

NETEVENTS

APAC PRESS & ANALYST SUMMIT

FINAL

Debate III: Seven Secrets of Optimal Data Centre Design

Chair: Clive Longbottom

Service Director, Business Process Analysis, Quocirca

Panellists:

Kevin Buckingham	Data Centre Portfolio Manager, BT Asia
Bruce Bateman	Networking Evangelist APJ, Dell Force10
Mark Pearson	Chief Technologist Data Center and Core, Advanced Technology Group, Networking, HP
Pranay Misra	CEO, Nanotel
Dr. Atsushi Iwata	Assistant General Manager of Cloud System Research Laboratories, NEC

Manek Dubash

Welcome back to NetEvents ladies and gentlemen. I hope you enjoyed your coffee and your Dim Sum. Now, we've talked a lot about networking, we've talked a lot about the Cloud, but none of this would be possible without the data centre, which is where it all happens. We are all talking about data centres today, forget individual computers. And data centres are complex ecosystems of infrastructures overlapped over infrastructures. So, how do you design a good data centre? How do you make it all work properly? We are going to find out the seven secrets of data centre design. Clive Longbottom, Quocirca, come on down.

Clive Longbottom

Okay, if my panel could please come and join me up here. So, yes, I've been asked to come along and talk about the seven secrets of data centre design. However, the first secret really is that a lot of what you're going to see should not be that secretive to you.

You should know most of this already. The second secret is that a lot of it is not about data centre design per se, it's the impact of various aspects on the networking within a data centre design. So I won't particularly be getting into things like free air cooling, how to actually go around the facilities management side of the data centre, what the building should look like or anything like that. So I apologise for that in advance.

But secret number one, Cloud Computing changes everything - or does it. Well, you could look at Cloud Computing and say if you moved to a full Cloud Computing model you may not need a data centre at all any longer. That's somebody else's responsibility. You could look at it and go; well we actually want to keep some parts of the compute capability without our own data centre, but some parts we'll outsource to a more public environment. And what does that impact?

And then you could say, well once we've decided that point should we then have our own data centre building or should we go for a co-location type environment. You've also got to look at what the impact is on usage. And I think that Cloud begins to really drive differences in how people approach their work.

So once you've got a Cloud model, then it's accessible from anywhere by anyone at any time. So again it does bring in the security issues that we've already spoken about, but it also means that the workforce can be far more mobile, it can be far more virtual; you can bring in the customers, the consultants, the contractors, the suppliers into that whole model.

So I do think that Cloud changes everything, except for when we look at what Quocirca's research shows, which shows that about one in seven large and very large organisations in the world are saying that Cloud has no part in their organisation's future - ever. There is far too much confusion out there about what Cloud really means, what Cloud really is.

So, Kevin, when looking at Cloud and the confusion that it's causing out there, what are you seeing out there in the market?

Kevin Buckingham

Yes that's quite an interesting conversation, because a lot of BTs customers are looking for different things. It depends which vertical we are actually looking at. If you are looking at financial, yes, there isn't a massive amount of interest in there, but if you're looking at the SME markets then there is a lot of interest in there because they are trying to take costs out of their organisation. They are looking for faster routes to market. They are looking for much greater agility.

And what that's driving is a behaviour where these customers are looking for optimal solutions. So we are looking purely from a data centre perspective, we are looking at initially smaller footprint but with higher power requirements. So, 10kVA per rack is starting to become the norm for Cloud deployments now. And that is a little bit of a headache, especially when you're in Asia Pac. You know, we have power challenges. We also have space challenges in Hong Kong, not to mention the cooling challenges. So it's quite a different behaviour than what I find in the rest of Europe, although the

virtualisation is approximately the same, the consumption of these services is uniquely different.

Clive Longbottom

Yes. And Pranay when we are talking about changes in the usage model, the way that people work, when you look at India and the complexity of the number of mobile operators that there are there – and how widespread mobile access isn't, as in it's only really around the cities. Do you see Cloud as being an enabler for India, or do you see this as being a fairly large stopper when going forwards?

Pranay Misra

Okay. I think Cloud Computing is in a very initial stage in India, even though the old conventional data centres are still running up. But the only issue is the energy and the OpEx cost is very high on the old conventional data centres. And I strongly believe that Cloud Computing would be coming up very soon in India.

