



Optimising datacentres

The latest *FStech* roundtable, held in association with F5, Fusion-io and Interxion, looked at datacentres and the hugely important issues of virtualisation, cloud computing and mobile device growth

Wil Cunningham: Let's start with optimisation, and discuss that in the context of maximising datacentre efficiency and resources in this new decade, virtualisation and increasing operational efficiency by migrating to the cloud, cost reduction through utilisation of the cloud, and then getting into compliance and environmental considerations.

Richard Norris: I'll start on virtualisation if I may. When I was at Tower Gate, running most of the change projects there and building up the datacentres, we built a Citrix XenDesktop solution with quite a large number of blade servers to deliver the capacity for our 1,000 users. And it soon became apparent that we were burning datacentre space almost faster than we could provide it. With the dynamics of the business, with it growing at such an exponential rate, what we ultimately did was work out the datacentre requirement as an ampage requirement per user as opposed to anything else. So from there we could forward forecast the number of users vs the amount of power we required to run the kit to service them. That leads me to another project...when I started building up the environment at Cullum Capital Ventures (CCV) I wanted to significantly reduce the datacentre exposure to the business and reduce capital costs and

Opex costs, so we started by virtualising all of the server infrastructure in the core which oddly enough meant that we had six physical servers delivering the core environment to approximately 500 users.

WC: What was your ratio?

RN: Probably around 4:1 - we could probably go significantly further than that but that seems like a nice, comfortable place at the moment. Virtualisation for us has really allowed the company to leverage some benefits. What I'm really interested in now is moving locally those virtualised instances out into an infrastructure as a service offering. I think there's some food for thought there as a lot of people I talk to haven't gone as far as we have down the virtualisation route. It gives us great flexibility with our suppliers.

WC: Other High Street banks have a ratio of at least 20: 1. They really pushed it and got to a situation where they only had to stop because of certain legacy and transition constraints.

Mat Young: One of the challenges we see and the position on datacentre efficiency and how you grow it, the CPU is hangstrung by its inability to get the data quickly. We assume they all respond in a milisecond and the processors are getting faster all the time. The reason that the workloads are sitting at relatively low CPU utilisation is because you can't feed them. The whole reason we exist is to basically feed those CPUs. We do that by making sure we're as close to the CPU as possible, that means inside is the server, outside is in the cloud. One of the biggest challenges to solving efficiencies in datacentres is that you can't drive the CPU fast enough using the architectures we have today.

WC: We've discussed at previous roundtables that the vendors always articulate the target environment but not necessarily the transition to the target. And for installations in largely legacy datacentres, unless you can tell us how to get there we're never going to go on that journey.

MY: I totally agree. From Fusion's perspective, the only way of getting from point A to point B, and this is one of the things we

Attendees:

Wil Cunningham, disaster recovery consultant, Lloyds Banking Group (chairman)
 Nic Evans, director, Evans Global Associates
 Nick Foster, Banking Division Head of IT, Close Brothers
 Terry Gigg, IT Infrastructure Programme Manager, Close Brothers
 Rob Handicott, Chair - Financial Services Group, BCS
 Rob Holman, Datacentre Specialist, Consultant
 Richard Norris, IT Director, Cullum Capital Ventures
 Pat O'Brien, Lead Designer IT Infrastructure Programme, Close Brothers
 Nathan Pierce, DC Solution Architect, F5
 Lilia Severina, Major Account Director, Finance and Capital Markets, Interxion
 Robert Threadgold, Head of Platform Infrastructure, Standard Bank
 Mat Young, Senior Director of Marketing, EMEA, Fusion-io

try to do, is by delivering efficiencies. You still have to spend money on consultancy to allow your applications and processes to move from point A to point B.

WC: I'll turn at this point to the guys at Close Brothers if I may - you're in the middle of this right now, is what's being discussed resonating with you?

NF: Yes, I guess it is. We're looking at converged stacks at the moment and clearly that's a big step for us in dragging it all together. I've lost count of the number of times I've been asked, should we be going into the cloud? But the organisation isn't ready for that.

