

NETEVENTS

## EMEA IT SPOTLIGHT

**FINAL**

### *Conference Debate Session IV – Faster, More Scalable, and Easier to Manage: Architectural Management Trends and Next Generation Data Centres*

Chair: Ksenia Efimova

**Senior Research Analyst, EMEA Telecoms and Networking, IDC**

Panellists:

Kevin Deierling

Chief Marketing Officer, Mellanox Technologies

Philip Griffiths

Head of EMEA Partnerships, NetFoundry

Joe Baguley

Vice President and Chief Technology Officer for EMEA, VMware

**Ksenia Efimova**

Hi everyone, I know that it's only our discussion that separates you and lunch. So we're going to make it quick but very productive. I'd like to welcome my panellists to enjoy the discussion about more scalable, faster and better managed data centres. I'd like them to introduce themselves very quickly, though you know them, most of them; they were on the panels before. So, Joe, do you want to start?

**Kevin Deierling**

I'm Kevin Deierling. I'm the VP of marketing of Mellanox Technologies.

**Joe Baguley**

Middle one.

**Kevin Deierling**

Okay.

**Ksenia Efimova**

Okay, good. So let's start. Perfect. So with IDC, I'm responsible for enterprise networks and so would be happy to discuss with you this fascinating theme. To start, I would like, by presenting what IDC thinks affects the overall ICT [unclear] market. According to IDC, this is the third platform; the third platform that consists of [feelers], cloud, social, mobility and data; and innovation accelerators, AR, VR, AI, robotics, block chain, et cetera.

But what is underlying all of that is significant demand of data that is being gathered, and that data needs to be processed and acted upon; but also what has to be done and has to be changed, is the management of the data, security of the data and the overall approach. Of course, data centres are the tools to process the data but the data centres need to change as well to adapt to those changes and how they can change? Well, first of all, and the first thing that comes to mind, right, is the increase of speeds of networking to process that amount of data.

We see the acceleration of 100 gigabits switch and of course it's not the limit. Several vendors have already introduced 400 GbE, but it doesn't mean that enterprises would go for that. In the very, very, very rare cases, more - mostly it's going to be driven by the hyperscalers; but also it's not - the increase of speed will not solve all the problems.

So infrastructure; when we talk about infrastructure, the rip and replace will not work. You need to have a different approach. You need to have a different approach to your storage. You need to have a different approach to your servers and the network, and altogether it has to work as one system. Energy consumption; increasing volumes of data require the [compute] systems to work more efficiently and require the use of energy power. Manageability; well, the complexity of the data centres on [PRAM] off [PRAM], cloud-based, traditional data centre, private, public.

It's impossible to manage it manually, so other systems have to be implemented. AI - AI will be used within the data centre environment, not only for networks, as we see right now, but also for storage and Huawei has introduced the chipsets [cheap sets] that are AI-based and they would be using them in their 400 gigabit [switcher]. So on that, I would like to ask the panellists on what would be the first steps of either enterprisers or hyperscalers that would want to implement that type of new data centres.

**Philip Griffiths**

Very open. Well, I think step number one is to define what your strategy is and what you're looking to achieve. Are you looking to build a hyperscale data centre for a [unclear] to be able to live in? Or are you building something just for your local German subsidiary, so you can keep data in-country? Or are you working to - you know, the data centre providers who are looking at Edge locations of doing low latency, IOT processing where there may be a couple of blades or racks within a telco tower or something like that?

Once you define that strategy, you can then start working out what you need to implement to deliver the outcomes that the customer wants. Now obviously, it brings in the general topics that we have up here around reducing energy consumption, using automation and AI so that everything is software-defined without having to send engineers onsite, or to fix things manually every time required.

But really, the key first step is, what am I trying to achieve? Who am I delivering this to? Where am I going to make money? Even if you - some of the data centre providers you look at, like Equinix for example, they've almost repositioned themselves as mostly a connectivity company to different clouds, who are providing places where data's [relocating]. So it's really that first step of the strategy for me.

### **Kevin Deierling**

Yeah, so I would say that, look at things holistically. So we like to talk about, think outside the computer. In the old days you could count on more is law, just sort of doubling performance every couple of years. That game is over. You know, they said that what Andy Grove giveth, Bill taketh away. You could just absorb all the software and just keep writing inefficient applications.

