



RUNAS RADIO



<http://www.runasradio.com>



Richard
Campbell

RunAs Radio is a weekly Internet Audio Talk Show for IT Professionals working with Microsoft products. The full range of IT topics is covered from a Microsoft-centric viewpoint.



Greg
Hughes

Text Transcript of Show #049
(Transcription services provided by [PWOP Productions](#))



Michael Manos on the Efficiency of Microsoft Data Centers!
March 19, 2008



Michael Manos on the Efficiency of Microsoft Data Centers!

March 19, 2008

[Music]

Carl Franklin: From runasradio.com, you're listening to RunAs Radio, the Internet audio talk show for IT professionals with Richard Campbell and Greg Hughes. This is Carl Franklin, introducing show #49, with guest Michael Manos, recorded March 10, 2008. RunAs Radio is brought to you each week by PWOP Productions, offering professional media and podcasting services online at pwop.com.

Richard Campbell: Hi, you're listening to RunAs Radio and I'm your host Richard Campbell. With me as always, my co-host, Greg Hughes.

Greg Hughes: Hey there, Richard. How are you today?

Richard Campbell: Things are good man and I'm plowing away at all the work we need to do around TechEd US. Of course, I'm stuck there for both weeks. I'm going to do the Dev Week for .NET Rocks and the IT week with you for RunAs Radio.

Greg Hughes: Right.

Richard Campbell: We've got a huge stage where we're going to be doing presentations pretty much the whole conference. So, I have an opportunity to put together a whole bunch of panel discussions.

Greg Hughes: Yeah. For me, for my money going to a conference, if you can get and listen to one smart person talking, that's great, but when you can get a slew of smart people all in one place having a conversation about really cool leading edge stuff, that's where you really get the most bang for the buckle a lot of the time. So, I'm pretty excited about that, the opportunity to bring some real thought leaders together and see what they have to say.

Richard Campbell: So, they're calling the area now TechEd Online. Last year, it was called Virtual TechEd, but now it's called TechEd Online and if you've got any ideas for a panel discussion or folks you would like to see up on stage with us to talk about IT-related topics, send us an email info@runasradio.com. All right, Greg. Let's introduce our guest. Michael Manos is a seasoned information systems management executive with over 15 years of industry experience and technical certifications. His strong background in IT Technology Platforms combined with telecommunications and network and service management has resulted in his work being recognized by IT industry publications and leading analyst firms including The Gartner Group and Meta Group. In his current role, Michael is responsible for the world-wide operations and construction efforts of all Internet and enterprise data centers for Microsoft Corporation. In addition to his

responsibility for the ongoing administrative and technical support of servers, network equipment, and data center equipment residing within these facilities, his role includes the design, construction and facility-related technology research as it relates to data center architecture. Welcome, Michael.

Greg Hughes: Hey Michael.

Michael Manos: Well, thank you very much.

Richard Campbell: You've got an impressive resume, sir. You've been a busy guy.

Michael Manos: Hopefully, I've done a good job. It's one thing to be a busy guy and one thing to do a good job I think.

Richard Campbell: I tend to agree and we bumped into each other at TechEd Europe in Barcelona a few months back and it was one of those conversations where I thought, "You know what? We need to make this a show," because working on the data centers from Microsoft, that's serious business. You guys have a lot of gear.

Michael Manos: Yeah, we certainly do. I think about the fact I've been here with the company now for about three-and-a-half years and in that three-and-a-half years our total footprint has grown cumulatively over 100% every year and it's on track to continue to do that for as far as I can see. So, it's definitely a pretty substantial challenge to try to 1) manage that kind of growth and 2) manage that kind of growth really efficiently.

Richard Campbell: Plus, with all the movements that are going on at the data center industry right now, if I remember correctly, Microsoft, you're building a new data center close to a hydroelectric dam.

Michael Manos: Actually, we've completed the first one. We're about to complete our second one.

Richard Campbell: Holy cow!

Michael Manos: So, I didn't think the first one was that that long ago, but we move at a pretty quick pace here at Microsoft. Yeah, we announced -- boy, when was it? It was December of 2005 and our first facility, our first really owned and built data center facility in Quincy, Washington. The Quincy, Washington location has since gone live and we are now wrapping up the second construction efforts in Quincy, Washington, as well as other fairly public announcements in San Antonio, in Dublin, and Chicago. So, we've been quite busy since the last time we've met.