But there are a few issues that's [calibrating] the security and the flexibility, and which model do you need to up there. Is it a sharing model, because Telco's have a mindset now to go everything sharing, because now if -- because if you have 10 operators in India where everything is being shared, their passive networks and their active networks, their internal IT has been outsourced to IBM and you know kind of extra. And in such a scenario I think a sharing business model on the Cloud Computing in India would really click very soon.

Clive Longbottom

Okay. So if we can go back to my slide. If we then look at the sharing model, not just the networking sharing but the services sharing as well that brings us onto secret number two - that the application architecture has to change.

If you start to move into a more dynamic model where bringing functions together on the fly to create the composite application is key, it brings us back to what we were talking about five, six years ago with service oriented architectures and web services, loose linking rather than hard linking together and then being able to aggregate everything together on the fly.

Now from the earlier discussions it sounds to me like SDN, OpenFlow has quite a bit to play here. Atsushi, what do you see happening when it comes to an application being built from multiple functions brought in from different service providers to create this composite application.

Dr. Atsushi Iwata

The -- actually SDN can become the good [inaudible] on this area. The OpenFlow can have a lot of API and then you can integrate the application to initiate some of the program on top of the switch equipment or [allowed] equipment. And then therefore it is easy to integrate application space to infrastructure space. That is the big changes of the industry, like the software application space and the infrastructure space is

completely independent. You have to [wait] network instalment and then you have to run the application on top. But from now on -- this is -- the order is different. Like, you should think about application first, then you can have customer made or order made network infrastructure based on the network application and the application requirements. And that is kind of the hybrid model or the integrated model.

Clive Longbottom

So, Mark when we look at bringing in the wide area network in the earlier discussion, in Manek's discussion, it sounded like that wide area network is not quite there with the OpenFlow SDN model as yet, but certainly when we start looking at hybrid Cloud that's going to be a necessity. How rapidly do you think that we are going to be able to get into an environment where the network end to end, across that hybrid model, will be able to support something like this seamlessly?

Mark Pearson

And so as we mentioned earlier the application deployment is now fairly automated, has lots of policies but a lot of the delay comes in the network infrastructure including the wide area for, especially the hybrid Cloud scenario. And so as we can speed up that automation of both the data centre connectivity to the network and the inter-data centre and hybrid connectivity, we can use the centralised policy tools to deploy that and to interact with the other networking systems.

Clive Longbottom

Okay. So then one of the big differences we are beginning to see within the data centre is the move from just buying a 19 inch rack and just populating it with pizza box, 1U, 2U, 3U units, putting a load of power distribution units at the bottom, putting storage into different racks or wherever it might be, to the vendors now coming through with highly modularised, highly engineered systems, whether that would be a single rack, whether it would be a single row, whether it would be two rows put together with a hot aisle, a cold aisle or whether it be a container or a pod or something like that.

This brings together in Quocirca's view the capability to optimise the CPUs, the storage but also the network. And so we are seeing people like yourselves as Dell and HP and NEC coming through with these systems with top of rack and end of rack networking fitting in and engineered as part of the system.

Bruce, do you see this as being the future for how to put together the computer environment within the data centre.

Bruce Bateman

You know, I think what we see is that the data centre requirements change from previous where there was more fixed environment. An IT manager or network manager would buy what he could forecast, because that's what he could get for a budget. And in today's market we want to build or we want to pay as we grow. And

the pay as we grow model I could put a data centre, container centre at a location today, it runs at 80%, then I can tomorrow very quickly, Dell provides fully containerised solutions, you just order the solution you want, container shows up, you plug that in. That kind of model, where I as the IT guy, don't have to think as hard about where I am today and where I'm going to be tomorrow. I can just say I live today, because I don't know what's going to happen tomorrow. And then tomorrow comes I can scale up.

But one of the things you also have to think about is scaling back where the market isn't here today anymore; I need to move that from Hong Kong to Macau. I need to move it from Hong Kong to [Shenzhen] or somewhere else. So that's the other nice thing about modular data centres is that we can move them around based upon traffic requirements, end user requirements.

