

## Q&A with Karim Wassef of GE Energy Management

Our increasingly connected world and the oceans of data we generate has put the datacenter under tremendous pressure to accommodate our demands efficiently and reliably. ECN recently interviewed Dr. Karim Wassef, general manager of the Embedded Products business unit for GE's Critical Power business about the trends and the challenges facing designers of datacenter devices.

**ECN:** How is the datacenter meeting the huge energy demands that mobile devices, cloud storage and the Internet of Things all require?

**KW:** So let's start with definitions. When we say datacenter, there's actually a lot of different types of datacenters and we can generally put them into two groups. There's datacenters that are sort of served and operated by an entity for their own use. For example, a hospital or bank will set up a datacenter to support their own infrastructure demands. Then there's another kind which is sort of the on-demand datacenter. So these are essentially datacenter infrastructure deployments that go out and then they're really rented out to different entities as they need networking or storage capabilities, so there's generally different kinds of datacenters.

We see that the majority of the demand is now being supported by these datacenters that are essentially available on demand. So there are installations that are put in place and then they're rented out to whatever commercial entities want to have access to a certain degree of content and connectivity. The deployments today are pretty traditional in terms of the infrastructure solutions that they can use depending on the size of the installation it's based on; either a UPS power backup system or DC systems with batteries or potentially a 12-V infrastructure which is sort of the lowest power grade.

**ECN:** I've read that companies that currently manage their own servers only use perhaps 10 percent to 20 percent of the available computing cycles available, which otherwise translates into wasted energy and cooling. Is cloud computing the best available option to address the inefficiency or are there other design-oriented solutions?

**KW:** It depends on the level of demand you're looking for. There's very few entities that need the full capability of a datacenter deployment. Really there's two kinds of needs. One of them is power and the other is computational power. So computational power is really talking about how many processors do you have available to serve your requirements and the other is how much electrical power is consumed by the datacenter in order to provide the level of services required.... Cloud computing is a way of using existing infrastructure and deploying it more intelligently both in terms of computational power and in terms of electrical power use. Cloud computing opens up the door not necessarily with new hardware but

simply with a more intelligent software solution to optimize the use of the infrastructure that's deployed; so from that perspective, it enables a new level of capability that the hardware has always had but is now really available to a lot more people.

**ECN:** How are datacenter infrastructure management software (DCIM) platforms changing the data center? What should server and storage systems designers know about DCIM?

**KW:** So let's take a classic issue the telecoms had to deal with for a while, which is backup. Part of the reason why there's a pretty substantial cost associated with the power infrastructure for telecom centers historically is that you need to be up all the time. You can't really suffer a lot of downtime. And to do that, you basically have to maintain the power delivery to the infrastructure that you're putting in place. Now with datacenters, historically you kind of had a similar instance where you simply could afford to have racks in a datacenter go down; and still, it's not desirable. The difference is that today with cloud computing you have the ability to move the content or move the delivery of the content around in such a way that even if you have localized power loss or you have some kind of localized weakness in the network, it doesn't hamper your ability to deliver the content to everybody who's looking for it. So the ability to intelligently move content, especially with just a little bit of warning, has changed the power requirements in the datacenter from being something that needs to be up all the time to needing to be powered intelligently, and you can now decide at what level it's OK to not be up. That's a concept that historically the telecoms couldn't even approach -- the idea of not being able to be up continuously and really having that extremely high uptime was unfathomable. Today with datacenters and cloud computing with intelligent resourcing you now have the ability to have power be substantially more cost effective and efficient because you have options like the ability to transfer that kind of networking connectivity intelligently, and that just didn't exist before.

**ECN:** When designing or scaling up a data center, how is power distribution and control considered?

**KW:** There is no one way that has today sort of been established as a core solution. So right now the industry is in kind of a state of experimentation where everybody's trying different things to see whether or not this or that are sort of better alternatives to move forward. The one that is sort of emerging as a potential conduit that kind of connects the pieces maybe is high voltage DC power infrastructure. So 380 V DC as sort of a step up when you go into really high consumption datacenters. Remember at the very beginning I said there's UPS systems, there's 48-V like traditional telecom DC systems, there's 12-V architecture systems. Generally it's a question of power consumption. If you're at the very high end of power consumption, you really do much better when you move up in terms of voltage for your distribution. So historically, 380 VDC was restrictive because the infrastructure solutions didn't really exist, but today we're starting to see a lot more optimism in terms of potentially using that as the future building block for really high power DC infrastructure.

**ECN:** How much of the use of DC infrastructure can be attributed to renewable energy which outputs power as DC, in the data center?

**KW:** It's regional. What you find is that in places where the grid is highly unreliable and where there is some penalty associated with using – whether it's the grid even diesel generators depending on what's available – you'll find that regionally you can have a tremendous demand for DC power that comes from renewable energy. I don't know if it's necessarily sort of the core piece or at least not the core piece that I have visibility to. I think it goes back to DC power potentially being the most efficient in terms of being able to distribute to very high density, very high power consumption solutions. The renewable portion depends more on the alternatives, so if you have a pretty stable grid and you have a low cost to using that grid, for example I'm in Texas and in Texas the grid is pretty reliable and electricity is quite inexpensive. But you go to other parts of the world – you go to India and the grid is extremely unreliable and there's few options. Diesel generators come with their own problems and so their renewables start to take more potential. You go to Europe and again there you have a lot more renewable energy consumption with a pretty strong green factor associated with it. So I think it's not a global activity; I think it's regional.

