as a customer would be: "how are you going to keep your other customer's ST loops from killing my link?" I probably want to hear words like "rate limiting" in their answer. My point here is that a customer broadcast storm at a site will propagate through a switch and into the MESP network.

About the Cisco Hardware

Some brief words about what the Cisco 15454 series does. This box provides aggregation of various media types onto optical (typically SONET) links. A traditional TDM vendor can use this box to aggregate T1's and T3's. This sort of thing is probably what Broadwing is doing with its national leased line service -- I'm not claiming they're doing it with a Cisco box.

The 15454 also provides Ethernet transport. This is done with various cards. There are three series of Ethernet cards for the 15454. From oldest to newest: the E-series, the G-series, and the ML-series.

The E-series supports Point-to-Point (P2P) or Multi-Point (MP) configuration, as well as trunk (carrier) side STP. It does L2 switching, especially within a card. Using the MP configuration costs some bandwidth. The E-series matches up with certain customer VLAN's and transports them. And the VLAN's have to be coordinated across customers.

The G series is simpler: it only does P2P, and it just passes bits through. If you want MP, you can do something like provision a ring or mesh using 2 ports per 15454, and then hang L2 switches off the two ports. In other words, to anything fancy, use a switch or a router alongside the G series blade. Practical designs right now use the 3550 switch for inexpensive designs, or the 7300 or 6500/7600 series switches/routers for more ports, and the 12000 GSR router for big designs.

The ML series is IOS-based, and does Integrated Routing and Bridging (IRB), and can do various routing protocols. It is supposed to do VRF-Lite for per-customer routing tables and light MPLS. VRF-Lite requires connection to a full MPLS router as Provider Edge (PE) device. The idea here is to do VLAN's per customer in a building, trunk the building switch to ML-series port either in the building or in a PoP within single mode GBIC range (10's of miles), and have the ML series act as router for all those VLAN's. I note the 4.0 documentation mentions VRF Lite, the 4.1/4.5 documentation does not, always a bad sign. I gather folks are also using external 3550's for the VRF Lite functionality.

My guess on the evolution here is that Cisco was finding development costs high and market tight. The E series represented dedicated and complex control code and perhaps chipsets. The G series is the core functionality that MESP's need and want most. And the ML series gets over to the IOS code base, getting ready for MPLS to provide the unifying L2 and L3 optical transport control going forward. And leveraging all the R&D that's gone into the Cisco IOS code.

About VPLS

I've discovered that I much more firmly believe in VPWS, the point-to-point wire emulation services, than I do in the VPLS, LAN services.

Cisco is trying to be agnostic, especially since some other providers (e.g. Juniper) are somewhat pushing VPLS, and may have sold some folks on it. The concerns that I have about VPLS are shared by several very knowledgeable people I've spoken with (who shall remain anonymous). Hmm, come to think of it, I haven't had the chance to debate or discuss this with my friends who now work for Juniper. You've been warned! Anyway, let me briefly run through the concerns. The excuse for this rant is to fill you in on what the MESP might be trying to sell you, and why it might be hazardous to your network's availability.

The quick description of VPLS is that the MESP network looks like a distributed Ethernet switch, built over MPLS technology. To imitate standard L2 switching, you either have to have flooding of customer MAC layer traffic to all points connected to the VPLS, or the MESP switches need a protocol to learn which edge device has a specific MAC address attached to them. The latter is basically what ATM LANE did, and the standards committee wisely chose not to go down that path again. Very few people may have ever understood LANE, troubleshooting it was a nightmare, etc. Lessons learned!

So VPLS Provider Edge (PE) devices flood unknown unicasts, but associate each SP peer device with a virtual port. That way, over time they do learn to associate MAC addresses with peers. A very large number of MAC addresses! This does trade MESP bandwidth for some degree of simplicity.

Think about number of MAC addresses for a moment. These devices need to learn many (but not all) MAC addresses for

each customer that attaches. And you need some form of protection in case one customer somehow acquires or configures devices with MAC addresses duplicating those of another. Per-customer (per-VPLS) MAC tables is one obvious way of doing that, and is mentioned in the IETF draft as a possible vendor implementation choice.

As far as BPDU's, they are tunneled over IP in MPLS VPLS. This is highly desirable for preventing self-inflicted ST loops. That protects both customer and MESP. But there are some scenarios with back door links where lack of MESP PE MAC address flushing might create issues, like the PE devices perhaps forwarding frames to the wrong peer for 5 minutes. So if I were buying VPLS services, I'd be listening for the MESP to discuss backdoor links and why I should not be doing them, at least not at Layer 2.