One of the things we talk about is that data centres today are about space, and space is about density. When we talk about speed, speed is about the latency and that's really about flat distributor core networks running now at 10 and 40 gigs. So, everything sort of merges together now in the data centre, and it has to be able to be fast, there has to be less power and it has to be able to expand when it need it, where I need it. And that's kind of the change from the old data centres that were fixed brick and mortar data centres.

Clive Longbottom

Okay. So Bruce brings up a fair set of different things there. Mark, one of the things that you mentioned earlier was OpenStack. When we are looking at moving these workloads around from Hong Kong to Macau or from Macau to Mumbai or wherever it might be, if we are going for a hybrid Cloud model and we've been working on the Amazon AWS and now we want to move over onto a Microsoft Azure-based system, it can't be done.

Now, all of sudden, a lot of vendors have jumped onto the OpenStack model as the underpinning to some as well, again, when you are then looking at bringing in open networking as well through the SDN and the OpenFlow model, is this really the future or is this something where there is too much revenue involved in maintaining proprietary approaches and keeping people on a specific platform that vendors are going to be going "No - We are going to put in something in place which makes it very difficult to move away from us."

Mark Pearson

Right, so I think the open system is important, because we need to give the choice to the Cloud manager in effect. The OpenStack connection here is that it provides an open system with lots of integration points, including the network integration point where we talked about the centralised control plan. And so the combination of openness, the API enables these Cloud stacks layers, as well as integration into the proprietary systems or the hypervisor vendor systems.

Clive Longbottom

So, Kevin just to pick up on that then, I mean a lot of where BT fits in is externalised data centres, BT has its own Cloud environment which people can host applications on. It's moving into the public environment where you'll be able to just self-service provision. But you've got to be able to work in this hybrid environment as well. So do you see that this is a major issue, the interoperability between data centres?

Kevin Buckingham

I actually do. And I mean to pick up on, sort of like the, OpenStack debate. I mean, traditionally the data centre business is quite a sticky business. Normally once you have a customer if you provide good service they will stay with you. Not just purely because it's a good service but they invest a lot of effort and a lot of money and a lot of time in configuring those services.

But if you was to try and make that service too sticky by using a closed stack, I think the customers these days are much more technically savvy than what they were say five, even 10 years ago. And they are looking to future-proof their systems. They are aware that at some point in the future they'll want to pour onto something else, whether that's the same kind of stack or whether it's going to be something completely new in the future that isn't out there at the moment. But they are desperate to make sure that whatever stack they chose they have got options for moving on in the future.

Clive Longbottom

Okay. So then when we start to look at how we design the data centre for optimum performance, particularly when we are looking at a hybridized environment, how have you got to build the components together, that networks have got to flatten. We cannot have multi-tier environments. We cannot have spanning trees in place. Every hub or switch introduces by the laws of physics, some latency.

And once you start to say, right, yes some of our services are going to be served out of Tokyo, some of them are going to be served out of London, getting round those laws of physics can be pretty tricky. But the more we can flatten the network the less latency we build in at source.

So, for when you are looking at mobile workers, Pranay, is the flattening of the network important or is it a case of well it's only such a small part of the jigsaw that there is so much else that needs to be looked at?

Pranay Misra

I think flattening of network is, in one way, it's important I believe that is for sure, because having a multi-tier one is not going to work out in terms of cost model. This is, I believe, very strongly. And as I told you that, because I am working very closely with many Telco's in India, and the only discussion point of view is it's all about sharing model, that how exactly the OpEx cost would be controlled there.

And the other area what we were seeing was having a local, a one [clear] language app store on the Cloud, rather than having on a conventional data centre. So there was an issue that when it goes [onset] on the Cloud and when the end consumer try to download it using its tablet or a Smartphone, so how the model is going to work.

Clive Longbottom

So if you then look at that shared model and you have somebody who has a highly performing network based upon a very flat topology, using an SDN type environment so abstracting, putting a lot of the rules outside of the hardware, being able to give more performance, and then next to that, running parallel to it you've got a three tiered basic model with all the latency and everything else, who should share and who shouldn't share?