We've seen companies that are really innovated at the networking level and I'm here with VMware; you know, obviously, Martin Casado has spoken here before. He was sort of the founder of network virtualisations. Incredible technologies but can consume all of the CPU cycles running that. What we've seen now, we're actually working - just in the old days we had a Cisco router that forwarded packets using software. Today no network in the world - we ship 100 gig, 400 gig switches; no network in the world can forward packets using software.

So you need a hybrid of ASIC hardware and software and this is where working with folks like VMware and the other OS and hypervisors, we actually accelerate virtual machine forwarding and firewall rules and load balancing and all of those things at hardware. When I first told Martin that - he's a brilliant guy, he's a Stanford PhD; and intellectually he was so excited; he goes, you mean you put a virtual switch into silicon in an ethernet NIC? I said, yeah. He goes, that's fantastic, that's incredible. He says, too bad I can't use it, because I'm VMware and I need to control that interface in order to have the software control.

It was only once we explained that, in fact, the control path is still controlled by the VMware, so really all of the discussion we've been hearing about security, for example, really the policy decisions still need to be at key maintenance management and all of those things. It's the data path that you accelerate. So the intelligence still needs to be at that higher level; and really, what we like to think about is optimising at the data centre level. So you're not optimising at individual [box] levels but really thinking across the entire platforms; and that is compute, storage, networking and application; and think of that holistically. That's one thing; the data centre is the computer and if you can solve things at that level, rather than getting buried down into the weeds, you'll do much better.

**Joe Baguley**

Yeah, so I think, you know Kevin makes some very valid points there. The looking at the whole data centre as one computer is the important thing. So you talked about hyperscalers, you talk about where things are going. Many years ago, I was involved in the sort of foundations of what became the code of conduct for data centres; the EU Code of Conduct in terms of energy efficiency and some recommendations there. A lot of what we thought was the fact that what was going into data centres that people were trying to power and cool was diverse and, in some cases, very badly designed and we were trying to retrofit that.

If you now look at what the industry's looking for, there's a couple of things. So a lot of the vendors are - a lot of the customers we're talking to aren't building new data centres; they're trying to get out of data centres; it's mainly about data centre evacuation. What they're looking to do is to move into someone else's data centre, like in Equinix, or maybe go and put all their stuff in somewhere like - you know, just move their VMs to one of the other providers that just hosts VMs.

Those that are building data centres are looking at what's generally known in the industry as GIFEE, right, which is Google's infrastructure for everyone else. So people look at how Google or the hyperscalers build a data centre and say, okay, how can I do that myself? That's what I want to do, right. If you look at how Google do that, everything's done in software; you have some [rack to centrally] laptops with their own on board power supplies, they have a local storage, and at the top they have a I-don't-give-a-crap layer two switch, and everything else is done in software.

So if you look at what's happening in our industry right now with hyper-converged infrastructure and what we're doing with virtualised networking, storage, compute, et cetera, it's essentially the same thing. So the customers that I'm working with now that are building large scale data centres, they're looking to rack essentially Lego building blocks of hyper-converged infrastructure that they then plug in via a 10 gig to a 10 gig spine, which then goes [10/40], or even 40 to 100, whatever; it doesn't really matter.

What that spine does is just flat layer to, don't really care because all the intelligence is done in software on the devices. The interesting thing when you look at how that evolves going forward is, that actually allows for a much more efficient use of hardware, which means, from an energy perspective, you can more intelligently lay workloads onto that. But it also allows organisations to do a rolling death of soft - of hardware. So if you look at most businesses and go to most data centres, there's a collection of stuff that was bought for a project.

So someone acquired some hardware that's for that project and five years later they go, my [unclear] is too slow, I need to buy new software for - you know, new hardware for my project. When in fact, if you run like a proper CloudScale vendor on a software virtualised platform, what you're doing is, on a weekly, monthly basis, you're buying new hardware, new Lego blocks; you're racking those up. Automatically, your most important workloads are moving onto your new [tin].

Slowly, over time, tin works its way down this process, so by the time that stuff's five years old, it's sat at the Edge running - you know, in your data centre concept, it's running the, I don't care overnight [unclear] crap, or whatever it is; and eventually it falls off the end. But you don't get this spikey refresh thing, so it allows people, if they look at that, to go through this rolling process, which, you know, has massive energy efficiency gains, et cetera, et cetera.