**KW:** For the datacenter, I can't speak to that level of detail.

**ECN:** One trend that comes up often is virtualized architectures for managing loads across multiple datacenters in different locations. What can designers do to maximize energy efficiency with the virtualized architecture in mind?

**KW:** When you're trying to make intelligent decisions as far as how to reallocate and optimize your power consumption, you need to link two elements together. One of them is the power consumption at the most granular level that you can, and the computational power you're able to achieve at that level. So think of them as lots and lots of building blocks. What you want is to have those blocks able to communicate to you so that you can intelligently identify how much power is being consumed by each of the blocks and how much computational capability that they can each perform. When you have that level of granularity and accuracy, you're now able to design systems that are more intelligent. You have the feedback from those building blocks and you can better optimize the systems to really maximize the throughput – the sum that all the computational blocks can deliver. The part that has historically been missing is that level of granular information feedback. Look, for example, at an industry standard like PMBus. PMBus is a communications protocol that allows power solutions to speak back to a central controller. And the benefit from that is you can now listen and dictate what you want driven on what board and rack in a system. And that level of granular reach through – you can imagine a mesh of communications and control – it really allows you that ability to optimize the system down to the particular boards that are doing the processing. And that opens the doors up for systems architects who are looking to optimize the system. Without that feedback, you're essentially kind of shooting in the dark.

**ECN:** What are some of the latest trends for backup power to handle utility power disruptions. What are some of the monitoring and control trends? Where does the

UPS system fit in?

**KW:** UPS is tried and true. You have a need for datacenters that have at least the legacy of being able to be up regardless of the grid's quality and the potential downtime associated with grid quality. So UPS is the established solution. I think it is for now sort of the defacto default for how you address this unknown potential associated with losing power. There are other pieces to the pie and this really depends on the size of the installation and the momentous time that the disruption can take. So if you need to be able to be up for hours or potentially days, that's significantly different than if you need to be out for seconds or microseconds. So part of what's happening is the virtualization of the network is enabling smarter technologies to be able to operate with longer actual disruptions in terms of power but emulating significantly shorter disruptions. So you can think of it as sort of time-shift where the disruption may actually last in real life hours, but the actual disruption that is experienced by the user is on the order of seconds because of the intelligent bridge - because you now have the ability to shift resources around and reroute your connectivity in such a way that you can bypass where you may have some latency. So this basically says that UPS is one solution but you also have the ability to go into DC batteries which give you the ability to stay up for a shorter period of time than you would have with, for example, the UPS but it gives you the ability to move the networking and the connectivity so you can emulate essentially an uninterrupted service level. Today we're also moving to the potential to use battery-free or battery-less options where you could use high-voltage solutions like supercaps that could enable a few seconds' worth of uptime. And today with intelligence network and virtualization you can essentially reroute your networking and your connectivity so that, again, it looks seamless to the users even though you may have a real power disruption. So the technology is giving users a number of different options. There's the highly reliable -- you know that they're going to have power available for hours at a time -- right down to the option of well, you may go down for hours but you'll only see it flicker because of the way the virtualization of the network allows you to emulate that connectivity.

**ECN:** What tips or techniques can you give for server/storage/networking device designers in terms of efficient use of power? Are there any misconceptions among designers in the telecom/datacom space? Are there any things designers need to know?

**KW:** There is a tradeoff between power continuity and power efficiency. If you're looking for a system that is so robust and so uptime reliable - you can be reliable but not necessarily uptime reliable; which means that you're not really designed to be able support a complete power failure for hours at a time. You're still a very reliable system but you're not really designed for that level of continuous power support in case of a failure on the grid or some other power supply. There's a tradeoff between the two and if you want a system that's always on, then the penalty is going to be efficiency. Because you're essentially trading off an energy storage mechanism for an energy throughput mechanism and the same goes in the other direction. If you're looking for the most efficient possible use of the energy that's available then you want to minimize the amount of backup that you require. So that tradeoff is something that's occasionally lost, where the idea is that while

## Q&A with Karim Wassef of GE Energy Management

Published on Electronic Component News (<http://www.ecnmag.com>)

---

you want to have the most uptime and the most efficient solution and those two are really pulling in two diametrically opposite directions. And the direction to achieve sort of the “greenest” solution is to find intelligent ways to not need 100 percent uptime or 99.999 percent uptime.

**ECN:** What additional telecom/datacom/networking trends do you expect to emerge in the near future?

**KW:** The three that I think will play a more dominant role: The first is digital power communication – that network of communication associated between the power supplies and the central controllers that are driving the computational and the power consumption across the datacenter. I think that intelligent digital power is one vector. The second is probably high voltage DC solutions because that has the potential for increased efficiency that today is probably unmatched in terms of existing infrastructure solutions. And the third is virtualization -- not intelligence in terms of power consumption, but virtualization minimizes the need for backup, and I think that element opens the door for being able to really improve the efficiency by not demanding as much uptime or backup in the systems.

**Source URL (retrieved on 04/19/2015 - 5:40am):**

<http://www.ecnmag.com/articles/2013/10/q-karim-wassef-ge-energy-management>