

NETEVENTS

APAC PRESS & ANALYST SUMMIT

First Draft

Guest Speaker Presentation Ethernet - bringing the Cloud down to earth

Robert Kuse

MEF Board Member & Director

Good morning, everybody. I'm here to provide some perspective and some commentary and update on Cloud services from an MEF perspective and provide some of my own characterizations and insights as a service provider.

When we talked about the subject of Cloud services within the MEF over the last year we really asked ourselves a question, which is what is a Cloud service. And we brought in within the community a whole series of consultants; we did birds of a feather; we did a whole lot of series of meetings to get ourselves educated on what it was and potentially what work streams or what work would need to be done by our group to have input and add value to this discussion.

And when we look at this historically, even as a service provider when I really look at what is a cloud, it doesn't seem tremendously different to me that something that we've done historically where we -- these are VPNs. Enterprise datacentres serving customers either through the Internet or through some Layer 3 VPN or through some other network access connectivity is a cloud. And so it seems like in some respects a clever rebranding or marketing effort because VPN and metro or wide area network services we've been doing this for ten years. So what's new here?

And also from the service provider's perspective it's like the word has been taken from us because we call the Cloud service basically all that the light touched. If I was a service provider for this metro, everything here is part of my data service cloud. Today when we talk about the cloud it's actually what happens in a datacentre. It's like that's the cloud location.

The difference here I think really is the achievement of what happens to scale over time. So the effort here is we're trying to move data services out of enterprise datacentres and we're trying to move and then to public or third party managed

datacentres. So somewhere out here there's a series of datacentres that we serve or somebody serves, Indosat or somebody, [COX], anybody but somewhere out there there's also a managed datacentre application or managed datacentre services organisation who then says come host with us, we provide these services in a broad manner. We can provide security and control and management and flexibility.

And so I think that's the big change. There was a statistic that was shared in our last MEF meeting that it indicated over the next few years some 80% of all Internet bandwidth or all bandwidth actually with services connectivity will begin and end in about 1000 facilities located around the world.

So what we see is some consolidation efforts here where we're going to reduce the number of corporate datacentres and collapse them and achieve some scale on a limited or a much smaller number of managed datacentre connectivity.

So the effort really from a service provider's perspective is what's new here, what's different and what's required to be successful in this space.

So today when we look at wide area connectivity services, the most predominant way to do that is a Layer 3 service and we use the Internet to do it. And it's cheap and it works. Its problem is it's cheap, it works and it's highly subject to threat and attack and we've had a series of discussions here. Once you allow open public access to a public network there are people with interest and motivation to do something with that and to secure advantage and to access to data that they would otherwise or should not otherwise have access to.

So large enterprises have always been hesitant to move mission critical apps to exposed transport. That's why the government has high sensitivity, corporate does, legal does, anybody with sensitive, time-critical information that when available to a broader audience could compromise their business or compromise their business plans.

So we've covered these in a series of other meetings, security vulnerabilities, data governance, regulatory compliance and from a service provider perspective we see a fairly high sensitivity around the performance attributes of the service itself. The Internet really doesn't come with an SLA. It comes with a statement of intent which says we are going to build a best effort service and I think cumulatively across service providers we build a pretty powerful best effort service network and it's global and it works.

There was a time, it was interesting when they first did video streaming apps there was a question you'll never be able to get high quality video over the Internet. Well, you can get HD quality video over the Internet; no one seems to be particularly disturbed by the quality by this point.

So it is capable to get a very high quality service through a best effort service. There does seem to be discussion when they talk about real time SLAs or real time service management that this is always an ongoing subject that we want access to higher performing bandwidth services.

But the net of it is corporate customers have done private. They'll talk about this concept then of we have public and we have private clouds. Private clouds essentially

creates secure environments where only a limited set of users are admitted for connectivity purposes. And you may use the Internet itself, you can build VPNs through the Internet or in a network that supports Internet services, but you limit access to a limited number of endpoint sites. And in that way you create separation and security.

