

NETEVENTS

EMEA PRESS SPOTLIGHT ON 'THE CLOUD'

FINAL

Debate Session 4 - Public/private/hybrid Cloud - choosing horses for courses

Peter Hall

Principal Analyst, Ovum

Panellists:

James Walker	President, CloudEthernet Forum
Khurram Ijaz	Head of Cloud Products - International, Rackspace
Mario Bianchetti	MEF

Okay. So good morning, everyone. Very pleased to be here to talk to you and rather to introduce the topic of public, private and hybrid cloud and then invite my panel to talk about some of the key issues across those.

So very quick agenda. I'm going to show a couple of charts initially about what's happening in terms of enterprises, in terms of the market dynamics. I'm drawing here on a couple of public surveys. Our own Ovum ones don't quite give the granularity I want here, so we'll talk about that.

And then I really want to introduce a bit further the concept of private - sorry, of hybrid cloud. I think this is something that's quite often misunderstood. I think Khurram actually did an excellent chart illustrating how that can operate between public, private and indeed a customer's own internal datacentre infrastructure, but we'll come back to that.

And then some questions that I'll ask of our panel.

So first, this is a survey which took place this year. Happens to be a survey undertaken by Ubuntu with large enterprises, asking them what their plans were for the next 12 months, so in other words their plans for 2014. And as you can see, quite a strong emphasis on private cloud and also people recognising this concept of hybrid

cloud which I want to come back to. And I think this is a - this is something that we're now seeing as the market for cloud computing really starts to mature.

This is another survey. This one happens to be done by RightScale, again with large enterprises, enterprise with over 1,000 employees. And this took a slightly different approach, quite an interesting approach. They asked what are your plans in terms of single public cloud, single private cloud and this thing that they call multi-cloud, which they break down in terms of multi-public, multi-private and hybrid. And again showing this - first an interest in companies now doing multiple clouds, which again is, I think, a sign of maturity of this and also the importance of this hybrid cloud model as well.

And in terms of hybrid cloud, I think there's a great misunderstanding in terms of what it really means. The 'gospel according to cloud' has been NIST for some time in terms of defining terms like private cloud and infrastructures of service and so on. But I don't find the definition that they've given of hybrid cloud particularly useful.

But quite before actually I knew that Khurram was joining the panel, I came across a definition from one of Khurram's colleagues in Rackspace which I think is quite useful. Hybrid cloud is the ability to mix and match different resources from factors like bare metal, dedicated servers, public cloud and private cloud into a common architecture solution for your application. It allows you to match the infrastructure to the workload, to optimise your security performance uptime and total cost for your specific application. So it doesn't just mean I'm using a bit of public and I'm using a bit of private. It's really sort of integrating those into delivering solutions using potentially public, private and indeed your own internal datacentre infrastructure. And I think we saw a good example of that.

Just for fun I quick had a look at Google Trends, which means you can put in a search term and see how commonly it's accessed. And interestingly, in fact, the - unfortunately the chart isn't showing the years very well, but I'm pretty sure that is - I think that's 2011, 2012, 2013 and we're just into showing half of 2013. The first term is cloud computing. And that maybe reflects, again, that at least the concept of cloud computing is maturing. I wouldn't say the industry is mature. As we've said before, it's still got a very, very long way to go in terms of adoption of cloud computing. But at least in terms of people asking, well, what does it mean?

If we look at hybrid cloud then you can see that graph is growing quite steeply. So we are seeing that increase in interest in hybrid cloud.

So in terms of some questions, I would like to understand the panel's view on the dynamics between public and private cloud, what is happening there today, and where the panellists see enterprises deploying this hybrid cloud model, which, as I say, is getting more and more traction. And one of the issues I think is quite interesting is portability between public and private cloud, how important is that to be able to take a workload and have it in one environment and potentially move it very easily to another environment? And then finally, how their organisations support the growing interest in hybrid cloud. We haven't got time to ask all panellists all of those questions, so we'll be a little selective about it.

My panel today, James Walker, who, in addition to the CloudEthernet Forum, is representing TATA Communications for the purposes of the panel today; Khurram Ijaz, who, as you know, has just given us an excellent presentation on Rackspace's activities across this whole area; and Mario Bianchetti from, again with dual responsibilities, MEF, but today representing Telecom Italia on the panel.