When I was at HSBC, there were a number of challenges around this. At the time, HSBC was one of the biggest virtualised companies in the world so if they could have virtualised it they would have. For starters, all the trading applications...most of the applications we had were legacy and to virtualise them was basically impossible. It's notoriously difficult to do capacity management and to design your infrastructure based on that - it was a very fluid and fluctuating situation. The other issue was consolidation - there was a programme to reduce costs by consolidating as much as possible - the problem there was that at times of maintenance and DR testing that brought other asset class applications down as well, which didn't fit with their maintenance. So there were a lot of internal problems and what we learned was that you do analysis of the business first and the IT process that supports that business and then you chunk them together. In terms of other places I've been to, virtualisation did cause lots of issues with putting them onto blades and increased power cooling.

WC: You're spot on. The technology is just a one dimensional consideration, the target operating model is how you have to address it. We developed a placement strategy/policy which mandated that testing and development should go to a specific 'testing' datacentre or even the cloud, but where should payment applications go? Unless you look at your processes and target operating model, it addresses one part of the problem but moves that problem somewhere else.

Nathan Pierce: Consolidation is one way of looking at it. But you can also get a lot more out of your existing datacentres. I've worked on a number of projects lately. Does everyone remember the SETI@home, Search for Extraterrestrial Intelligence project? It used your idle CPU time but it was run as a low priority process so whenever your application needed CPU it backed off. I'm

working with a bank now with their debt risk calculation team - about 400 blades a year, so you can imagine the cost of it to them, but they have a grid controller that sends a job out to a machine and that machine responds with an answer. So instead of buying 400 servers a year, we're looking at their virtualised estate and we're giving back any idle CPU to that grid team so every machine and every datacentre should be at 100 per cent CPU 24/7. And the savings are enormous. These are the kind of things you can deliver today.

WC: Sounds like you first have to do a lot of technology and process discovery in order for the solution to work?

NP: We've already done a lot of those things and a lot of the deployment we do now we've got it down to just templates. We've done it at so many financial organisations that we've written all of the work.

WC: I'm trying to understand the benefits here. Virtualisation I can see, I've recovered some more datacentre power and space that I can now use for different technologies etc. Do you decommission what you don't need or do you buy another server?

NP: I do a lot of work on dynamic services models. Let's say you've got 10 racks in a row at the end of the day when the service winds down, running at 10 per cent CPU, why have them all turned on? What we can do is start moving virtual machines to one end of the rack, keep those two running at 70/80 per cent, turn the rest off, then you can set thresholds - when it gets to 80 per cent turn the third one on. We can turn a server off that's in production without downtime because the user doesn't connect to the service - they connect to our technology and we do rest for them.



The mobile security threat - how can banks employ technology that will keep customer data safe?



WC: Again, the guys from Close Brothers, would that help with what you're doing at the moment? My problem was to defer a £200 million datacentre spend. I'm not sure what your problems are at the moment?

PB: I think we have a long way to go before we reach the level of sophistication offered by F5.

WC: I guess it's a case of small steps as there are many things that optimise a datacentre and what has been described is just one of them.

NP: To be honest, few deploy the level of sophistication that I just outlined. That's the end goal. I've got two or three very large global organisations that are going for the end goal. Most are deploying the technology now so that next year or the year after they can get there. But if companies keep going with traditional ways of doing things, where everything is rigid and static, they won't be able to back away from that until they get to the next datacentre refit.

RN: You have to engineer it out, don't you?

NP: You have to plan now, yes.

MY: The bit that links all this together is no matter how smart our technology is, it's having the architects and teams that can bring it together for your target operating model to work out what you can do today, with both the medium term and end goal in mind. The risk is that if you stick with conventional wisdom, the opportunity to change is often gone and won't be back for some time. But I'm not playing down the fact that going to a new architecture, one that is slightly more outside of the box, has larger risk associated with it. But as everyone in finance knows, where there is greater risk there is greater potential reward and it's a case of striking the balance for your company.

WC: We brought the vendors in and said, tell us what you've got. We are vendor agnostic, we don't care who takes who's business as long as you can deliver and our platform guys will sign off the solution because when you walk away they have to support it. So you can create a lot of traction in terms of costs and stuff like that just by opening the doors to them.

NP: It's not just a cost perspective because these technologies give you agility as well. You can run things in a leaner way and turn on the servers when they need to be on. The next steps are visibility and orchestration, you can't have one without the other.