Richard Campbell: No kidding. So, what's the real advantage to being up close to the hydroelectric facility like that?

Michael Manos: Well, there are a couple of things. One of the big challenges the industry is going to face is that of increased scrutiny and regulation around green legislation. How green are you? How efficient are you as a company in using the energy resources that you consume? A large burden of that or a big portion of that burden is going to fall to individual IT departments everywhere to start thinking about how efficient and how green am I actually running my data center? How green am I really running my IT environments? As a result of that's we started thinking sort of long-term in our site selection process and I would be happy to cover that more as well, but in our site selection process, we really started thinking about what should we be thinking about in terms of sustainability or green aspect of building our data centers? One of the things that we've discovered is that power generated from hydroelectric dams has a zero carbon footprint. You could probably argue that the rivers themselves give off some carbon emissions, but by and large, the carbon emissions from a river being dammed up is pretty minimal and so to build a data center, a couple of power poles away from the dam allows you to ensure two things. One that you're a completely carbon neutral facility in terms of your overall consumption and use of power, but it's also sort of a nice green thing to do in terms of drawing your energy from a renewable energy source.

Richard Campbell: Now, doesn't being close to the dam mean you're also consuming power that might have otherwise been wasted transmitting over longer distances?

Michael Manos: That's true. There's a certain level of inefficiency. Every time you try to generate electricity and send out long ways down the poles, there's always a transformational loss at every station so being closer to the generation point definitely has its advantages that you don't have those significant losses upfront.

Richard Campbell: So, did you guys got a deal on your power?

Michael Manos: No, we didn't get a specific deal on the power in Grand County. The local rates are actually very favorable as it is and as they were. So, we didn't have to do much in terms of negotiating power rates or anything like that. The base fact of the matter is a lot of those dams were sort of coming off of depreciation if I could use the business term. They were built 50 years ago and as the depreciation of those dams come off, the cost advantage to using power from them is pretty significant. It's a fact that

the local governments and the local utilities definitely saw us as an advantage to attract companies like ourselves.

Richard Campbell: Now, I got to imagine the land around the dam was pretty cheap. It's not exactly high end commercial property up there.

Michael Manos: That's correct, yeah. You know, you're talking about an area that's mostly rural and as a matter of course land costs aren't necessarily significant compared to the Seattle, Washington, or Redmond area, it's definitely not even in the ballpark.

Richard Campbell: Yeah, no kidding, but then I guess the expensive part has got to be you had to pull all your own bandwidth to your building.

Michael Manos: Well, yes and no. One of the interesting things about our site selection process, I began to talk about this a little bit earlier, we have mapped out the world in about 35 different factors in terms of where we want to site select our data centers and those factors are everything from what's the makeup and generation and the power in that area, how much power capacity is there, what's the cost of power, what's the availability of IT personnel in that area. One of the things that you can imagine is sort of one of the key drivers of our site selection process is availability of fiber and bandwidth in proximity to pairing locations.

Richard Campbell: Right, right.

Michael Manos: One of the interesting things about Quincy, Washington, and Grand County in general was that fact that they invested some significant money and infrastructure to kind of build a fairly substantial fiber optic network there and as well as being sort of a crossover point of sort of cross continental fiber connectivity coming into the Seattle area. So, there was actually quite a bit of fiber optic connectivity that we take advantage of in Quincy, Washington, more so than you would think.

Richard Campbell: So, essentially pipelines that were traveling through the county to connect to other endpoints further on, you were able to tap into those and then take advantage of them.

Michael Manos: That's correct. Yes.

Richard Campbell: I guess the labor market is an interesting challenge when you're talking about largely a rural area.

Michael Manos: It is. One of the nice things about this area is there is a Big Bend Community College. It's a community college in the area that's



driving towards more technical degrees and so by having us come in there, we were able to work very closely with the community college to try to figure out what's the right set of courses and curriculum to make sure that folks coming out of that institution have an opportunity to work in our Quincy facility. I would say today a vast majority of the people that work in our facilities in that area are all based out of Quincy or out of that area.

