



Announcer You're listening to the Sun Microsystems Podcast Network. Welcome to another edition of Innovating@Sun with your host, Hal Stern. Today's topic: Sun Eco Innovation. And now, here's Hal Stern.

Hal:

Hello, and welcome to another episode of Innovating@Sun. I'm your host, Hal Stern, Vice President of Global Systems Engineering. And today I'm going to be talking about Eco Responsibility with Dave Douglas who is Sun's Vice President of Eco Responsibility. So, Dave, welcome to the show.

Dave:

Hi, Hal. Glad to be here.

Hal:

So, my first take on this is: Eco Responsibility sounds very trendy. It's sort of today's politically correct theme. What kinds of things are you working on?

Dave:

Well, there's really a wide range of things. A lot of it is around energy, which we'll get back to in a second, but e-waste is also a big topic. We all know how fast we go through electronics these days. 150 million cellphones are thrown away every year now. So, what happens to all these servers, storage, PCs, cellphones. That's a big topic. We're also working hard on Sun's own internal operations about how much water do we use, all the materials, paper and all of that stuff, but really the big issue for our customers these days is energy. And what's interesting about energy is it's hitting our customers in so many different ways at one time. The first is that they're running out of space, power, and cooling in their data centers. You wouldn't think that space would be on that list, but the way most people deal with space is they put in denser servers until they run out of power and cooling eventually. So, it all comes back to power and cooling. The second thing is in many metropolitan areas, people are having trouble getting additional power to data centers. So, if somebody needs to buy new equipment, they need more power, they can't get it. New York City is an example of that today. Also London is having issues as well. And that problem is going to spread farther and farther. There's growing economic costs of energy. Energy prices are rising. But more importantly, the amount of money you pay when you buy a piece of computing equipment has been going up. So, if you think back five years, if you bought a PC for 500 bucks, they used 40 watts. If you buy one today, it's still 500 bucks, but now it uses 140 watts. And that's happening throughout computing. Each dollar you spend on equipment is causing you to spend more and more money on energy over time. And the final area is the one that's near and dear to me is really the environment. And all of this power that's being used for computing has to come from somewhere. Most of it's coming from coal burning power plants, and that's generating environmental issues. So, all these things are kind of coming to the forefront at one time for our customers. So, energy efficiency is hot on everybody's mind today.

Hal:

So, if I can channel the Tower of Power, direct out of Oakland for a minute, they had a song about there only being so much oil in the ground that from a cultural perspective drove awareness of energy and energy crisis in the 70s. From a technology perspective, what you just described I think are outstanding business climate issues. From a technology perspective, what's going on that's disruptive about Eco Responsibility?

Dave:

Well, we need computing to get where we want to go. It's hard to think of a sustainable world without it having more computing than today to help optimize everything, but we're so reliant on it, and it's pretty – it's increasingly clear that the way we build computers today and storage and networking and data centers themselves isn't sustainable itself. And so, that's really where the opportunity to go innovative is. It's to make more energy efficient products, make computers that realize maybe they're not fully loaded. So, they power part of themselves down and don't operate at full speed, use software like Solaris that is able to make better use of the hardware it's given and tape and on and on and on. I think there's just tons of opportunity opening up for innovation around energy efficiency.

Hal:

Well, you used a good phrase there. You talked about thinking about sustainable computing, because I think that much of the attention has been paid to the current and short term operations. What's it cost to power and cool the data center? And we haven't really thought about: What's our data look like in five years or ten years or fifty years? I was talking to a customer a couple of months ago who said that they try very hard not to write contracts using proprietary word processing software, because they're afraid that 50 years from now when someone wants to open the contract and hold them accountable for something that they'll won't be able to read it. [laughter] How much do you see the sustainable computing aspects being part of the Eco Responsibility?

