

Transcription details:

Date: 16-Nov-2014
Input sound file: coralpodcast

Transcription results:

S1 00:00 Hi. I'm Rich with insideHPC, and we're here at SC14 in New Orleans. And this is the Radio Free HPC team: Dan Olds, Henry Newman. I realize you guys normally don't see us, but welcome to our radio show and we are going to talk about some very pertinent things to supercomputing today. First of all, how you guys doing?

S2 00:20 Well, one thing we want to talk about is how we plan to do this show. We talked about it last year. We said, "Okay, let's get a date and arrange it," and all that. One thing that wasn't taken into account - I think we'll all agree - is that just because a day is a Monday in one year, the next year it's not automatically a Monday.

S1 00:41 So, November 17th doesn't fall on the same day of the week year to year, does it?

S2 00:45 No. We're a Tuesday.

S1 00:46 I think that's where we screwed up badly.

S2 00:48 Yeah. So, actually the show hasn't started and as as you'll see in the video--

S1 00:52 Yeah, there's four [?] here.

S2 00:54 It's beeping and all of that. And so as a result, fewer guests for us to have and interview. Less show.

S1 01:03 There's nobody here but the working folks.

S2 01:06 But that said, we're here.

S1 01:08 We're here and we do have stuff to talk about.

S3 01:11 We have a lot to talk about.

S1 01:12 There is some big news this week and I want to just get to the nitty gritty. Coral is three National Lab deals, two of them got announced this week for 2017. These are 150 Petaflop scale machines - big-ass machines, okay?

S2 01:29 Well, 150 to maybe higher.

S3 01:31 Maybe higher, yes.

S1 01:33 So let's talk about the win. Both wins went to IBM/Nvidia.

S2 01:38 IBM POWER and Nvidia.

S1 01:40 IBM POWER with accelerators from Nvidia. Now did they announce dollar numbers on this? I can't remember. It's in the hundreds of millions, right?

S2 01:50 Yeah, it's in the press release, but I don't remember the number. I think it is in the hundreds of millions.

S1 01:54 Yes. Anyway, these machines won't be delivered until 2017. In fact, the Nvidia technology called Volta is two generations out.

S2 02:02 Yes. Just past Pascal. And IBM's talking POWER9.

S1 02:06 POWER9, which is next gen?

S2 02:09 Next gen.

S3 02:09 Next gen.

S2 02:10 They're POWER8 now, or POWER8 plus.

S1 02:10 They're shipping POWER8 now, but this will be POWER9 based. These nodes in this thing are incredibly spec'ed up. There aren't that many nodes for a machine of this size. I want to say 3,400, something like that. These are big, fat nodes, fat with a capital F.

S2 02:26 Yes.

S1 02:27 Right?

S2 02:27 Right.

S1 02:28 With accelerators.

S2 02:30 And by the way, those accelerators are going to be hard-wired into the CPUs using Nvidia.

S1 02:35 Yeah. The NV link is Nvidia's future technology that bypasses the PCIe bus. It goes direct memory access, et cetera, right?

S2 02:45 Yeah.

S1 02:46 But I do want to mention that Mellanox is part of this too, because system interconnect is some future generation of Mellanox.

S2 02:52 It's their whole rebel alliance deal.

S1 02:55 So, is this open power, or is that just a tag line? Is Open Power the same as POWER9, Dan?

S2 03:02 Yes. Open Power is the Open Power Foundation, is what it refers to. [crosstalk]

S1 03:09 Mellanox and Nvidia are a part of that little family.

S2 03:13 No. There's probably 20 or 30 other companies. And there's going to be folks making other systems including Inspur, and Tyan's got a board for power.

S1 03:21 Anybody that joins it can architect systems based on this. Just like ARM, kind of. Except it's not open source. But it's whatever, right?

S3 03:32 Guys, guys! Do you realize the amount of new technology that's going to happen here in three years potentially?

S1 03:37 None of this-- this is all white-board [inaudible] from what I can tell.

S3 03:41 Rich, I would like just to remind the audience, and the two of you, remember NCSA. The last time IBM signed up for something--

S1 03:52 What happened then?

S3 03:53 They did not deliver.

S2 03:55 They backed out of the deal.

S1 03:55 They pulled out of the deal, and Cray came in and saved the day, and built Blue Waters.

S3 04:00 Correct.

S1 04:01 And that, from everything I understand, is a functional machine with incredible [?] in it's delivering real science, okay?