I think we always have to caution ourselves when we start talking about energy; energy and sustainability, so a big thing of mine been for a long time. We have to be aware of Jevons paradox in that the more we make something efficient, the cheaper it eventually becomes to run, therefore the more people look for ways to use it, therefore we use more of it. Cloud, right? No one's knocking data centres down, they're building more and more and more. So, you know, going around randomly claiming that using the cloud is energy efficient is a very, very long, dark road that you don't want to go down.

### **Ksenia Efimova**

Joe, you mentioned the hyper-converged, so I would like to address that question; do you think it's going to be for everyone? Because we see the migration of workloads - because before it was for VDI almost, only, but right now we see the migration of CRM systems there, we see the migration of ERP systems to hyper-converged. So will that trend be continuous all the way to the end where we will have just the hyper-converged, or...

### **Joe Baguley**

Yes. If you can virtualise it, it will run on hyper-converged and you can pretty much virtualise anything. I think the challenge here is actually organisational and people and process. I mean, the early hyper-converged weren't really hyper-converged, it was just converged; it was those full 42 rack - the full 42-year rack Vblocks from VCE, right. I remember going to an [unclear] once to show me their new Vblock. So they'd actually taken the Vblock apart and put the networking pieces in the networking rack, and the storage bits in the - and it was like, okay, you've missed the whole converged bit here; because their AWK can support it.

So actually the barrier to HCI is very rarely anything to do with the apps or the software we're trying to run; it's the people not understanding how to take advantage of HCI. Because, if you're using hyper-converged infrastructure and properly building a data centre as one big machine, there's a whole bunch of people, like I said before, that need to give up their fiefdoms and understand they're playing a bigger game. Networking, compute, storage, all one.

### **Philip Griffiths**

Yeah, I mean, if you can go further, some people might say, oh well, you know, I need some bare metal over here; there are companies now which do bare metal via API, which is based upon virtualisation, hyper-converged underneath that that - so fundamentally, regardless of what the app requires, you can provide it using that virtualisation technology and whatever the [tin] is underneath.

**Joe Baguley**

There's very few cases now that require bare metal to...

[Over speaking]

**Philip Griffiths**

Yeah, yeah, but there are still some.

**Kevin Deierling**

Yeah, yeah. But just one other thing; we also see it now moving into the Edge, which is really interesting. So you know, VMware showed ESXi and running on, or SmartNIC on the Edge as an Edge device...

**Joe Baguley**

That was my team [laughs].

**Kevin Deierling**

Yeah, so it was pretty cool stuff to actually - this is a SmartNIC we have that combines 25 gig, 100 gig networking connectivity with some ARM cores for Edge applications; and now that's running ESXi. So we're starting to see hypervisor running on these tiny little machines that can go out. So I think the whole notion of hyper-convergence, the key - somebody was describing earlier - I think it was you, that you move a VM and the network would break.

**Joe Baguley**

Yeah.

**Kevin Deierling**

So all the guys had done their great work and, you know, people talk about hyper-converged infrastructure as invisible infrastructure. You just plug it in, it's easy to deploy and it just works and you get rid of these little fiefdoms of fibre channel; you have a converged network with compute storage all working together, until you move to VM, and then everything broke.

But now we've done the work. So together with switchers, we're monitoring what's happening in the network and when something moves, we get a notification and we can figure the network to do everything. So we've done that integration with NSX, which is the infrastructure layer that sits on top of all the networking, and so now we've made the network invisible too. So it's really a powerful combination, I think. The main challenge is that behavioural challenge. If you're the fibre channel tree hugger and you want to control your SAN, and somebody says there is no SAN, there is no storage area networking, that's hard. You need to re-train yourself and become a hyper-converged guy.

**Joe Baguley**

Yeah, the virtualisation on ARM thing's an interesting one because a lot of people look to this virtualisation on ARM thinking we were going to go for those massive, multicore

ARM service and data centres, where no one's really using those yet. They are the Edge cases, right. Where we're seeing ARM is at the Edge; I'm using Edge too much here but yeah, so we're seeing ARM at the Edge and, yeah, one of the reasons we're looking at it as well is, I can build an incredibly tightly hyper-converged unit with, say, a couple of Supermicro E200s as a twin pair for a vSAN node. Then I can put a Raspberry Pi on top running ARM as a witness.