I'd like to just ask each of them to briefly introduce themselves and the role of their business in terms of cloud computing. I won't ask Khurram to do that because he's just given a whole presentation on Rackspace, so James, can you tell us a bit about TATA?

James Walker - CloudEthernet Forum

So TATA has a whole variety of different solutions in this space. We operate 42 datacentres globally and then on top of that I provide a range of services around datacentre interconnect and that also increasingly includes private connections for cloud to public cloud environments. We already have connectivity through our own internet backbone. We're the world's fourth largest internet service provider, so we have private peering with most key cloud operators. But this is also private connectivity. So we're increasingly looking at how we can provide a blend of public and private cloud infrastructures and we ourselves also provide cloud services under the TATA banner.

Peter Hall

Thanks, James. Mario?

Mario Bianchetti - MEF

I'm Mario Bianchetti and I'm representing both Telecom Italia and MEF technical committee.

Peter Hall

Have you got your microphone on? You're using -?

Mario Bianchetti

Better now? Okay. I'm representing here both MEF technical committee and Telecom Italia. I would like to emphasise the role of cloud provider or the cloud broker, as sometimes they're named, in the sense that obviously you need a very strong networking between the layer of the cloud service provider and the cloud consumer to get the availability and the possibility of dynamically changing the connection and the bandwidth. And I hope to have time to describe what is done now in MEF and what we have done in TI on this part.

Peter Hall

Okay. Thanks, Mario. I'll come back to you on that. Khurram, can you - my first question is really in terms of what you're seeing in terms of the dynamics between

public and private cloud. Can you comment on that? You're clearly - Rackspace is a major player in this industry and so you have, in this room anyway, quite a unique insight into those dynamics.

Khurram Ijaz

Sure. So can you hear me okay? Basically what we are experiencing with our customers is that they would like to choose the right solution for their use case. It's not advocating just it's either public cloud or nothing, it's either private cloud or nothing. So to give you an example, let's take the example where the customer has higher requirements on better security, more flexibility and control over their environment, then something like a private cloud deployment on OpenStack in either their own DCs or in Rackspace datacentres would address those needs.

However, if a customer has a use case on which they have requirements around utility billing, low per-unit cost and time to market or time to provision or deployment is really important, then public cloud helps them address that.

And I hope that layers out the different dynamics that customers are considering when they are trying to choose between a private or a public or a dedicated environment. If it is a use case in which customer knows that they would probably be utilising their resources at a 90% level about 80% of the time then it is best for that customer to use some dedicated physical resources to help meet those needs.

Again, I think it lays out, and I think you pointed it out, Peter, earlier in one of your slides that the hybrid cloud is all about having a single interface, having an environment that can work seamlessly across dedicated public and private for the same application or the workload. And again, so I think it's down to what the customer's looking for, security, flexibility, better control, private cloud helps fit that bill. If it is time to market, lower cost and utility billing pay as you go, then public cloud helps them meet that.

Peter Hall

And, Khurram, do customers always get it right first time or do you have situations where they might start in a public cloud and say look, I really think this needs to be a private environment?

Khurram Ijaz

I think that's a really good question, Peter. So - and again, that's where customers need help. That's where a hybrid environment helps customers actually experiment, test, run proof of concept in a safe environment. So going back to something earlier that I said around our private cloud distribution being downloaded close to about 35,000 individual downloads since May or April this year, that is customers playing with the OpenStack environment in their DCs and trying to experience and experiment with that environment.

On the top of that, with somebody like Rackspace, where we are able to provide fanatical support to that customer, we are then able to help them understand what is

the best migration route for deploying applications on the cloud. And that's where customers need help. That's where customers are saying, and I think to a point that Manek raised earlier around the scepticism in the industry, that scepticism is actually not entirely unfounded because customers are saying, well, you're asking me to go on the cloud and let go of my environment completely which I'm running in house right now. But how, as a service provider, are you able to help me do that? And that's where I think Rackspace, with its fanatical support promise, helps customers navigate through some of those challenges.

Peter Hall

Thanks. James, can I ask you what is TATA seeing in terms of the dynamics between public and private cloud? And to pick up a point that Khurram made, how are you helping customers in terms of that journey themselves?

James Walker

I think what we're seeing is that customers really are doing cloud things irrespective of whether it's officially sanctioned within the organisation or not. So there are often parts of the IT department that have gone out and spun up a few VMs somewhere and they're doing things anyway. So for a lot of existing IT departments it's a matter of trying to bring some of that under some level of control, or it's part of the company's formal strategy or a part of their formal infrastructure and therefore they're bringing it in a much more integrated and controlled sort of way.