S3 04:11 It is delivering real science, but IBM did not deliver the machine.

S2 04:16 That one? Absolutely.

S1 04:17 They shipped something on the way to that, didn't they? And it just-- well, we're talking about what's the public record here. We'll talk what? Four years ago?

S3 04:29 I don't remember how many years ago.

S2 04:29 No, it's like three.

S3 04:31 All I'll say - and all of it is, I think, public - is the idea of a machine that was supposed to be delivered to NCSA, the machine never got delivered, and it never got installed. I have concerns about the commitment. There was a lack of commitment. I mean--

S1 04:51 It looked from the outside that IBM realized they were going to lose big money on the deal. And they're in business to make money, and they pulled out and took the penalty and left town. Does that sound accurate? That characterization?

S3 05:04 I don't remember the details.

S1 05:05 That's the way I remember it.

S3 05:06 Let's assume that's true. How is that good if they do the same thing and misjudge this time?

S1 05:13 Well, this is two big a deals.

S3 05:14 They have two big deals.

S1 05:15 This is Oakridge. These are serious people. The Coral thing was a big collaboration. It wasn't like some backroom thing that got to announced. This had rules, dude.

S3 05:26 The NCSA same machine was the DARPA HPCS system.

S2 05:31 The thing is, the--

S1 05:32 You're right. So it wasn't a fleeting affair in the night either. I get it.

S3 05:40 It wasn't a fleeting moment. There was a bunch of money put on the table by the government to build that system and make a commercially viable system for HPCS - high productivity computer system. I'm concerned--

S1 05:56 You are saying you're concerned because IBM is coming back and say, "I'm a better man now." Right? They're Jack Nicholson.

S2 06:00 Wait a minute. You're pointing to one deal. We're not talking anything about the deals they did deliver, which I think was [crosstalk].

S3 06:07 They have delivered on a lot of big deals.

S1 06:10 You're going to call out that sour note? I get it and I think you have a valid point, but that is one of many.

S2 06:17 But I would also point to the fact that if-- I don't know if you have any doubts about

POWER9 being delivered?

S3 06:24 I do.

S1 06:24 They bet the farm on that. They jettisoned X86, they have no choice. They can either deliver POWER9 or [inaudible] fade into the sunset. [crosstalk]

S3 06:36 You're 100% correct, but they don't have the integrated ecosystem to control their destiny anymore without the fab.

S1 06:47 They don't have the fab, either.

S3 06:48 That's what I mean.

S1 06:49 What's that company they went to, what's it called?

S2 06:51 GlobalFoundries.

S3 06:51 GlobalFoundries.

S1 06:53 They didn't even sell it, they gave it to them like a IOU.

S3 06:57 Right, so they don't have the integrated ecosystem to say, "Look guys, we need to change our production line and this has to go first."

S2 07:06 But they are paying GlobalFoundries. I want to say it's--

S1 07:12 Shit, man, are you okay? Did you drink a lot of caffeine this morning?

S2 07:14 Yeah, I've been down on the caffeine pretty hard.

S1 07:16 I'm a little worried. You're not going to have a stroke this week, are you?

S2 07:19 No, I'm good. I'm good.

S1 07:19 All right.

S2 07:20 Back to that--

S3 07:20 It isn't-- wait. By the way, and listeners, guess whose birthday it is today? It is Dan's birthday today.

S1 07:28 It is, that's right, at midnight.

S3 07:29 54.

S1 07:30 We went out for dinner last night. I forgot that.

S3 07:32 His 54th birthday?

S1 07:32 Five-four?

S3 07:33 Yeah.

S1 07:35 Well, as the junior member of-- I am only 52, so you old guys are really an inspiration to me.

S2 07:43 But you seem like 17 in a lot of ways, in too many ways.

S1 07:48 I've got good genes.

S2 07:49 Anyway, there's a big commitment that IBM has given to global foundries in terms of cash for delivering chips. Now, you have to have some doubts. If they don't own it anymore, of course there's more doubts. But I think GlobalFoundries can do it. And I know IBM absolutely has to have those chips and has to have them on time.

S3 08:10 It's a lot of fingers to be pointing around. You've got lots of fingers to get this is--

S1 08:15 Well it's two, or three.

S3 08:16 No, it's IBM GlobalFoundries, Nvidia too who's producing their chips, it's Mellanox who's producing their chips, getting this all together in three years, guys - challenging.