Greg Hughes: I'm curious what kind of an impact does this have on a typical community from a people standpoint.

Michael Manos: Well, I'm not quite sure from a people standpoint. There are a couple of different ways to kind of interpret people's standpoint. From a job's perspective, the data center itself had upwards of about 400 electricians on the job at its peak so there is a lot of construction activity. The building itself was pretty substantial, measured in the hundreds of millions of dollars worth of construction. That obviously parlayed its way into sort of construction-related jobs. Moving on to sort of the long-term ongoing facility support, there is about 50 permanent jobs that the data center itself creates in that area. More so, when you think about sort of the impacts of the data center in that area, what does it mean to have a data center? Is it good or bad? I saw an interesting report out of the local Grand County paper there that said that the result of the data center going in everyone's property tax bill would be reduced by \$100 or \$200 just by having us there. So, I think the impact to the local community can be significant by putting these types of high technology buildings in place in a particular area.

Greg Hughes: Well, you certainly have to be good at your job to build these massive data centers and to do it multiple times. What are some of the things that you've learned over the last several years that maybe IT pros and media and smaller and a little bit larger businesses might be able to take advantage of?

Michael Manos: Well, one of the challenges that I had when I first came on board here at Microsoft to bring a Moore's Law to the data center. If you think about anybody's data center across the industry Moore's Law has had a dramatic impact on everything from your processors, to your memory, to your motherboards to just every aspect of everything that lives inside the data center but very little advancements in driving towards efficiency gains have actually been targeted at the data center itself and so one of the key charters I have on my organization is that let's bring Moore's Law to the data center. Let's bring Moore's Law to the rooms that all of these gear which has such hyper growth associated with it. Let's make sure that those rooms

can keep up specifically in terms of efficiency because when you think about running data centers as an industry executive or as an industry insider you really want to think about, how do I insure that more work is being done per watt than just watts in support of that work? There is an industry metric called the PUE, Power Utilization Effect or Power Utilization Efficiency. I heard it used both ways but it is a way of measuring the overall utilization efficiency and utilization of your data centers. How much power is going in for critical work? Work either serving bits out to the internet or performing some internal IT function and how much of it is sort of in support of that, the cooling, the electrical distribution systems? You really want to try to optimize your data center to be as efficient as you possibly can. We talked a little bit earlier about the coming green or just legislation and sort of regulatory oversight can be coming to this space. The IT pros today are going to have to start thinking about that efficiency metric because one a data center is sited and once you have a data center in an existing area, your overall sort of carbon emission footprint for example, how much carbon emissions does your data center give off? It is a question I think if you post it to most IT professionals today they wouldn't necessarily have a good answer if they have any answer at all but it is a question that's going to be coming their way. In situation where your data centers are already been in place and they have been on operation for a long time, there is not much you can do to drive sort of your carbon emissions other than driving for efficiency in the data centers. So you're using less power to do the same amount of work. It is going to be a key driving component. I can tell you that in our own facilities design, in our own construction optimizing for efficiency whether that be efficiency and power or efficiency in work flow are key in critical elements on how to design and build our data centers today. If you're to walk into a data center today at Microsoft you would see some interesting things even from a work flow perspective. I would like to say that the days of the hand crafted master artesian silver builder are long dead at Microsoft. The days were the box would come on the loading dock and they would open up the box and they might install some memory or some hard drives or what have you and then lovingly install it into a rack and dress up with network cables and potentially load the operating system on the machine itself through CDs or through some other means that has long since died and moved on. We have moved from a single sort of work unit from a server to a rack of servers to whole rooms of servers at a time to really entire data centers at a time. So you can go to a facility today and you would see the truck backing up. When the doors of the truck roll up you will see that the racks are already pre populated with servers. They are on wheels we roll them into the data center. The data center itself has been designed for moving gear in and out. We roll the racks of equipment unto the data center floor. We



plug in the power. We plug in the network and the machine self image themselves, etc, etc, etc. So the driving towards automation is really been a key for us at Microsoft to really figure out how do we streamline this process and turn sort of IT drive automation in sort of IT practices. The scale is pretty amazing if you think about the fact that we deliver tens of thousands of brand new systems every month. It is a pretty substantial scale challenge for us as a company.