With VPLS, you still have flooding and bandwidth issues with multicast and broadcast. They go everywhere. There are some other management-related issues. E.g. QoS and availability SLA enforcement.

One possible design to work around this is to only allow routers to attach to the VPLS. That would certainly help cut down on MAC addresses. The drawbacks there are peering adjacency numbers and routing protocol behavior, multicast behavior ("spraying"), and traffic aggregation/QoS controls.

All of the above discussion really applies to VPLS built over a MPLS core. In short, this is not-very-proven technology that attempts to go where problems have been observed in the past, at just enterprise scale. As it evolves, its limitations may be better understood, and workarounds may be found for the known issues.

It appears conceivable that VPLS could also be offered based on some combination of optical devices (such as the Cisco 15454 and E-series blades), QinQ trunking (see below), and L2 switches as Provider Edge (PE) devices. This scares me, as it would be based on pure L2 technology, which just does not scale very well. Problem isolation, modularity, and troubleshooting all seem to have potential problems. When I design networks, I try to avoid building campus-wide VLAN's. So exactly why would doing multi-campus-sized VLAN's remotely resemble a Good Idea? (I've had this question for 3 years now, and have yet to see a good answer to it.)

In summary, if a vendor is trying to sell you VPLS, stop and ask some questions. If they're talking full mesh between many devices, ask if those are L2 or L3 devices. If L2 switches, ask why massive scale STP is remotely desirable, and how they're going to keep your network from melting down. If L3 routers, ask why full meshing is desirable, and why you want 30 or 100 or $n \times 100$ routers forming adjacencies. (Note that in MPLS L3 VPN, routers form adjacencies with the SP PE router, not with each other. So traffic follows a full mesh of paths but your routers do not have anything resembling a full mesh of adjacencies.)

Some providers are pointing out that full mesh provides lower latency. At optical speeds, this is a non-issue: shaving 10 msec off 40-60 msec just doesn't buy you very much. If it does, you've got applications or protocols with throughput problems, and I'm not sure the network is the place to fix that. Admittedly, sometimes we don't have much choice, but trading stability for throughput may not be a good trade-off. I personally very much like the idea of routers and some hierarchy. If you're still worried about latency, you can still cut way down on latency with hierarchy: create 4 pairs of regional routers with meshing, and have everything else feed into that core.

Now that I've talked you out of VPLS, what about ERMS? Well, ERMS involves mixing VPWS and VPLS, perhaps several virtual VPLS's, with various VLAN's getting either P2P or VPLS treatment. I.e. VPLS + some provisioning complexity. Enough said?

Underlying Technologies

Optical: Optical switches have Ethernet ports. They usually allow point-to-point transport of Ethernet over SONET or CWDM or DWDM optical media. They may allow multi-point connections. The Cisco equipment does so at some cost to the MESP in terms of bandwidth they can offer. VPLS like services could be provided via a set of L2 switches outboard of a ring of point-to-point optical connections. Scaling that would be problematic. No Service Provider wants to manage hoardes of little boxes. So I look to optical more for the P2P type services, perhaps with some QinQ transport.

VLAN-based: Smaller MESP's may offer VLAN-based services, e.g. based on Cisco 15454's with the E-series card. They may provision P2P optical links between POP's, and give each customer a dedicated VLAN for each P2P link. Your Ethernet gets optically carried to the POP, plugged into a L2 switch port assigned to a VLAN, and then that VLAN gets bridged to another POP, trunked to attached L2 switch, and the port in that VLAN connected to your equipment. That gets such a MESP up to a max of 4000 VLAN's (and probably less), which is OK as a starter business model. I've certainly seen some high operational costs by MESP's overbuilding for huge numbers of customers, which is reminiscent of the dotcom era. Building for 1000's and having 10's or 100's of customers can be a Bad Thing financially.

QinQ: This is the short name for tunneling 802.1q VLAN's inside 802.1q VLAN's. The idea is to apply two 801.1q VLAN header tags to a frame. The outer one identifies the service customer, your company. The inner one is your VLAN number. The way Cisco implements this, you tie a trunk port to a specific QinQ port with an assigned VLAN. All traffic in the trunk port gets sent out the QinQ port with the specified extra VLAN tag "wrapper". Inbound traffic gets the extra VLAN tag removed, and is then forwarded out the specified trunk port. The Cisco implementation is strictly port-to-port "virtual connection", i.e. you cannot bridge the traffic back out some other port(s) on the same switch, and you cannot tie multiple trunk ports to the QinQ port, or multiple QinQ ports to a trunk. Well, you can sort of do that by getting tricky, by cabling the trunk port back into a normal trunk port on the same switch. That trunk might connect to yet another QinQ port, and so on.