So my general view is pretty much everybody we talk to has an extended or hybrid cloud environment. It's just a question of whether the company knows about it or doesn't know about it.

I think also we've talked a lot here about things which are to do with infrastructure sort of things, that is moving VMs and storage potentially and that sort of stuff into cloud. A huge part of it nowadays is also software as a service. And having reliable connectivity to a service which - to a software service that the company relies on, like Salesforce.com or something else, SAP and so on, that you've got to have a predictable performance if you're going to run your company on top of that.

So things like Office 365 and Azure are interesting that they come out of the same company but they're completely different approaches to something the company might want. So integrating those external applications in a way that every user within that company's estate can use it reliably, predictably, not be beholden to vagaries of internet performance and all that sort of thing, that's a big question that's being asked a lot now.

So perhaps, I suppose, to sum all of that up, there is - hybrid cloud is a reality for almost every organisation we deal with, so that's whether it's official or not official. The question is more about how can I look after it? How can I enforce port policy? How can I enforce security? How can I enforce regulation that I need to be compliant to? How can I make sure that I'm not putting my company's secrets out on the public cloud and then potentially having them at risk? And then if I need to use a software

as a service platform then I want to make sure that I can have performance that my users can tolerate and that I can somehow manage. So it's really a lot of it comes under control management visibility, those sorts of things.

Peter Hall

Thanks. Mario, I know you wanted to say something about how Telecom Italia views this area and also how you're helping customers deliver a hybrid cloud model.

Mario Bianchetti

Yes. Public, private, hybrid that's the point. Within Telecom Italia now, we have both the role of cloud service provider and cloud carrier. And we think that it's very difficult to foresee which will be the model, deployment model for cloud activities. And from a telco point of view, I think that we must be prepared to be able to manage all the models that will be adopted case by case as Khurram and James explained.

So going back also to the MEF activity and technical committee, which are very important and we are very in line, I just want to tell you that in the technical committee MEF has already published the a document. I would suggest you to have a look at this. It is called Carrier Ethernet for Delivering Cloud Services. This is more or less an overview of how the already standardised MEF services can be used in the cloud environment.

More of this in MEF, we have an active project on cloud that is named Dynamic Responsive Ethernet. DRE refers to the managing Ethernet service on demand. This is very important because hybrid, private, whatever, I think that the customers, the end customer, for sure don't want to be bottlenecked by the network. So the network must be open and flexible to be able to move between datacentres in the case of workload, as Khurram showed in the panel before, and so on.

And to do this you must be ready to dynamically manage your connection in that sense. For instance, we have deployed within Telecom Italia a prototype solution. We call it elastic cloud networking and [Crystal Pipe5] in which our cloud service providers can manage by themselves the provisioning of the service. So they can decide to which datacentre has access and change it. Obviously seamlessly from the point of view of the network. They only thing they have to do is to change the parameters in the two end points. That's a very powerful way of being able for the future scenarios, but I think also during these days it's not completely clear. So from the point of view of the company, of the telco company, it's very important to be prepared to all the scenarios.

Peter Hall

And a follow-up question again to you, clearly as a telco you're relatively new in this space, certainly compared with the companies in the industry, compared with the big player we don't name, and many of the other players that have come from an IT background or a hosting background. To what extent do you think, as a telco, you can really help enterprises in terms of formulating their journey to cloud and what they

really want of it? That's the sort of added value that I think enterprises really need to be able to do cloud effectively. As a telco, do you believe that you can do that? Do you have the skills to do that? So Mario first. And also I'd like to ask James to comment on that as well.

Mario Bianchetti

Yes. I think that the better way to manage this is to work together, and MEF and CEF is a good example of converging on the same table for having common discussion because, from one side, we have the IT requirements. On the other side you have the networking capabilities. You have to match them. And the only way to do it is to share the thing that can be done now and the thing that can be done in the roadmap of our vendor in a few months, and so trying to figure out the better way to reach the final goal.

Peter Hall

And James, do you believe that you really have the skills to be able to sit down with CIOs and really help them both strategically and tactically map out their cloud journey?