Richard Campbell: So, we hear the Steve Ballmers of the world talking about the virtualization technologies utilizing our servers more efficiently. Has that actually started happening on the data center level inside of Microsoft?

Michael Manos: Oh, you can bet that it has. We are talking about significant amounts have spend for us as a company both in terms of the initial server purchases and then if you think about the ongoing power draw over the life of that server driving efficiency and how you utilize that resource is really key. If you think about the skill challenge, if you really need architecture applications in such a way that they really work very closely with virtualization strategies. It is important to know that virtualization strategies mean different things to different people. Sort of the spec out of the box virtualization stuff is really good at CPU intensive applications but maybe there are some challenges in storage or memory bound applications or network bound applications so you really have to address the virtualization challenge very holistically sort of an entire ecosystem layer but once you can kind of get your hands around that and handles around how to address that with virtualization it is a pretty key driver for us internally in terms of how we are approaching the market space. Both in terms of virtualization and geo diversity being able to run an application in multiple data centers but having them rely on each other. Those are key components to our strategy as a business.

Greg Hughes: You know, virtualization is one way that one team of mine used to try to solve the classic perennial problem of too much heat. We have plenty of power we actually have plenty of floor space but heat seems to be the repeat itself data center issue. Is virtualization one of the ways that you combated on that way? I think a lot of the people that are listening have probably had to deal with that the rooms are just getting too hot, I can't add anything else to it. What other ways can they combat that problem?

Michael Manos: From a cooling perspective?

Greg Hughes: Right.

Michael Manos: From a cooling perspective, virtualization is going to drive some benefits there. I

think as you would mention as I think about one of the benefits of virtualization can drive there if the base draws my machine is 100 watts or 160 watts or let's say it is 200 watts but it fully loaded its 400 watts and I have been able to virtualize six machines or have been able to take six physical machines and collapse that down to that one you obviously have a pretty substantial and reduction and total overall power being utilized for those facilities and as a result the net reduction in heat in those facilities as well. One of the challenges though I just want to bring this up is data center managers today are dealing with an ever increasing density per rack. If I go back three or four years ago you're looking at three kilowatt or four kilowatt of power per rack but I think you're now you're seeing per rack density somewhere only 8 to 12 range with some of them at the 16-32 kilowatt range so that's a function of a couple of things. Sometimes it is a function of space that you have to be able to work in. You know a data center in downtown Manhattan as they are expensive they are going to get more space but just in general the amount of power per rack has been increasing, increasing and that has been causing very substantial challenges on the cooling side in data centers to try to figure out how do I cool that? How do I architect my data centers in such a way that I can? If I have the ability to build them brand new, how do I architect my data center so I have the most flexibility on the cooling side to make sure that I can keep their cool? A lot of strategies that I see today with sort of ever increasing density in the rack is that people are moving further and further out there and the space between rack rows is getting definitely much further out of spaces in a constraint for you but it is a challenge. It is a pretty significant challenge as a whole.

Richard Campbell: I wonder if, you know, there is a double battle going on here. It is existing infrastructure, existing racks and existing cooling and existing power allocations are getting stressed because we are increasing the power in competing density per rack. You're putting more demands on that entire infrastructure but then the other side of this is looking at new gear. I get a sense that we are trying to go to cooler running gear. I think the perfect data center is one that has no active cooling at all. I just don't know we can actually run that cool yet.

Michael Manos: Well, I can tell you that that's spot on the truth. There are two points I will make to that. One is sort of backwards looking and the other is forward looking. If you look at most of the data center inventory across the US and abroad today, it was designed between 80 watts of square foot in about 110 watts of square foot somewhere in that range. If you think about sort of the density challenge that we are faced with today we are easily drawing racks sort of in the 200, 225 sometimes up at 300