Yeah, and it was one of my team, well, one of the team in the [OCTO] that looked at the NIC and when, hang on, that's got ARM on it. If we could run that - wow, and it kind of came together. So people are discovering that these convergent Edge pieces are becoming very, very important to the big mix. So having that consistent hyper-converged platform that you can say, okay, from - well, that's the same operations model for that as it is for the stuff in my data centre, as it is for the stuff in the cloud; that's massive for companies. That common operating model is huge; whereas, opposed to now which is how I operate it when it's in my data centre is different to how I operate it with my cloud and the Edge models, different again. That consistency's a big differentiator.

### **Philip Griffiths**

I fully agree with that and I think it's interesting because the topic earlier on 5G and people - you know, the original concepts around IOT and the need for 5G was that; just everything's going to be connected directly to the cloud. As it's evolved over the last five, 10 years, everyone's started going, well, actually no, that's madness; we're going to have multiple tiers of architecture. So we've got on one end, the hyperscale, massive cloud data centres; and then we've got regional cloud data centres; and then we've got Edge on premise data centres; and then we've got maybe even IOT gateways doing local process; and then we've got southbound devices connect into that, creating the ability to have one ubiquitous layer, so that you can move applications and data from one end to the other.

That is where we need to focus because that's really the challenge. I mean, can you imagine taking an MPLS network all the way there? A lot of our customers we work with, it's not possible. So in fact we're working with Dell and VMware on an IOT opportunity at the moment where we're deploying onto gateways to provide that connectivity back into various levels of compute, because that traditional MPLS network, where you would have regional carriers connecting in all to one central point; that's not possible as we're pushing these data centres processing data absolutely everywhere.

### **Joe Baguley**

Those 5G telco networks that are being built out now, they're being built out on hyper-converged infrastructure; and the 5G network itself is being run in NFV on top of the hyper-converged infrastructure. The conversations I'm having with telcos is how they can put those hyper-converged lumps out in the networks. You saw the picture earlier of the 5G base station that was up on the roof; well, one, you have to have 16 times as many of those as a 4G, but no one mentioned that. But behind that, people are saying, how do I put compute next to that? If I want that low latency gaming platform, what I

want to do is, I want to have hyper-converged infrastructure running right there so I can run VMs, et cetera, in it.

So actually, hyper-converged infrastructure and that hyperscaling mechanism is going closer and closer, and further and further out. You know, we've had this talk for years; at one point the MC claimed that the compute was going to go to the storage, if you know what I mean, and vice versa. But ultimately it's all coming together into one lump and it's that management layer on top of that, the operating model that's probably more important.

### **Ksenia Efimova**

Talking about obstacles, what are those from the point of view of hyperscalers and from the point of view of enterprisers? Because, for example, one of the obstacles that we see very obviously on the market is the lack of skills; the lack of skills that can implement those changes and we see the increasing need for the consulting and Network Lifecycle Services, not only for Western Europe but for the whole world. So what are your takes on that?

### **Joe Baguley**

Kevin said it, [unclear].

### **Kevin Deierling**

Yeah, I think, you know, one of the fundamental things here is that the hyperscalers and the enterprisers are different and both of those people are our customers. When we look at the hyperscalers, they're willing to spend time to save money. So they have PhDs, hundreds of them, and they will orchestrate and automate everything, they will write their own network operating systems that run on top of switchers, they will do massive investments in order to save money; and it goes back to having nice simple hardware.

Enterprisers traditionally have done the opposite. They will spend money to save time, so they will actually pay up front to have somebody else manage them, and I think that's a false trade-off that enterprisers have made. I think what they've bought is complexity. So when you go buy something, you think you're getting a turnkey platform from Cisco and the complexity there is so high that you actually pay for much of managed services to manage all this stuff that you're trying to run. We think it's a false choice; we think that the right approach is to take us [unclear] approach - buy simple hardware.

Now, you may want to go and buy the additional things -so a vSAN, so that you don't have to go right from scratch like the hyperscalers do, write their software to find storage. You can get it turnkey off the shelf now; you can get hyper-converged infrastructure. So I think five years ago it wasn't an option for enterprisers to go off and use the technologies that hyperscalers have. Today it is an option. It's available off shelf; we've seen all that [integration] happen.