So from the MEF perspective if you look today, when we go to build services and we're an access services provider, so connectivity between datacentres today if these were two Amazon centres, these were two anybody's two centres that connectivity today is typically done with Ethernet. It's done with the carrier Ethernet, CE2 that old compliance service. It's 1 gig, it's 10 gig, it's n by 10 gig and it's usually done with multiple providers.

It's usually done as a point to point service and it's openly transparent so that the datacentre operator has full control to creating and managing VLAN's. He can change and manage the service within this tunnel which is essentially a tunnel or pipe. And he's able to as he needs to move applications back and forth between the two sites or other sites or to other partners, and he typically does that at Layer two as a pure Ethernet service or at least with an Ethernet service that may be done over a direct optical connection. But they're high speed connections.

But they're also -- they're deterministic since he knows where its centres are. There's a limited number of them and you're able to build and manage bandwidth and infrastructure appropriately.

The other area of interest really is -- so this is typically when they talk about Cloud services they're up here. They're not down here. This is even when you look at what SDN proposes this is just called infrastructure. Everything that happens outside is just a line and we somehow magically we reach end point customers. And that's really the space that we live in and it pervasively is how we get to tens of millions of endpoint sites.

So as important as it is between datacentre connectivity and there's massive bandwidth between these sites, I mean the overwhelming majority of it here, and then actually all the individual clients sessions that are trying to get to these centres. This is also an area where Ethernet services can and are used today. You see an interesting mix of pure Layer 2 solutions that MEF services would be completely applicable or the Internet services itself.

There are an awful lot of -- we see a lot of large service providers, they will actually run a Layer 3 network. They'll build a private IP network on top of the Layer 2 service. So they may build a multipoint service or a VPLS like service or MEF would call it an E-LAN like service in which they would have all of these attached endpoints. They come back into what would be a service provider provided cloud or service area and then would provide connectivity back here.

In flexibility terms this is more of a fixed configuration. It doesn't change as much dynamically unless the host provider, that's the company itself or the cloud provider has control of the cloud and in fact is the customer of record for it.

But in terms of -- if you look at it like what does it really take to be successful for connectivity to the cloud using Ethernet services or Layer 2 services addresses a number of issues. It removes the security concerns. You have logical separation from other customers with malicious intent. You have SLAs for these services today typically. You can get those in Layer 3 but you always get them in a Layer 2 service.

And it's a large and growing area for customers because previously while it's easy to build high capacity fibre bandwidth to a limited number of sites in your network as you guys have reported frequently on what's the fibre penetration to commercial buildings throughout the world. And while there's some variations it's typically low. In North America it was in the low mid-teens and maybe at this point it's wandered into the 20% range.

And while it seems artificially low it actually beats current demand requests. If we don't have this ground swell of customers saying I need fibre to my building and why can't I get it is if you live in a fairly dense metro there's somebody there that can get you fibre. It may take a little bit of time but it's not a year, it's not six or nine months. It can be one to three months. So it just comes down to your patience and your willingness to want that service. But we continue to build fibre and the number of on-net buildings for us and every other provider continues to grow year over year.

A kind of interesting work is the large body of work that was completed in Carrier Ethernet 2.0 I think substantially finishes it, finishes the service development phase and it give service providers something compelling in that we're able to talk about the services in and between us. We're able with the wholesale Ethernet access certification to build and maintain, to buy services from each other, and to actually what it is we're going to get.

When you look at what's next, and I think at times it really coordinates with what's required for Cloud services, we talk about service extensions or we talk about APIs, we talk about what would be like the need for elastic bandwidth where a customer could then make a API or on-demand request for a change in bandwidth, the addition of PVCs or EVCs, the addition of tags or the removal of tags on services.

And I think when you listen to cloud presentations I think what you'll see going forward is a healthy combination of public you know via the Internet and then private Layer 2 based access solutions to cloud. I think MEF's role in this is what can we do to assist in the where does the API, what does the management interface look like. When we look at SDN and open flow and what's our involvement in this, how are we going to fit with this, what will be required of us to do this.

We had a similar investigation when we were looking at what would the implications be for LTE. We'd done Carrier Ethernet and specified. We were like what's going to be different about LTE that may require additional work streams on our part.