Dave:

I think it's critical. And like you say, most people are really thinking about the here and now. What's this month's energy bill like, maybe? But what's coming more to the forefront is things like CEOs standing up and setting five and ten year goals for greenhouse gas reductions for their organizations. You see organizations now grappling with "Okay, I need twice as much data center space as I have today. Maybe I should stop putting it in my office building like I traditionally have, because I can't get any more power here any more. Where should I put it?" So, that's forcing this long term view. And then like you say, the whole stability of my data, of the infrastructure around that data is getting people's mindset out longer than I think it has been for quite a while. The here and now issues, yeah, there's lots of ways to tackle them, but the 10 year issues are actually bigger and much hairier right now.

Hal:

So, how much of us driving transparency of compute – of really thinking about decoupling the data center as a place you go to versus the data center as something that you consume over the network is going to change our perception there?

Dave:

Yeah, and it's funny you mentioned the data center as the place to go to. That's been a model we've had probably since the 60s, right, where the data center had to be close, because you needed to walk down the hall with your stack of punch cards, and then they handed you the printout at the end. Those of us who remember that from way back, I mean, we've had those data centers in the building with us for a long time, and haven't really challenged that model. And I think as you said, the network is really getting to the point in the way we can deliver services that in a good fraction of the situations in the world, we need to ask ourselves, "Do we really need that thing in the same building or is it better off being in North Dakota next to the wind farms or somewhere?"

Hal:

And there's a sense of being able to see it and touch it and point to it, almost a sense of ownership and pride that goes into that. But there's also, I think a practical sense of how we evolve the way that we manage and build these data centers that will make this possible.

Dave:

<http://blogs.sun.com/innovation>

I think that's right. If you think of the power grid, it's really run by some very large power utilities and buildings that sit in the middle and drive most of the power. A lot of that happens, just because it's more efficient to build a 500 megawatt generator than five 100 megawatt generators or 500 one megawatt generators. And the question's going to be, "Do we reach the same thing with computing that if I really get serious and think about efficiency of building large complexes, can I build much more efficient large data centers?" And does that become a skill that certain organizations have, and maybe we're even starting to see it in the Googles of the world right now, that they're getting just plain better than everybody else at building large scale computing. In their case, it's a very specific application, but does that start to generalize?

Hal:

And if you look at Microsoft, for example, Microsoft Live, Debra Chaprati who has been brought in to go run that organization is spending a lot of her time figuring out where they're going to go plop down that new building so that it can be scalable to be what they hope is going to be one of the world's largest network services. This is not just an enterprise computing thing. This is affecting the consumer market as well.

Dave:

And really as an IT guy, eventually it's going to come down to asking: Here's this outside organization that's specialized in running, say, Salesforce.com and doing CRM so much more efficiently. Is it worth my while to run that internally? Or have they just taken such a quantum leap from me doing it as a side thing that I can't afford to run it myself any more? I really need to go use it from these large, effective utilities.

Hal:

I want to go back for a moment to some things that you said about power and efficiency here. I guess one of the issues we're looking at is density in the data center. There's a space density of how tightly we can pack our computing, and coupled with that, the power density and the cooling density of what we're going to be able to just make the physical space in the machine room a bit more efficient. One of my customers was telling me a story of how he had gone through a virtualization project, consolidated a number of older servers down to a much smaller -- an order of magnitude smaller number of physical machines. But when he got done with the virtualization and the software licenses, it was costing him about the same to operate from a people perspective and a software perspective. When he looked at his power bills, the new servers were consuming more in aggregate than the old ones were. So, he saved on space, but that was about it.

Dave:

In general, today, if you refresh systems that are more than a couple of years old, you should be able to drop your energy bill quite a bit. But it is getting more complicated. Again, if you're taking the historical perspective, we started building data centers just in regular rooms. We pretty soon figured out we couldn't cool stuff as it got denser. So, we put in raised floors and pressurized cold air underneath them and blew that up, and we've been doing that model for quite a while. We're actually now getting to the point where that model's about to break down, and we're seeing power density as well over 10 kilowatts are actually getting over 20 in some cases now. And that means you can't just blow air around any more and get the stuff cool. So, there is a new kind of expertise in design methodology that's going to start showing up. Black box is one instance of that, but you're going to see ideas like that in the data center itself as well.