James Walker

No, not particular, so that's why there's a sister organisation called TATA Consultancy Services that's got 270,000 people that do that for a job. So as a communications organisation we're focused on communications, which may not sound so exciting, but it is a very fundamental bit of connecting cloud environments and customer environments together. So we don't want to downplay that role.

In terms of the strategy that I run in my area, my objective is to be as much of an honest broker as possible. In other words, I want to give the customer every option to use their own internal resources, to see external public cloud resources as an extension of their own datacentre environment and make that as seamless as possible, and then in turn then give them the ability to reach to whatever software as service or cloud-based services they want to reach from third parties and make that again something that feels like it's on-net to them and has all the performance and particulars around that.

Within that, at the end of the day, I am equally happy if they take a cloud service from TATA or they take a cloud service from Salesforce or they take a cloud service from Rackspace or whatever. I don't want to be the one that's forcing them down a particular track. And I think the danger that a lot of telecommunications operators have had is to say I've got a cloud service, I'm going to force you to use that because there is no organisation that I've come across that has 100% of its use cases covered by a single provider, so it's unrealistic to do that.

So if you start from an opening position that says hey look, I've got to work with everybody else, then things are a lot easier and, indeed, as Mario said, extending that out across the industry through the use of forums to ensure interoperability, and as I

said yesterday, particularly end-to-end operation visibility and management across multiple pieces of the puzzle, that's really the challenge that, as an industry, we need to overcome in order to make cloud truly viable, I think, for enterprises.

Peter Hall

That's a refreshing honest answer. Khurram, I alluded to the issue of portability between clouds, and I guess in a sense the Dominos example was actually portability between a cloud environment and the in-house or the dedicated hosted environment actually wasn't it? But in terms of public and private cloud, how important is portability? And is it something that customers need to be able to do in real time? Do you need to be able to take an application and move it from one environment to another on the fly almost? How is portability - how is it used in practice?

Khurram Ijaz

So by portability we are - just to clarify, right, and just to level set on that, by portability we mean that customers have a particular workload that they're running on a public cloud environment and then they're able to implement the same thing and move that across to a private cloud environment if needed, and vice versa.

Yes, the dominoes example that you've quoted, Peter, is more a case of how can the customer use a dedicated and a cloud environment for a single application. With regards to portability or workloads across cloud environments, it is quite an important consideration for customers because that helps them again choose the right platform for the use case or the environment that they're looking to build.

Having said that, in terms of the way customers do it right now or the considerations around that, more often than not it's not really done on the fly. I think it's just an option that customers like to have in their back pocket so they know they can use that feature or function if needed. So, for instance, you can take an example of using a public and private cloud environment, a combination of the two, for implementing some sort of very basic disaster recovery mechanisms within the organisation. So you are running something internally, you take a snapshot of that environment and implement a base level of that environment onto your public and private cloud, so if disaster were to strike your primary site you could still get some sort of a reduced service, which will keep things going from within the cloud while people work and fix your primary environment.

So yes, it's not really looking to get stuff done on the fly, but it's more around an option that is available to customers. What customers do work on in terms of immediacy and the need to do it in real time is the ability to scale their environment. So for instance, within the cloud environment our customers are able to set in thresholds and say well, you know what, if my environment now reaches 80% capacity utilisation, then there are API calls that are made to effectively throw in another [five] cloud services, which help address and meet that peak in demand from whatever application the customer is running. So that's more stuff to do with auto scaling, which is done in real time, and that's spinning up resources and scaling down.

Peter Hall

Very last question, because I know we're running out of time, just as a follow-up to that, right at the start of your presentation, Khurram, you said you alluded to portability between suppliers. How real is that today and where have we got to go in terms of customers having the reassurance that if they need to move a workload from one supplier to another, they can do it without a huge amount of effort?

Khurram Ijaz

So that's asking a really good question, Peter. On that I think we need to consider a number of things. First, given where customers are in terms of using cloud and given that it's very early days, I think providers need to, at the end of the day, still work with one another to make sure that that migration happens if a customer makes a request for it. Absolutely, the platforms can work with one another if the customer has developed a workload using a standard API. But at the end of the day it will still be a case of human interaction between the customer, the existing supplier and the new supplier all coming together to make sure that the customer still gets value out of it.

And again, I think it's in the interest of the cloud industry to make sure that we can work with one another because traffic will go both ways. You will have a customer that leaves you and goes to another provider and probably you'll get new customers coming onto your platform as well. So yes, I think the industry needs to work more on it. But it must be viewed in the context of this is still early days within the current industry.