S1 08:31 I'm not saying it's a cake-walk, but Nvidia has a track-record of delivering on time, and so does Mellanox.

S2 08:38 And the performance of this thing should be astounding if it all comes through.

S3 08:42 If it all comes through.

S1 08:43 If it all comes due.

S3 08:44 If it integrated-- there's a lot of integration happening, both in Nvidia and Mellanox have new things they've got to build, because it's not PCI Express anymore. I mean there's--

S2 08:56 But they've got a pretty decent runway, to do this--

S3 08:59 Three years is not a lot runway.

S1 09:00 Here's the point, here's the point, guys. Machines of hyperscale are hard. This is the hardest thing anyone's even tried. Even a 10 Petaflop system is still hard even though they're out there.

S2 09:14 Absolutely.

S1 09:14 This is 15x or whatever, right? That's going to be really hard to deliver.

S2 09:20 But the thing is we're talking about three years. They've been working on this for a couple of years already. NVlink didn't just come out of the ether. They didn't just present the idea. It's already been worked on with their little legion of doom - with IBM and with--

S3 09:35 Legion of doomed [chuckles].

S2 09:34 Well, because you know Intel isn't going to incorporate NVlink into their CPUs.

S1 09:41 But you know they've got their own plan. They just can't seem to brand it consistently.

S2 09:41 They've been working on it. Same thing with-- Mellanox has been working with them, too. That [trudge?] has been working together for quite a while.

S3 09:51 We shall see, guys.

S2 09:52 We shall see.

S1 09:55 Here's an important point I want to make about this now because here's the nature of Coral. It's a threesome. Three deals. Two got announced this week. The third one, by statute, or whatever they drew up, has to be a different architecture. They said, "You will not buy three machines of the same flavor."

S2 10:13 Rich-- is this finally Apple finally getting into HPC like you've always talked about and predicted?

S1 10:18 No. My point is-- let's just do that math here, guys. Two shoes drop, there's a third shoe. I don't know when it's going to happen, but who's it going to be? Come on, all right. What is the alternative? It's going to be Intel, I'm sorry.

S2 10:33 Well, well, MIPS.

S1 10:34 Don't even. No. Look at what Intel's been up to.

S2 10:40 What about Spark, Rich [laughter]?

S1 10:45 I'm biased there. If it was up to the Japanese. Well, plus this little triage here, they are not going to buy Japanese machines, so let's just cast that aside, the merits of Spark aside. It's going to be Intel. I think it's going to be Intel as the prime.

S3 11:02 Yeah.

S1 11:03 And it's going to be with the--

S2 11:04 Intel itself as the prime?

S1 11:06 -- some offshoot of Knights Landing. And I'm guessing that the sub will be Cray and a bunch of other companies, but Intel's going to be the prime. And that's my bet. And if you want to wager with me right now, I will bet dinner in 2017.

S2 11:20 Let me ask you this. How happy is a Cray or an SGI or somebody like that about Intel being the prime?

S1 11:27 I can tell you right now they are already not happy. The moves Intel is making is like they're going back to the Paragon days, and people like HP are feeling bruised.

S2 11:37 I can imagine.

S1 11:38 I don't think they're going to do that. I don't think it's going to be Intel Paragon business. I really don't. I have no inside information as to that. But they're making all the moves to become everything that that was, except better. And they did their one-off and they bailed. You talk about IBM bailing. Look at Intel.

S2 11:56 Yeah, talk about that.

S1 11:56 They bailed that big machine and they pulled out of the market completely in the '90s. Does that mean Intel is not a reliable suitor?

S2 12:05 Yeah, what about that, Henry?

S1 12:06 Why is it different? Why is that different, Henry? Why is it different?

S3 12:09 Why is it different? Because they built the machine, they installed it, and they said, "We're not going to be in this market." That's different than saying, "We're going to build a machine for you, we're going to install it," and then not doing either.

S1 12:22 Point taken.

S2 12:23 Is there a supercomputer vendor that hasn't broken their pick on a particular deal ever?

S1 12:29 I don't know the answer to that question.

S3 12:31 I don't know the answer to that.

S2 12:31 I'm going to saying that there are. I'm going to say that there's definitely [crosstalk].

S3 12:36 Sure there have been, Dan. Thinking machines went out of business in the 90s. The ones that went out of business, yeah. The ones that are in business, no. I don't think so.

S2 12:47 I'm going to bet everybody's got a skeleton in their closet, at least one.