The net of it was we did an implementation agreement. We talked about best practices on what you would need to do to implement Carrier Ethernet or Carrier Ethernet 2.0 specifications in LTE backhaul.

We did recent work and finished the multiservice class of service specifications. And we see interest now from carriers for just that hierarchical bandwidth profiles and multiservice cost policies. And we'll start to see I think some of that in datacentre connectivity.

So in terms of -- again if you look at within the cloud itself, and SDN would call all the connectivity within the cloud, infrastructure, within that, within the datacentre and between datacentres Carrier Ethernet exists and it is [inaudible] and the service attributions specifications and work that was done by MEF exists there today. It exists on the connection between those centres and it exists on the connectivity to the customers in and outside of those centres.

And I think I'm particularly interested when we have these discussions about network security and in threat management. There's also this whole need of bandwidth management. Customers increasingly -- one of the first, one of the early, early movements towards cloud services were customers that had requirements to do offsite storage and data disaster recovery. So if you look at the west coast of the United States and California the overwhelming majority of the population lives in a fault zone. And every so often there's an earthquake and things are damaged and people ask questions about the viability of the infrastructure.

The overwhelming majority of the actual disaster recovery sites are in Las Vegas and Phoenix and have ample fibre connectivity back to those sites. And they're secure desert sites with ample access to hydropower and they're fairly secure and somewhat remote. But it meets those needs, it meets those needs. And those sites are overwhelmingly connected back and forth between those two geographic areas with high speed Ethernet services.

In terms of MEF's relationship to SDN we're certainly committed to the revenue growth of Carrier Ethernet and we understand or are trying to understand as best we can as the rest of you are to what the migration to the cloud means. We really do though we see it as an evolution of existing services and practices and it brings with it certain opportunities which we're keenly interested in. And we're engaged on both administrative and technical committees. Over half the equipment provider entities that are members of MEF are also members of SDN and so we share these same resources and we do have a continuing dialogue on these issues.

We had envisioned beyond CE 2.0 the idea of a dynamic automated metro Ethernet service and it seems like our timing was good on this one because it looks like it maybe something that will be required as part of both wholesale interconnection services as well as relationships with actual Cloud service providers who in many respects want to behave like a service provider, like a wholesale interface partner. It happens to be a fixed site location as opposed to a guy who actually owns infrastructure across a geography; he owns a particular asset that has a particular function within a geography.

But in terms of what he does and what he wants to do with us he wants to build and provision services; he wants to generate a large volume of orders; he wants to do it

continuously and he'd like us fulfil it very quickly. And we understand that. We understand that as carriers, we understand that as MEF.

So we talked about now Carrier Ethernet extensions. They are top priority. They were top priority because of wholesale services. I think we're starting to see Carrier Ethernet reach its maturity phase and a growth phase now. If you look at -- we had shared some slides in our briefing that the total amount of provisioned Ethernet bandwidth in the access network now exceeds that of all other access types and we hit this inflection point in 2012. So even though the total number of circuits and service is dramatically lower for Ethernet, the total amount of provisioned bandwidth that's carried by those circuits is now in excess of all other access types in both consumer or metro or wholesale services access, which is an amazing accomplishment.

But if you look at the implied slope of the growth it'll be twice that by the end of 2016. So we're in a massive growth phase. And when you compound that with -- it's not an unrealistic expectation. When you talk about what the expected number of mobile devices is, mobile data and mobile data backhaul is in its absolute infancy at this point. The amount of bandwidth that it's actually using while it's experiencing phenomenal growth is low relative to wire line. And it may take -- you know we tend to think of bandwidth growth on twelve month cycles, it may take -- and we call that a generation. It may take three generations of mobile growth and it may in fact at that point eclipse wire line growth, but that's speculative on my part for sure.

And I would just like to close. MEF and ONF are in discussion on technical alignment. We have an interest in the transport working group and some of the output functions of that. And if there's work that we need to do in terms of implementation agreements or the technical specifications or other liaison work, we're either doing it now or will be doing it shortly. That's what I have, thank you so much.

[End]