Peter Hall

Okay. Thank you. So questions from the audience?

Unidentified Audience Member

One question that hasn't been discussed since we've been talking about this is confidentiality and integrity. One of the challenges is in the [post-node] era is that when networks are now becoming untrustable and the integrity of the carriers is being questioned because their networks are no longer secured through to now that the US government has been shown to be performing commercial theft to pass data from one carrier to another. In the panel, have you got any statements in your - in terms of how the cloud Ethernet organisation's going to work about creating security and integrity and confidentiality as part of your processes? I haven't seen anything in the briefing material I've received.

James Walker

No is the short answer, partially because a lot of what you talk about there is really activity that takes place at layer three. And so from a layer two point of view, we're not actively looking at the payload today. There is some consideration or there's been some discussion that's been had around sorts of things that Google's starting to do now,

for example where they encrypt connectivity between datacentres. That we've considered.

Within TATA, from our point of view, when we look at certain multi-cloud or in fact private cloud connectivity, private connections into public cloud and all that sort of thing, the security is really something that, at the moment, we're leaving up to the customer. However, I do see a need for it and it's a very specialised area of knowledge that requires quite a lot of work, often on the customer's side.

And one of the biggest problems we have there is the customers actually understanding the criticality of their information. So a lot of customers don't understand how important their information is or what value indeed is associated with it and so they very much struggle to understand how much money to spend protecting it. So they quite often will say I've got all this critical data, it's absolutely important for my company and I don't want anybody to see it, and you say great, it's going to cost three quarters of a million dollars to protect it, and they go oh, I don't want to spend that.

So until we can get to a more informed view within organisations of what the value of data actually is, the only vertical I've come across that's very aware of the value of its data is the banking and finance industry. So they can accurately quantify, because many of them have been fined, with what the dangers are of exposing particular types of data to being taken and stolen and so on.

Khurram Ijaz

Sorry, if I could just add to that. So as a service provider, Rackspace is network-agnostic, right? We work with multiple carriers on that. As far as our customers' data is concerned, our policy is very, very simple, that our customer's data is their data. They choose where they can host that data. We do not automatically put it in different DCs. Our customers choose it. And our position is very simple, that we only respond to requests from courts which have jurisdiction in the country of operation where Rackspace is operating from and where that data is held.

So yes, we are assuring our customers that things that you may have seen around [Prism], our position on that is very, very simple, that we will not respond from within the UK to requests that we receive from the US.

Peter Hall

I think there was a question on the second row.

Brian Levy - Juniper

Brian Levy, Juniper. Apologies for my cold first, so you can understand my voice. I was looking at some of the recent analyst reports, talking about the barriers to cloud adoption. A few years ago it was security, and actually now security was number two and compliance was number one. So I think things are changing. And I'd like to ask the panel just one example.

My wife uses DropBox. A lot of us do. It's an application on many tablets and PCs, you put your data. And I said to my wife, do you know where this DropBox is? She hadn't got a clue. It could be anywhere in the world. Now ask yourself about a CIO, because when you go to the cloud, and certain with software as a service, you are outsourcing not just the application but the data. You ask a CIO to talk to his board when he's got fiduciary responsibility for customer data and you're putting it anywhere, and you go to the board when it's lost, how long are you going to survive as a CIO?

So I actually think you talked about these things, but it seems to me that if you're going to really make this thing scale, you've got to really address those problems of security compliance in these services. What's your views?

James Walker

Yes. So I think that's very much what my view and the position that I hold is, is that there was a very big growth in terms of public internet-based access to public cloud services. One of the interesting things I've found about working in cloud services the last three, four years is that almost everything that people stridently tell you is true about cloud is almost instantaneously proven to be wrong, usually almost immediately.

So people have stridently told me that there's no interest whatsoever in private connectivity to public cloud environments, and yet I've got tons of customers who demand exactly that. Because of the compliance requirement, they have to understand where this data goes. And what's more, it gives them an environment where they can force data to be seen by the compliance engines that they have in place. As soon as it goes onto the internet then they don't know what happens to it and they can put in a block in place that says anything that's going to DropBox, is going to any of the other competing similar services instead is going to be routed internally and is going to pass through some kind of inspection and we're going to decide what we want to do with it. So that's one specific area that we see huge demand for.