Is it a case of, well I know I've got a really bad network, I'm very happy to open up and share as in use somebody else's network. Or am I sitting there saying I've got a really good network. There is no way I'm going to share with you because you're not investing with your environment. How do you manage that sharing?

Pranay Misra

One thing is the investment, but the other concern is also the security one. So today the security is a concern there because its data base of the subscriber, especially the Telco [is mine]. Let's assume there is one service provider who has a very good network on the Cloud, and if I go and share with him the concern for the Telco is my data base, because the data base is a strength for Telco, so will this be leaked. That is a concern. But, nevertheless I think the multi-tiered is not going to work, the flattening of network, what we discussed, because in terms of OpEx model it's going to work out.

Clive Longbottom

Sorry, Bruce yes.

Bruce Bateman

Yes, sorry. I think one of the big differences has been the traffic is now more East/West than before. A query will come in and I need to send it to a cluster, and so bringing, flattening the network, bringing the user closer to the cluster where the actual data resides, whereas the old days we saw a lot of North/South traffic. And today we see -- because our users want faster responses, so yes we had to take the physics out of it with the cabling issue, but because the traffic is now -- and we've now gone to cluster servers to cut that down, so that's why its East/West traffic is much faster.

Clive Longbottom

Yes. We see that it takes us back to that application architectures have to change. If you have an "n" tier application architecture, it's never going to be performant in a Cloud model. If you say, right, if we can keep the data and the business logic next to

each other, if we can use essentially a data bus at the back end of that and only put the presentation layer to people, then we'll get the best performance - which is based on the mainframe model of 1970.

Bruce Bateman

That was a [inaudible] programming.

Mark Pearson

Clive, can I just come in on sort of like flattening the network? I mean from an end customer perspective, and we bump into this all the time. So we can speed up the infrastructure within the actual data centre itself both the servers and the network. But what's been pushed onto the service providers is the latency between the different countries and the different continents.

For example, if you provide a service in Hong Kong and you want to link up with the UK typically 350 milliseconds round trip. And that is crippling some of our customers. Now, okay, so if you're a banking customer or if you have confidential information then you will have to store that or you may have to store it, it's likely that you will have to store that information within the country of origin. However, you still need those links back to a central processing and that is really difficult for our customers.

So what their having to do, and this is where we could take cost out by flattening the network, if we could provide a low latency link that will get back to certain continents that will take the cost out for the customer rather than they having to duplicate infrastructures within different countries.

Clive Longbottom

And there's all sorts of other things which add to this, data dedupe can mean you are making better use of your network.

Mark Pearson

Oh yes.

Clive Longbottom

Packet shaping, so using large packets, data caching for more static information means that there is less information going backwards and forwards, keeping everything less chatty.

Mark Pearson

Yes. I think that the customers and the suppliers have done basically everything that we can by having networks that are optimised for the different applications. We've put in so many different kinds of things to help our customers. So we've now got to a saturation point, we've got to that bottleneck of trying to get out of the data centre.

Clive Longbottom

Okay, which I think brings us through to the next thing which is static SLAs are no longer any use. Performance creep is something we've had for a long time. This is where you agree a service level with your users; we'll say you know 350 milliseconds is fine. And they start off going, wow this is great. In a year's time they are coming to you and saying everything has slowed down, it's truly awful we can't cope with this any longer, what's gone wrong.

So you measure everything that's happening and you are still well within the 350 milliseconds. They've got used to it; it's no longer fast enough. So being able to say, well let's be far more deterministic. What is the performance that's required and how do we meet that performance? And if the performance changes, how do we dynamically and automatically meet those new requirements which in the past we haven't been able to do. We've engineered to meet a service level that's been required, now we need to automate to create a service level.

So, Atsushi how do we do that? How do we move to a value level agreement rather than a service level?

Dr. Atsushi Iwata

Okay, that's important questions. Actually in the last week as I said in the previous session the Open Networking Summit cover a very interesting keynote from Nick McKeown in Stanford. He is talking about how we can provide the guaranteed service for the networking side. So far the, let's say, an NSI design for example do we have a lot of [inaudible] to verify the circuit design and the simulation in how we can guarantee the performance for the circuit design label.