So first you build your datacentre, virtualise everything where possible, after that you need visibility and then you go to the final step of orchestration.

Lilia Severina: Nathan, in our datacentres we have quite a few vendors that allow clients to try out their systems and platforms...IaaS and SaaS is becoming more, well perhaps not more popular but a lot of people are talking about it.

NP: It's just been rebranded cloud - it's the same thing they have been doing for years.

LS: Having said that, there are latency measurement tools that vendors allow us to try and various market data stuff. Your software sounds cutting edge to me and you've said that there are cutting edge organisations using it but for firms that are not cutting edge but would be interested in trying it out, how do you educate the market?

NP: I know some of the examples I've given are very big organisations doing incredible things. But we start right down at £1,000 and we start on a virtualisation platform so it's all ends. And as far as cutting edge is concerned, I don't think it should be limited to companies with the most money and biggest datacentres. All datacentres could be deployed in this way.

LS: So you sell boxes and do you have a service type offering?

NP: We work with a lot of hosted service providers who offer that kind of thing, so they offer the pay-as-you-go model but we do the delivery, resilience and performance solution in front of that.

Into the cloud

WC: OK, let's move the conversation on a bit. In your experience, how many traditional datacentres are moving to the cloud? And by traditional, let's say banks and insurance companies.

NP: It's interesting because a lot of them will say they're not but they forget that with services like anti-virus it is cloud, it has been outsourced, it's just not in their datacentre. The issue is around the terminology, that's the big problem. Banks are more reluctant because of course governance is the key issue but technology is being developed to help with that.

WC: Lilia, from your perspective, how many companies use your datacentres for hosting their own stuff - testing and development being put in an expensive datacentre, what's the point of that?

LS: Testing and development, from what I see in banking, security concerns are not as high as other issues. But we do have financial organisations using IaaS. To be honest, it's sometimes quite hard to tell what they're using it for - some use it where there are certain strategies they don't need to run for a long time - process wise it's easier for them to grab it and dump it. If you go to a cloud event, when you start looking in detail and who is using these services for short-term projects, in my experience there aren't that many.

WC: For me, testing and development is a safe option. But none of the cloud users are coming out with a safe path for companies which leads the whole way.

LS: Like any large procurement of such an important asset within your business, you do due diligence. One of the things that came up at a forum I was at last week, you should absolutely do the due diligence to make sure the service provider is robust and the security protocol's in place and make sure you've got an exit clause in there as well on getting your data back, moved, destroyed especially if you're under regulatory compliance regimes where your data has to sit in the UK.

WC: We're drifting into the area of data integrity so let's discuss that in more detail and also bring in business continuity and reducing risk. Your thoughts on that...

Perhaps to link in with what we've been talking about, part of the challenge is being able to profile your applications, your data, users to identify exactly what their requirements are, in terms of security and processing power, optimising the location where you are. How can you keep track of these thousands of apps and users and identify their requirements?

The key to a business continuity programme is the business impact analysis which looks at the business process and what resources are required to support that business process including the IT services and apps that sit underneath there. And that you can map down to the hardware, databases and physical location of the datacentre. So by prioritising those processes and looking at the impact and value of those processes in a disaster situation you can prioritise your mapping - a gold platinum plated, tier one solution, that's where you spend the money.

WC: It's interesting that the BIA doesn't say who should supply that service.

MY: From my perspective, it seems like that's a dying art. Forget all the technology and hardware, understand what the business needs are if things go wrong, I come across less of those than when I started 15 years ago. Is that just me or do other people see the same thing?

There are an increasing number of business continuity and IT service continuity analysts around, judging by members of the BCI and looking at LinkedIn groups - do a search and you'll find alot more than when I first started.

I was on a datacentre floor recently where they went for a failover test and it went screamingly wrong - 4,000 users offline. And they didn't have a clue when you starting looking into it.

WC: The business experience is really important as all the business people want to know is, is it going to work at the end of the day? That brings us directly onto the issue where, when you talk about continuity, it's not just about the loss of something, it's also the sudden realisation of extra processing you might need at short notice and being able to bring the capacity on very easily by buying what you need over the cloud to supplement your processing needs.