You talked about compliance around where things are being stored as being important, but it's also the security of those organisations. Do you know that DropBox employs the best security people for your data? That's a risk that the people have to assess. And it all comes back down to the value of the data that you hold there. How much are you willing to spend to protect it?

The issue to date has been really convenience, that DropBox is so convenient and it works really, really, really well and there's not been a competing enterprise product that's been made available yet. So anyway, where we're coming at it from is given the option of private connectivity into the bulk of the public services, at least given the option back to the IT department who enforce compliance through that methodology.

Peter Hall

I suppose the counter argument is that Salesforce has become the most successful SaaS company with many, many customers just accessing over public IP. And that must be the most precious data that a company has, its customer base.

James Walker

Except that - I agree with that aspect of it, but there's also if you go to any organisation today that is global in nature and ask them about using Salesforce.com, then their biggest complaint will be, well yes, I signed up with Salesforce, I got associated with the EMEA server and as a result, everybody using either Asia or the west coast of the United States has a terrible performance experience because it's not properly globally distributed. So you can alleviate a lot of that again through looking at much more controlled access to it.

If the data is so important to you then the question of their ability to protect it from the internet and all that sort of thing aside, predictable performance which is fast and known on a global basis, that becomes really important. So I think people get too focused on hey, it's all out there in the cloud and I don't really need to understand how I get to it. You do absolutely need to understand how you get to it if you're going to run your business on top of it.

Peter Hall

Excellent plug for TATA's networking. Last question?

Thierry Auton - Atlantic News

[Thierry Auton] from [Atlantic News]. I just wanted to know if you have any plan to sign an agreement with some insurance company, because with simplification usually you say we are just arranging the freeway to let your truck run, but it's a different freeway across the jungle and you could duplicate the track. So is there any program to create like an insurance service of data or anything like that?

Peter Hall

Khurram, can you make a comment on that?

Khurram Ijaz

So if I've understood the question correctly, was your question more around the network or the actual platform?

Thierry Auton

[Inaudible]. Usually we speak about risk management now more than security, and anyone is aware about the difficulty to manage that. So just have you any plan to give insurances around management of data?

Khurram Ijaz

So we address this issue on two levels, right, and I hope that answers your question. So we address that issue on two levels. So first of all it's the service-level agreements that we base our services on, both on dedicated and in the cloud environments. And then the other one is to make sure that when the customer is actually deploying their environment onto our platforms, we help them architect in the best possible way. So for instance, on the cloud they're deploying a cloud aware application, an application that can help, in a way, detect and pre-empt failure if it's coming to them. It can take some actions to make sure that the application stays up and running.

See in our case it's a combination of service-level agreement coupled with best advice before the customer actually deploys their platform.

James Walker

I think it's interesting because we quite often, in RFPs and in networking terms, used to get things that said unlimited liability if the network goes down. But that's not really my problem; that's a budget limitation problem on your side where you've decided that you don't have enough money to buy services from two or more service providers, but you want to outsource all the risk of your enterprise onto the network service provider. It's not a realistic model. There's no other industry in the world that accepts that as a way to operate. And indeed, if it ever came to the court, I can't see the court abdicating all fiduciary responsibility to the service provider in the same way that it's not done to consulting firms and so on.

I would apply that to cloud also. I think it's completely unrealistic for an organisation to say okay, I'm going to actively make a decision to outsource a whole great big pile of my internal IT infrastructure and then we're going to abdicate all responsibility myself for actually looking after that data and rely entirely on the third party service provider to be responsible for that and to provide all the guarantees back to me.

For 20, 30, 40, 50 bucks a month for a virtual machine, or less than that, why on earth would I take on 20, 30, 40, 50 million, 150 million, 200 million worth of liability? It doesn't make any sense to do that. So these things need to be in proportion. If you spent \$500 million with me because you're outsourcing your IT to me, am I going to provide guarantees around the security and management and all that sort of thing? Yes. If you're going to approach me and may be \$30,000 a month for a global MPLS network and then say it was my fault if your business fails because you can only afford a single connection to Dubai, that's ludicrous. Nobody's going to take that model.

So somewhere in between that is a collaborative model that says if you tell me what your business objectives are, I will design the network, the cloud environment, the solution, whatever it is to meet that business objective and then I'll tell you how much it's going to cost. Your decision then needs to be whether you can afford that or not.

Peter Hall

Okay. Thank you. I think that's, okay, the end of this panel. Thank you.

[End]