And that -- the -- if we think about such a [inaudible] analogy from a circuit design on the world to a network design world, then if we can have a certain kind or same kind of tool in the network world then we can guarantee the SLA an end to end. And therefore, what is missing pieces in a networking world is that kind of what we have is just a pin what [inaudible] and then you have a lot of human, the imagination that where the problem is, where the performance bottleneck is. It's kind of like just a human skill set. And we have to have we have to provide such a good network design tool, ammunition tool, verification tool in advance. Then install the [router] for configuration. Then you don't have any scare, that kind of yes.

Clive Longbottom

So, Bruce if you'll pardon the analogy, well it's not really an analogy it's a statement. Dell used to be a pile 'em high, sell 'em cheap organisation. And it really just concentrated on the computer area. If you wanted storage it was EMC, if you wanted networking it was go out and get it from wherever you want.

There has been a few acquisitions going on within Dell, you know, Compellent and Equallogic and then Force10 and then Ocarina we could go on, there is 30 odd small and large acquisitions, which have made you into a far more rounded player which I think has made it that service level agreements and knitting everything together has

got to be far more of a focus for you. You've bought systems management company Kace. How is it going, and how do you now look at being able to meet this sort of a problem?

Bruce Bateman

Yes, I think Kace and Boomy those are really good examples of how you manage a server or PC desktop level. And that approach -- the nice thing about those products, with those products you can upgrade a Dell laptop server, we manage and we manage. That's the tools now we are moving into the server side, sorry into the network side. So that you can manage things remotely, you don't have to be there when patches are done. So a lot of the tools that are coming out are about management. I think the differentiator between the panel members will be about management tools and about automation and about how I manage all of this.

One of the issues you brought up here about data centres and why the speed 350 may or may not work anymore is because now we have data blocks. It was okay when we only had 10 users on it, and now we've got 50 users on it. And people, whereas the speed at the link is still there but the bandwidth may be consumed 100%.

Mark Pearson

Well I mean, when we are looking at convergence as well, so all of a sudden we brought voice on, we are increasingly having video. Now a lot of organisations we talk to they say oh we don't use video. And we go so do you block YouTube? And so they go, oh we don't count that as being video. [We tell] them you'd better do.

Bruce Bateman

So I think that comes up -- so the tools now have to tell me as the network administrator that I've exceeded what my original plans were. So the design tools that have been tested are only valid for the installation implementation. Those tools need to come back and [ping] tells you nothing. So we need to have true monitoring tools that tell us the true state is the sequel server up? Is it responding properly or did someone hack my site? And it's coming back with a different image.

Clive Longbottom

Okay, so Mark I mean HP has been in this market a lot longer as a full service provider and you had OpenView, you bought Mercury and had a few problems with bringing the two together.

Mark Pearson

That's right.

Clive Longbottom

And now you're talking about ICM, and saying that that's the single pane of glass. So OpenView and Mercury have disappeared, ICM is the new way of managing a total system or just managing the network?

Mark Pearson

Okay, first to correct its IMC.

Clive Longbottom

Sorry, IMC.

Mark Pearson

And second of all I don't know that we've had problems integrating the other management tools. We definitely try to bring these together as a full ecosystem.

Clive Longbottom

I apologise.

Mark Pearson

But what I want to comment on about is these dynamics in the data centre. And the only way you are going to accomplish these dynamics is to have some form of a central control plane that abstracts the all the networking resources. There is no way we are going to do this with device by device management. With the central control plane we can now start to incorporate feedback. We can have an abstract view of the network. We can advertise those views of the network as being templates. Those templates can be consumed by these new forms of dynamic applications.

If we consider another trend in IT which is an idea called DevOps in which we bring together operations and developers into a single model that drives fast deployment of applications. These guys also use templates. They drive off of templates to rapidly deploy the applications. So if we bring this common paradigm of templates in the network, templates in the application deployment we'll get these kinds of dynamics.

Clive Longbottom

He didn't particularly answer the question, so IMC; I'll get it right this time, is this now the way of managing a whole end to end environment? I understand it understands the application contextuality, but --.