NP: You can use standby datacentres to deal with that and continuance. In the case of one organisation I work with, they would do a flip every month and not just a test flip - they would flip back and forth so six months of the year it was in one datacentre and six months in the other. They then evolved onto a 70/30 split so there was never one that was flat out. And then for capacity on demand they had another company doing the 70/30 but the remainder of the lower utilised datacentres was burst capacity.

NF: We're of a size now, we've gained efficiencies, because everything's cheaper due to virtualisation, instead of having to worry about which services am I going to provide HA to, I can actually provide HA to everything. I can have two datacentres, they don't cost me as much in terms of space and power, I can have one here and one over there, my network's becoming quicker and cheaper year-on-year, so instead of having to worry about complexities, we can stick everything in one datacentre and for a small cost infrastructure we can run an either or, or both at the same time.

That was always the big cost disadvantage - you've got idle resources and you're paying for two lots of disk.



From a physical perspective, you've virtualised and everything's becoming cheaper. If you're making more use of your CPUs, you can use them for a much longer period in virtualisation.

MY: We think of datacentres as places where questions get answered. If you think about that from a company perspective, in an ideal world you want the minimum amount of space and resources to answer all the questions you have in your SIs. That breathing space we initially give them they eliminate very quickly by asking even more questions - my point is that we have to make a break from that, which comes from a brave group of people, architects, implementers etc. Understanding the business becomes key.

WC: The logistics in that scenario becomes your problem, unless you've got orchestration.

NP: Orchestration can only happen with abstraction and visibility. As long as you've got all three working together...

Mobile security issues

WC: OK, moving onto the big mobile security issues. The mobile security threat - how can banks employ technology that will keep customer data safe? And what technology still needs to be developed to increase mobile banking take-up?

RN: A bit of encryption between the aerial and the device would be good.

NP: We've developed some technology built into IOS, Android and Windows mobile devices where we've built the encryption communication from that operating system to the app in your datacentre so no-one can see the traffic. We can cover that issue of mobility but the next problem is that technology is not driven by the IT department anymore - it all starts out with a couple of execs being given an iPad and saying, I want my mail on this and because they're execs you do it. I know only one organisation who was brave enough to turn around and say no, we're not doing it. Once you get to that level of iPads and iPhones, people have left the company and gone back and sued the company because all their photos, personal contacts etc were on the device and they were wiped when they left. So what's happening now and I think Android is going to have this first, phones with dual core CPUs and the first mobile phone hypervisor where you can switch from calls from one number then over to calls from your work number.

LS: That's a brilliant idea.

I have trouble answering one phone!

You can't have technology consistently solving what are human process issues. It's people who want it all together, you can't keep blaming technology. If you allow data on a mobile device, then someone can steal it. There isn't any technology, any guaranteed way of stopping someone with a bit of smarts who sets out to circumvent and steal it.

NP: Organisations that want this are writing their own apps that will use part of the storage on the phone and encrypt that, so that not only is it encrypted on the fly but then when it's on the phone it's also encrypted.

What happens if there's a legal issue where you need to look at the data but you're deploying military grade encryption that even your guys in Cheltenham can't crack?

NP: Actually, the best way to get through it is to buy a bunch of Playstations - the GPUs are apparently brilliant for hacking crypto. Going back to the human issue, mobile devices should be a way of carrying data, not viewing data.

LS: Shouldn't there be some separation between what is business critical and what isn't? We live in a democracy and if someone leaves a company and works in sales and marketing, for example, you can't just say everything about you belongs to the company. People get hired because of their contacts.

RN: It's all about the data at the end of the day, isn't it? You can build datacentres or use cloud services, you have to analyse the data and see how critical it is to the business and work out where you feel comfortable placing it.

That's the problem - most people don't have that kind of control over their data. That's really difficult for a small business - it's one of the issues against moving to the cloud.

Crucial data should be contained centrally and not released to mobile devices. That runs the risk of losing it and also you control back ups of the data if it's held centrally - it's crucial to know where things are being kept.

WC: We've come to the end of the allotted 90 minutes, I'm afraid. Thanks to everyone for your contributions and I'm sure we can continue the discussion over dinner.