**Joe Baguley**

[Unclear]. The key word in what you've just talked about is automation. I think if you look at what the hyperscalers do, hyperscalers look at automation as a fundamental design requirement of any platform that they build. You can't be a hyperscaler if you don't have automation as a fundamental design requirement. Enterprisers have yet to wake up to the fact that automation is a fundamental design requirement, not a bolt on, and that's the problem. I see these people building systems and then working out afterwards how to automate, as opposed to working out how to build an automated system. The only way you get to scale is to build an automated system.

So we do an awful lot of people seeing - you know, as I say, this after management sale thing where we build a system then we look at automation afterwards. No, the key challenge; they're not building automation into things.

**Philip Griffiths**

Yeah, I really fully agree with that. It's the key challenge, particularly for the enterprisers, or the people consuming this, is how to make it as simple as possible for them. Everyone wants this AWS click API experience but there's still so many parts of the IT landscape and the infrastructure stack, which does not comply to that. You know, whether it's deploying firewalls and security, or connecting the network; like, that all needs to be API-driven and some of the stuff we're doing at NetFoundry, so that anyone can create it in minutes and doesn't - you don't have to have a network engineer who knows how to build a VPN with CLI experience.

Instead, the guys that are deploying their applications and their instances, they can create - you know, using, let's say, Terraform, or whatever tool they're using, they can also create those other inputs they need so that you can get to the point where fundamentally every business is becoming an IT business, or a massive function of the business is IT because the whole world is becoming digital. They therefore need to embrace that and look at how they can use tools which are available on the market and then automate around it to pull everything together in a seamless workflow, rather than manually pulling stuff together every single time.

**Joe Baguley**

That comes back to having an inherent policy-driven platform. I mean, if you want to talk about the two differences and get a little bit into the weeds, you know, old school deploy VM, phone up, talk to network team, tell network team UVM has been deployed, ask them to set up the firewall rules, which means they'll then go and configure two or three [unclear], then they'll set up the right [switch ports] and blah de blah, blah, blah and it happens.

New school, create VM. That's it. Literally, create VM and because you've created VM in Pool Z, Pool Z as part of Application Y has a set of policy rules around flow control to do with data and associated traffic that goes with that VM, and it's done. Kill VM, old world. Tell a whole bunch of people you've killed VM, free stuff up, [delete]

firewall rules, rarely happened. New world, kill VM, done. Now you can cross out the word, VM, write in container; it's the same deal, but the point...

### **Philip Griffiths**

So before I was at NetFoundry I was at a systems integrator, so I've seen many Excel sheets, which go through those [x] steps of people's daily business, and it's horrendous seeing organisations still operating in that environment 15 years ago.

### **Joe Baguley**

I spent the first part of my career as the DHCP spreadsheet for HP globally, before DHCP existed. That was my job. I loved it when DHCP came out, but it still took 15 years for companies to actually embrace and use DHCP, and I was the manager of the 15 dot spreadsheet. You know, scary.

So yeah, I mean, what funny - what makes me laugh and this still blows my mind; there is a company, a customer of VMware that still has a manual paper-based approval process for Vmotion. All right, you want to VMotion a server from one place to another, you have to get it approved and signed by however many people, because they still have processes based on the old world applied to new tech. Big, big, big problem.

### **Philip Griffiths**

Yeah, I mean, maybe we're going off topic a little bit but it's an interesting area where, like Devvix is clashing with [ITIL] and SMC. As of yet, I haven't seen a very convincing bringing together of the two so that you can have the agility and also the operational stability at the same time. I think that's still being worked out; but bring in ITIL to new deployments with data centres and cloud, it doesn't work.

### **Joe Baguley**

ITIL's gone, sorry.

### **Philip Griffiths**

Yeah.

### **Joe Baguley**

It's a bit old. It's a good way of slowing things down.

### **Ksenia Efimova**

Good. The last question from me and then we can bring it to the audience. Partnerships, because no man can do that alone, right?

### **Joe Baguley**

Or woman.

**Ksenia Efimova**

Or woman, yes. So what's your take on that? How do you approach the market? Who do you see as your partners when going to customers and saying, we can help you to transform your data centres?

**Joe Baguley**

Are you talking from the perspective of our companies? Or just generally, if we were in a [unclear]?

**Ksenia Efimova**

Generally. Generally.

**Joe Baguley**

That's an interesting [unclear].