Mark Pearson

So I'm seeing it as an aspect of managing the network. The HP portfolio includes several other things that we've now combined under what we call the converged Cloud. And within the converged Cloud we have other tools that together will bring the complete Cloud ecosystem together.

Clive Longbottom

Okay, so when we look at what's been said there, automation is king, which brings us onto the sixth point and abstracting has great value. We love it. We think it's the way forward. But if it's active abstraction then we run into the problems of as we move from 1 gig to 10 gigs to 40 gigs and beyond, trying to make sure that everything still runs at line speed then becomes an issue.

From what you've said about SDN it sounds like the whole idea is it's the policies which are abstracted not the actual action, the action is still going to be done by either ASICs, FPGAs or basic CPUs. So there it's down to pure engineering to make sure that we get the line speed. But can we do everything that way? Or are things going to become an active environment? And when we are looking at the wide area network environment, when we are looking at things like mobile access points and so on and so forth can we do line speed everything, Pranay? Shall I move on?

Pranay Misra

I think we can do it. It's all about how fast can we do. And other thing is, especially in the Indian market where dynamics are very high, where transferring of the life data from the old conventional data centre because it's a [mindset], it's a legacy, so it's a migration, it's a live migration from a point A to point B. So I think it's [take] time consuming.

Clive Longbottom

But Kevin when we look at BT, going slightly out of your area, but BT's backbone is MPLS, which allows you to optimise traffic streams as they are going through. Is it up to what's happening now? Is it up to the future where, not only do we have to say "oh, this is video it's got to be real time" or "this is voice it's got to be real time", or here is an application we can do something different. But it's a case of this is YouTube video so it doesn't have to be real time. It has to all appear at the same time, but if it's all two seconds out it doesn't make any difference. Can MPLS do that?

Kevin Buckingham

Yes, it's a good point. I mean BT is investing heavily in the network, and again from our customers their demanding to do more with less. So we are looking at trying to put active traffic shapers on and agreeing with customers what is their prioritised traffic actually going out to the network.

Because some of our customers don't have the expertise and the experience to be able to keep up with developments such as, you know, what's the next YouTube going to be? What's the next bandwidth grabber going to be?

So they have to have somebody actively monitoring that. And obviously you can get tools to do it, but there is a lag. So BT is looking to provide that as a service but investing heavily as well in the bandwidth.

Unfortunately when you look at the latency within Asia Pacific it is a problem like we mentioned before. Bandwidth is now becoming a problem, as I think we discussed earlier.

Clive Longbottom

Or how that bandwidth is used is the main problem not bandwidth itself.

Kevin Buckingham

It is, but I think, you know, we discussed earlier the explosion of data now going back from Asia to the North, and I think that's become an issue because I think originally the traffic sizing was actually done for the other way round. And now the business is migrated over to here, and we have businesses that actually want to send large amounts of data back to the North or even the West, then that is becoming an issue.

Clive Longbottom

Okay, so I just want to give a chance for one question from the floor. I've gone on a bit too long unfortunately. Has anybody got a burning question that we can just ask to the panel? It looks like a no, so that brings me back pretty much to time, which is quite good.

So you'll have noticed that I've only done six secrets. Here is the seventh one just to actually tie everything together. Networking has become commoditized as far as a lot of organisations are concerned. For the vast majority of business people the network is that piece of cable that goes between their PC and the wall, or to people who don't use that any longer who just use wireless connectivity, networking has become almost transparent. Although as we've all seen today once that network disappears all of a sudden it becomes very, very important again.

But here what we are seeing is as we move towards the new design of data centres all of a sudden networking becomes exceedingly, extremely important again. If we don't get that basic building block correct everything else is like building on sand. And we have to get it right at the data centre level. We have to get it right at the local area network level, the campus and the branch as Mark was talking about earlier.

But we also need to get it correct at the wide area network as well, because more and more of the people who are working with our organisation, not necessarily for but with, the consultants, the contractors and so on, aren't coming in over a highly deterministic environment. So unless we take a complete holistic view of everything the network will hold us back in a competitive environment.

So I would like to thank my panel very much for entering into the spirit of things, and giving some really interesting views. And thank you very much everybody.

Manek Dubash

Thank you Clive, thank you to the panel, I suspect this also is another topic that will run.

[End]