Hal:

And to take that even a bit further, I would say one of the age old debates in computer science circles is always about time and space trade-offs, that if you have more space, you can take less time because you do more things in memory. If space is the critical thing, then you solve the problem by extending time. And now I think we have to be looking at time, space, and power. And one of the big players there, believe it or not, I think is tape, because it has all these wonderful power properties like it doesn't use any power when you're not actually reading from it.

Dave:

Yes, in the U.S. now, the law says that your medical data has to hang around for like seven years after you pass away, and I kind of often muse about that. If all that's on disk, then your carbon footprint's going to continue on long after you actually aren't breathing any more, because your medical data is going to be spinning around, consuming power.

Hal:

And furthermore, if you're not waiting for it in real time, because you're not here, a couple of seconds to get it on tape, is it really a big deal? [laughs]

Dave:

Exactly. So, medical data is just a great example. When I checked in the hospital, it's great. It takes a little while to read it off tape, get it on disk, access it rapidly and repeatedly while I'm in the hospital. When I check out, after I've been out for, say, a week, and it's clear that I'm not going to be back, then go ahead and move it all back on the tape and let it sit there and not spin around at 5400 rpms. I think tape may really be a wild card in all of this, and it's just got some great properties from that point of view.

Hal:

I think this is one of those things where it's been a forgotten media for a long time. Tape itself is still tape. It comes in cartridges, and it spins around. But the mechanics of finding individual bytes on tape in terms of what the interface is or how it looks like a file system, there's just been dramatic improvements. It starts to feel like a slow file system. It depends on what you're doing online. It may not feel a whole lot different than accessing a large process from over the net these days. So, I think in terms of architectural change, I hope we're on the cusp of seeing people really rethink what they need available to them in the millisecond response time and what they need in the multiple second response time. It's a couple of orders of magnitude, but it's not the traditional 20 minutes to go find the tape mounted up and then go find the bytes, and then convert them into some usable format that we were used to, I would say, with the tape drives to go along with your punch card decks.

Dave:

It's true. And now especially with all the things we've seen about data being stolen and stuff and looking at where real time encryption of tape has gone in addition really I think opens up some new opportunities as well. So, I think you're right on the money on tape.

Hal:

And to kind of close the loop on the governmental side of that, the more that we encrypt on tape, obviously the more we're going to be looking at policy around encryption and key escrow, and just making sure that data's recoverable in any number of scenarios in the future.

Dave:

Yeah.

Hal:

Great. Anything else you want to say? Anything you want people to go think about as your parting shot here as we really begin to ramp up our efforts on eco responsibility?

Dave:

Yeah, to close with my common joke that the eco in my title is as much around economics as it is about the ecology. And I think there's big business opportunities in here either as a provider of things. So, if you're

somebody who sells widgets or sells a service, if you can figure out how to do that more efficiently than the next guy, that's going to be a sustainable advantage you can keep for a long time. And as a user of all of this stuff, if there's a real savings to be had around energy efficiency, I think the last time anybody took it seriously is around when Jimmy Carter was sitting there in a sweater next to the fireplace talking about how we should use less oil. Since then, it's really kind of had a bad name. But I think it's time to think about energy efficiency again.

Hal:

And again, what we think about there will drive how we build data centers, how we look at networking, where we go and put our data, what kind of media formats we use. These are certainly the top conversations I'm hearing from customers right now. And at some level, it will all boil down to the physical aspects of just building a scalable computing environment. That's really what we talk about, David, is doing that with a sense of something that's going to last longer than three years before we panic about it again.

Dave:

That's right.

Hal:

Well, Dave, thanks for joining us today. Again, you've been listening to another episode of *Innovating@Sun*, talking with Dave Douglas about Eco Responsibility, and I'm your host, Hal Stern.

Announcer You've been listening to *Innovating@Sun*. Join us next time for the latest in innovation from Sun Microsystems only on the Sun Microsystems Podcast network.