Be aware that Cisco device support for QinQ is sparse. You can do it in 3550 and 6500 model switches, and in certain routers.

Example use: you have multiple municipal departments in one building. Each has a VLAN assigned. You trunk these up, hand off QinQ to your MESP, they transport it over EPL to your data center or HQ building. This keeps each department isolated, without having to put a router at the remote site and run MPLS VPN routing for separate per-department routing tables. (I personally would prefer to troubleshoot the latter, but that's to some extent my personal preferences).

AToM: Cisco's term for L2 over MPLS transport in an IP-centric network. AToM stands for Anything over MPLS. Ethernet over MPLS (EoMPLS) is the portion of AToM most likely to be used for various Metro Ethernet services. See for example EWS and the discussion about VPLS above.

Considerations

Cost for bandwidth is probably the big thing. If it doesn't save you money or get you gobs of bandwidth cheaply, who needs it?

One cost factor I've noticed is whether the MESP wants to put optical gear at the building. If so, there are two issues you'll encounter. If they haven't already put gear in your multi-tenant building, they may want several tenants signed up, or it may take time to get the gear and get it installed. And if you're single-tenant, good luck, your contract is going to end up buying them the optical gear over time, several times over, and it isn't cheap, although SONET access gear is rapidly decreasing in cost. I've still been wondering why MESP's don't want to use GBIC's and backhaul to a metro region POP (30 km is easily doable) before hitting the SONET or DWDM gear. The puzzle to me is that various parties still seem to be acting as if fiber is scarce, even though in many regions there is a vast glut of fiber. I've heard of dark fiber outside DC for \$1500/month on a 20 year lease. 20 year anything is of course a problem for some enterprises. Then put the costlier gear in where and when fiber gets scarce. Cost of engineering and leased fiber paths is also non-trivial, but a good Metro provider ought to know what's available to where, within their coverage area.

Service Level Agreements (SLA's) are a consideration. You'll want financial teeth in case of non-availability, and you'll want some assurance that you're not just up but getting usable bandwidth. No oversubscription, no getting clobbered by somebody else's ST loop and ensuing broadcast storm, etc. Since this is new technology, I'd still be thinking about getting two MESP's, for redundancy. Did I mention, checking out diversity of paths?

Type of service you buy. As you can see above, I'm in favor of EPN and EWS. ERS also doesn't bother me, although you're trusting the MESP to pair up the VLAN's right. That's about like trusting a FR Service Provider to get the DLCI's right. It doesn't take a huge leap of faith to believe they can mostly get that part right.

Design. As you can see, we're all still learning. I think one rule of thumb may be that having a router on at least one end of each P2P Metro Ethernet link is a good thing, for containing any L2 and STP problems (are they called STPP?). Central L3 switch/router connected to multiple remote sites with only L2 switches still seems ok. I don't think you want to be troubleshooting multi-site VLAN's, that gets nasty fast. The router in the middle means you check if the EPL or EWS is up and working bidirectionally, and if so, the problem's got to be at the other end.

Colocation? By this I mean you want to think about whether colocated or managed equipment is involved, and how you feel about that. How much control do you want or need to retain? What's your comfort level with somebody else managing it for you. If the MESP is going to give you free colocation of gear, see also the next paragraph.

Colocation business issues. You really need to think about business/legal issues if colocated or managed equipment is involved. Who owns it? Who manages it? Do you get SNMP RW, or SNMP RO, or no management access to the device? If not, you can't help the MESP resolve problems, you also can't see things like link loss data that show commands might shed light on. Who provides the configurations? If you do, what's the charge for every change you make? If you have some management access (SSH? IPsec?) then do you also get to hang a modem and analog line on the AUX port? Details and costs of doing so? Can you get physical access? Who's responsible for physical security of the device? What security

controls apply to entering the area of the device? You want to be sure your connections or cabling aren't going to be messed up when somebody else is in the cage and catches on a cable or something. Suppose you terminate your service from the MESP. Does your gear get tossed out? What's the fee for keeping the gear there? Do managed device fees change? Etc.