watts of square foot range. So that's putting unbelievable pressures in sort of these existing data center inventory and there is a demand for that inventory as well. As you can imagine data center inventory across the globe is pretty sparse and as a result of that the rates for that space is getting higher and higher as well so it is a tremendous amount of pressure that how do we optimize for that. There are a couple of key things sort of looking forward that I think kind of play into this; one, we had announced our data center in Dublin, Ireland. One of the things that that facility takes advantage of is something called airside atomization, where we will not cool that facility with sort of internal systems we will make use of the outside air and use that outside air to cool that facility continually. I think going back for over the last 50 years worth of data we've saw sort of the worst year you're looking at eight hours that you've to actually run your data center on some sort of cooling solutions so it is an interesting statement that you have made saying, "Can we even run it that way?" I think the answer is yes we can run it that way. A lot of the developments that Intel and AMD are driving towards the chip perspective are reducing the overall power per core which is having a pretty significant impact as well as the data center. I have mentioned that we are growing 100% year over year. We have definitely seen impacts of the chip manufacturers reducing the power per core requirements as a result of it. I think as an industry we are moving on the right direction on a lot of these things but there is still a long way to go.

Richard Campbell: My first thought there was... I think this is more of a statement of how cold it is in Ireland.

Michael Manos: In talking with the Prime Minister of Ireland I think it is the first time he has ever heard of the climate in Ireland ever being described as something positive.

Richard Campbell: Yeah, beneficial to the process. I definitely agree with you that we're getting cooler cores these days but I have also seen interesting movement in the two and half inch drive space as enterprise class drives for SANs because they also run cooler and consume less power.

Michael Manos: I think the drive for efficiency and the drive for reducing the amount of power in the data center is pretty key. At the end the day, it's all about efficiency in my mind. If I can run the same amount of storage or if I can run the same amount of compute cycles and a smaller power footprint ultimately my bill will be less. I pay a fairly hefty power bill every month I can tell you that.

Richard Campbell: Yeah, that's a bit of electricity there.

Michael Manos: So, very small gains in efficiency over time had a pretty big impact. It is sort of the law of large numbers applies. It is certainly something worth doing for everyone across the industry whether you're a small IT professional or you're managing cloud compute services that we are building on. It is certainly you want to focus on and you're doing the most efficient job you possibly can.

Richard Campbell: A lot of the folks that I have talked are running existing data centers are running out of electricity especially in the urban centers. They simply have no more power to be bought in that given building and so they have entire empty racks because the rack density has gotten so high. In some ways you've got it easy because you're doubling up every year. You're always building new gear rather than constantly tuning and tweaking your existing infrastructure. What is happening to the old data centers from Microsoft?

Michael Manos: Well we are going through the same processes that any standard IT professional would on their own data centers. We have a stable of data centers that were built prior to this growth period that we are managing to and when you build a data center whether it is now or 15 years ago that building is there and the design requirements that you built into that building are set forward at that point in time and you live with those until that building is depreciated off the books or forever and we have the same challenges. We are trying to optimize and drive efficiency in our data centers with our ever increasing power drop or rack and as a result we are doing all the things that everyone else is doing. We are widening the roads between our rack roads we are redistributing some power elements or some cooling elements in our facilities to make sure that we get the most efficient design kind of giving the changing technology requirements those are things that are all happening right now and the drive for efficiency in our facilities is pretty significant especially in our older facilities. All of our data center managers, the folks who are responsible for managing our buildings and managing our facilities have as one of their commitments, one of their key things that they need to deliver every year. How they are measured is all around the efficiency of the building. Are you using the power efficiently? Are you using stranding any power? The worst sort of cardinal sin that any data center manager can do in our world is that if you're stranding power you're not doing the right thing for the business. If you think about the data center itself about 80 to 85% of the total capital costs on the data center both on construction and sort of on an ongoing perspective are all geared towards the electrical and mechanical bill bout. If you're wasting power you're wasting 85%, you're wasting some percentage of that 85 building to some degree concrete is cheap so



spaces is a little bit to waste than power or being able to cool the data center. So we measure all of our data center managers on their overall efficiency or whether or not they are utilizing every last watt that they possibly can.

Richard Campbell: I love this concept of measuring efficiency and transaction per watt. That's a pretty deep down measurement getting right into how does an application work and how much power did it take to get that given piece of work done. I just haven't seen much in the way of integration from the application level going down to the sort of core power consumption aspects of a machine. I have yet to see a really great feed where computer is telling its operating system, "Here's how much power I'm consuming right now."