**Kevin Deierling**

I'll start off. So I think one of the big challenges we have is we sell to the big hyperscalers, so if you're using the internet today you're running on our gear; it's buried inside. That's very straightforward. You know, you go into these big guys, like a Microsoft, or Facebook, or whoever. Those are very straightforward engagements. Where it gets more complicated is where there's a multi-vendor solution. So VMware is a good partner; we have a ton of activities. I have a lot of people on my team and in R&D resources supporting VMware and the challenge there is we say, well, how do we get to market? How do we get the solution to market? I call it, then a miracle happens; that somehow we get to the market together.

But the reality is it's very complicated. You actually need to create the solution, you need to get the integration to have it happen, you need partners, [unclear], system integrators. So to me that's one of the big challenges, and then also tracking it. I have to justify to my management, to my boss, the CEO, and say, why are you spending all these engineering hours doing integration with [unclear, where's the [revenues]; and I can't count it. I can't count it. It's hard for me to point to it. I know it's there. I know things are happening.

So sometimes it's just business challenges to make that happen and we have lots of partners that we come together because the stack - again, we're optimising across the whole solution stack. So that's the infrastructure layer, the application layer, networking, storage; everything has to just be fully integrated, and that's the biggest challenge we face.

**Philip Griffiths**

I fully agree with that. We work with a lot of different partners creating different solutions and doing things. But the first question is always, okay, well, what deal is leading to that? What opportunity is going to drive that? How are we going to get back that revenue? There's many interesting cool things we can do, but ultimately you have

to square it back to, what's good to the business and how is this going to drive things forward?

I think the challenge in doing that then is really sitting down with a customer because it all depends upon sitting down with a customer and understanding, what is your roadmap for your business? Where can we bring in some capabilities, together with partners, where that solves a problem for you, which, today you don't know how to solve because you don't know the world of the possible; but that we're able to pull that together.

Now, doing that requires people that have unique skillsets that are able to understand both customer side and then different technologies. I think the biggest challenge there is - therefore two things, one, being able to drive it towards the sale side of the business, because all business - well, the majority of IT businesses are built to drive profit, therefore have to be able to show that; but then, two, having the skilled people who are able to bring together those and see the values in partnerships.

I don't know about you guys but I've worked with many people in IT who are great salespeople but you describe a technology to them and they go, okay; and then a week later they go, can you describe that technology to me again? You're thinking, you've been - so you've been selling this for 10 years, how do you not understand it, let alone pulling together other partner technologies to solve a customer challenge? They know how to map an organisation but doing the technical innovation stuff, that's a rarer skill.

### **Joe Baguley**

So if I was an enterpriser - I mean, obviously VMware we've got - well, we partner with everyone. I mean, if I was an enterpriser looking to sort of say, okay, fine, how do I step into the next level of data centre, so hyper-converged, et cetera, et cetera, I'd be looking for [BRETS]. The challenge you have when you meet most organisations is there's a certain narrowmindedness or bias based on where they've come from. They might be experts in networking, so everything will be seen through the network lens, or it will be seen through the storage lens, or it will be seen through the compute lens, or the whichever lens they come from. Or the managed service lens, or - however they think that the world is going to end up and it's actually understanding that they have a broader viewpoint, or at least gathering a few of them around a table to build something bigger and broader.

The difference we made at VMware was that we collected together an incredibly diverse mind set, an incredibly diverse group of people to build, many years ago, this vision around how you should build data centres and how you should build infrastructure going forward. It's only because we had people like Martin Casado and others on the team that would sort of remind us that we had to look at things differently or broader is very important.

I remember in the early days we were talking about VMs and one of the people on the team said, oh no, but you should be really calling them containers; and everyone around the table was like, containers? This is 2011. What are you on about? It was the guy that invented zones on Solaris, so he kind of knew what he was talking about in containers;

and that then made us think, actually, yeah, we're not just talking about a future platform for VMs, it's a future platform for Apps in containers; they just happen to be VMs right now and in the future they're containers and next is unikernels and so on. But it is; it's that BRETSA thinking. That's what holds people back.

**Ksenia Efimova**

Thank you. Audience, I would like to address you, any questions from your side to the panellists?

Well, if there are no questions, I'd like to wrap up and thank you everyone for this discussion, and if you have any other questions either to me or to the panellists, I think we are around to address them. Thank you.

[end]