Security exposures. You ought to understand business risk issues. What controls are there to make sure you don't accidentally get connected to somebody else, either initially or as the MESP adds customers? What social controls are there, so that a hacker or competitor can't just call up and get a connection into your network? What design or other controls are there to make sure somebody else's problems don't become your problems? What security measures has the MESP put in place as far as L2 anti-spoofing protections and protections against other L2 exploits? Routing and MPLS protections (if their network is MPLS based)?

Links

An acknowledgement is due at this point. I found the Cisco Networkers 2003 talks on the subject very interesting, as were some after-session discussions with the authors/presenters. The presentations continue to be useful and were a primary reference source in writing this article. All opinions and errors of course are mine. Here are some URL's relating to these Cisco presentations:

Title	URL
Networkers 2003 Metro Presentations	http://www.cisco.com/networkers/nw03/post/presos/metro.html
Emerging Standards in Metro Ethernet	http://www.cisco.com/networkers/nw03/presos/docs/OPT-2043.pdf
Deploying Metro Ethernet: Architecture and Services	http://www.cisco.com/networkers/nw03/presos/docs/OPT-2045.pdf

Note that the Presentations page doesn't mention OPT-2043 for some reason, but parts of it are extremely relevant.

Providers mentioned in this article:

Provider	URL
Broadwing	http://www.broadwing.com/
LightPath	http://www.lightpath.net/
OpenAccess	http://www.openaccessinc.com/

IETF L2VPN and PWE Drafts

Title	URL
Layer 2 VPN (l2vpn) Working Group	http://www.ietf.org/html.charters/l2vpn-charter.html
Virtual Private LAN Service	http://www.ietf.org/internet-drafts/draft-ietf-l2vpn-vpls-bgp-00.txt
Virtual Private LAN Services over MPLS	http://www.ietf.org/internet-drafts/draft-ietf-l2vpn-vpls-ldp-01.txt
Pseudo Wire Emulation Edge to Edge (pwe3) Working Group	http://www.ietf.org/html.charters/pwe3-charter.html

Metro Ethernet Forum

I haven't yet found a web site that tells me who is offering Metro Ethernet services, say by zip code. The following site may help some: <http://www.metroethernetforum.org/>. It lists equipment vendors, and it also lists some MESP's who are

members. SBC, BellSouth, BT, Verizon, QWest are on the list. I know AT&T is offering the service (but not a lot of info, the last time I looked). And a Cisco Press release from 3/2003 says Time-Warner is a MESP.

Summary

Dennis Hartmann is going to be working with Chesapeake Netcraftsmen on one of the large governmental Metro Ethernet and L3 MPLS VPN networks I mentioned above. I actually had him in an MPLS class, and was very impressed with him at that time. Dennis is a detail-oriented guy who knows MPLS and Optical extremely well. Dennis is co-author of the Cisco Press book *Building Cisco Metro Optical Networks (METRO)*. See also the URL <http://www.amazon.com/exec/obidos/tg/detail/-/1587050706/>. I know Dennis put in a lot of work on the book, filling in or refining details. I've been reading through the book and enjoying the clarity. It provides the optical, SONET, etc. details that I didn't have room for above. All in all, highly recommended. If you're a potential Metro Ethernet Service Provider, or responsible for the fiber and other engineering of a Metro optical network based on Cisco gear, buy this book now!

I wish you good luck with your Metro Ethernet and optical networking. Please drop me an email if you disagree (or agree) with anything I've said above, or if you can help me fill in some of the many gaps as to who's offering Metro Ethernet services in various geographic areas. If you would like network design or design review, or help with optical networks or MPLS, please drop me an email.

Dr. Peter J. Welcher (CCIE #1773, CCSI #94014, CCIP) is a Senior Consultant with Chesapeake NetCraftsmen. NetCraftsmen is a high-end consulting firm and Cisco Premier Partner dedicated to quality consulting and knowledge transfer. NetCraftsmen has eight CCIE's, with expertise including large network high-availability routing/switching and design, VoIP, QoS, MPLS, IPSec VPN, wireless LAN and bridging, network management, security, IP multicast, and other areas. See <http://www.netcraftsmen.net> for more information about NetCraftsmen. Pete's links start at <http://www.netcraftsmen.net/welcher>. New articles will be posted under the Articles link. Questions, suggestions for articles, etc. can be sent to [pjw <at> netcraftsmen <dot> net](mailto:pjw@netcraftsmen.net).

1/4/2004

Copyright (C) 2004 Peter J. Welcher