Michael Manos: Yeah, it's an interesting -- at Microsoft, we measure everything against power, how we normalize, how we look at it financially these facilities. Everything is normalized against power. With the amount of focus that we place on power in our facilities in terms of how we run them, what the operational pay back is, everything is all around power. We don't look at square footage of data centers. We look at what is the total power or footprint of that data center. When we normalize and compare data centers to learn other it is all over the cost per kilowatt per month, that's the metric we use, the cost per kilowatt per month to make sure how much what is our cost to deliver one kilowatt of capacity per month. That's kind of our base metrics work. So when you start pushing and start measuring everything based off of power eventually the application developers begin to take note. I can tell you that there has been many instances internally where our developers especially on our online and live products are looking at things and they say, "Boy you know this SQL query seems to draw more power than this other SQL query and one might be a little less efficient but it saves us more power." It ultimately saves us more power so we are starting to see sort of the very beginning of our developers starting to think about what the overall power impact of their applications are and we are using sort of internal billing mechanisms to really highlight that power is a significant factor that everyone needs to be aware of.

Greg Hughes: Yeah, if the developers are thinking about it, then clearly it's an important factor for Microsoft.

Michael Manos: Yes.

Richard Campbell: I just suddenly get the sense of a data center operator wanting to stay away from certain apps because he knows how power inefficient they are and it is going to mess with his performance

metrics. That's an interesting pressure to exert on someone.

Michael Manos: It is, it is. It's really a pressure that we are applying on ourselves as a company not so much and at the end of the day I can't go in and say this particular application or that particular application is more or less power hungry. There is a natural within the way we architected our infrastructure for people to drive towards more efficiency. The data center managers are measured by if you have 300 kilowatts of extra space let's figure out a way to fill that 300 kilowatt of extra space and use it as effectively and sufficiently as we possibly can.

Richard Campbell: So it strikes me that it is largely been a set of evolutionary tactics you have used to improve data centers as a whole, is there a jumping point coming? Is there something on the horizon here that's going to totally revolutionize how data centers are built?

Michael Manos: I think this is a revolution in many respects but there are many things that are sort of coming on sort of on the horizon. There is always so much a data center manager can do to drive efficiency. We really need to drive awareness into the manufacturers of the switch gear of the cooling systems of all these things to really start putting more of a focus on efficiency perhaps more than they have done so before. So we are working together with many large manufacturers and sort of large data center or heavy infrastructure, heavy industry types of players to figure out how do we drive this sort of the next step but in general by focusing on driving efficiency you will find all sorts of areas to win low hanging fruit to continue to drive data center design moving forward. I would argue that a lot of the work that has happened to date is kind of happening in interesting little silos. You know, the cloud services companies that are kind of dealing at one scale or driving significant innovation in their space. You've got groups like Uptime Institute and Green Grid and others who really driving towards industry and wide awareness towards efficiency gains but there is a couple of challenging factors when you think about data centers. When you go and buy a data center or when you go and build a data center every data center is a little different. There is an ecosystem there and every ecosystem is different you might have lever cooling mechanisms or you might have another brand or you might have this electrical distribution gear or that electrical distribution gear or you may have a mix of a data center manager or data center professionals can have a mix and match of sort of big heterogeneous environment and how do you optimize that particular environment that something that's going to take time and focus on. Not everyone has the ability to go build green field data centers every



couple of months. It is a challenge as an industry but the challenge is evolutionary and when someone figures out here is the right way to think about this problem in my particular ecosystem it might be leveraged elsewhere as well and you have to take the complexity in these thousands and thousands of different heterogeneous environments and what are sort of the best practices. In fact, Steve Ballmer at CBIId a couple of weeks ago just announced sort of a data center best practices, things to think about for best practices in data centers so why pay for that my team and partners and others here internally put together to sort of driving some of that thought, that thought leadership and the space so that people are going to start thinking, "Boy, I'm I really thinking about this problem space on the right way or at least on a different way than I am today."

Richard Campbell: Michael Manos, you have blown our minds I think. It's been a half-hour. We're already out of time and I have a pile of questions sitting on the tip of my tongue. I'd like to keep going. I think we're just going to have to bring you back.

Michael Manos: All right. Sounds good.

Greg Hughes: Thanks Michael.

Richard Campbell: All right. Thanks so much for your time and we will talk to you next week on RunAs Radio.