S1 12:50 Look, they've all installed stuff that went badly and they had to go back. Even Titan, right? They had to replace every single GPU in that thing, and redo the soldering or something, right? And it was on Cray from what I understand. But that was not a trivial amount of money that got spent to make that right. Cray went in and did it. That's the point.

S2 13:13 Yeah.

S1 13:14 That happens a lot. These machines are hard. That's my point, all right?

S3 13:19 Yes.

S1 13:19 And Titan is a productive, big ass machine that's working.

S2 13:24 You're still drunk, aren't you?

S1 13:27 We don't need to go there right now, but yeah.

S3 13:29 Guys, we shall see. I'd put dinner at the restaurant of your choice where is 2017 supercomputing? Restaurant of your choice.

S1 13:37 Let's do the math here. Next year is Austin. That's 2015. 2016 is Salt Lake. So I'm guessing Denver.

S3 13:47 You pick the restaurant. I'll buy the dinner if Coral and IBM will not be a happy situation.

S1 13:58 I'm guessing it's going to get delayed 18 months if not two years when it comes down to it. They'll deliver some partial machine.

S3 14:04 Then I'll be getting dinner at 2017, 2018, 2019.

S1 14:09 Who wins the bet if they deliver something partial in 2017, and the end production thing is 2019, which I'm guessing is the case?

S2 14:17 Partial of what?

S3 14:19 Partial of the system.

S1 14:20 They're going to strip some nodes that are useful, is my point, not the full monty.

S2 14:25 Well, you define useful.

S3 14:28 Yeah, that's a really good [inaudible].

S1 14:30 It runs science and engineering jobs for people.

S2 14:32 You can boot Linux.

S1 14:34 Yes. A Linux boot [laughter]. It boots up and says, "Hello, world."

S3 14:40 I'd bet against that.

S1 14:41 In 2017?

S2 14:42 In 2017?

S3 14:43 Yeah.

S2 14:44 Okay. I believe there's a bet.

S1 14:47 Guys, we could talk about this for hours, but I want to wrap this up and just ask for closing thoughts. This is a big deal. To me it is a sign that the US is making commitments to Exascale. This is a big step towards that. Exascale was supposed to

happen in 2018 based on something that is never going to happen.

- S2 15:07 It's more like '21, '22.
- S1 15:09 But, these machines get built, we will learn a lot from them, no matter who makes the threesome.
- S3 15:14 Yes, we will. If the machines come out, we will learn.
- S1 15:17 Because 150, 200 Petaflops, whatever it is, it might be 20%, Exascale is going to be way hard, but these will be the stepping stones that our tax dollars are making towards that and Exascale matters. There is stuff that our government and our scientists can do in an Exascale machine that needs to be done, including global warming, et cetera. I'm really encouraged by this news. And that's my thing, and I'll go to you.
- S2 15:44 I would agree with that. Well, any thoughts about the show, anything you're particularly interested in seeing?
- S1 15:50 Well, the theme of the show is HPC matters.
- S2 15:52 It does. It does, yes.
- S3 15:55 And it always has.
- S1 15:56 And it always has.
- S2 15:57 I'd agree with that.
- S3 16:00 Rich, I hope you're right, I don't think you are. I think we're going to have some vendor swap outs before this is all over.
- S2 16:09 I don't know, I'm with Rich on this one. I don't think they would've taken this on if they weren't going to actually do it. And again, these vendors have been working together at least a couple of years before this has been announced.
- S1 16:22 Absolutely.
- S3 16:24 I don't disagree with that, I'm just saying. We'll see.
- S2 16:27 They've got a fair amount of intellectual horsepower within those three boxes, too.
- S1 16:31 There's a lot of smart people involved and no one would argue that.
- S3 16:34 No. No arguments.
- S1 16:35 On all sides of this table, including the people that made these choices to buy this.
- S2 16:37 Not this table. Not this table.
- S1 16:40 No, well we are what we are. We're Radio Free HBC.
- S2 16:45 You talk about that table.
- S1 16:47 That metaphorical "that table". So, anyway.
- S2 16:52 On that sad note...
- S1 16:55 Can we wrap it here?
- S2 16:56 On that self deprecating note...
- S3 16:57 You guys have a good show. Good seeing you.

S2 16:59

Yes, you too. Thank you.

S1 17:00

I'm Rich Brueckner from insideHPC, thank you for watching and listening to us. We're Radio Free HPC. We'll see you next time.