LLNL Readies to Deploy 5.4 petaFLOP Cluster

Transcript from March 9, 2020
Customer Spotlight Webinar

intel

Mary Killelea: Welcome, everyone. Thank you for joining us for Intel's customer spotlight series. This series highlights innovation, industry leading companies that are undergoing digital transformation have tackled business and technology challenges and created new opportunities using Intel data central technologies and platforms. On today's customer spotlight, we are excited to feature Lawrence Livermore National Laboratory. Today's host is Tim Crawford. Tim is a strategic CIO advisor that works with enterprise organizations. Tim, I'm going to turn it over to you to kick off today's conversation.

Tim Crawford: Great Mary, thanks so much for the introduction and welcome to everyone joining us live on the presentation today. I'm excited to talk about what Lawrence Livermore National Laboratory is doing. And let me kind of move forward so you can see who's going to be joining me on the presentation today. We have Matt Leininger, who is with the labs, and in just a minute I'll ask Matt to introduce himself. So first off, Matt, welcome to the program.

Matt Leininger: Thanks. Thanks for having me, Tim.

Tim Crawford: So Matt, maybe is a level set for folks. We could just spend a minute or two having you introduce yourself, who the laboratory is, and your mission there at the laboratory?

Matt Leininger: Sure. First of all, I work in the computing division of Lawrence Livermore National Laboratory there's an office in there called the advanced technology office, which is basically the Office of the CTO for high performance computing.

So, in that office, we're busy tracking technologies, seeing which ones can apply to our mission, and integrating those technologies into the laboratories production computing facilities and providing them to our users to accomplish our national security mission. And more Livermore National Laboratory is part of the National Nuclear Security Administration or NNSA. Our core mission responsibility is the stewardship of the American nuclear stockpile, and the computer simulations and high performance computing resources necessary for those simulations are very critical with all the work done by our scientist engineers at the lab and supporting that mission.

Tim Crawford: That's great. And so if we start to kind of maybe drill a little further into that, maybe you can talk a little bit about your perspective on the problem that you were working to solve, what were some of the challenges you were having, but more importantly, what is it that you're really trying to do with these computing resources?

Matt Leininger: Sure, yeah. The primary mission I mentioned—that stockpile stewardship—we have secondary missions and everything from non-proliferation and counterterrorism all the way down to just fundamental science. The high performance computing aspects of that mission involve the development of predictive physics based models. Those models are models for such areas as material science, molecular dynamics, particle transport, hydrodynamics, mathematical solvers, and other areas. The computational workloads that use those physics-based models and simulations ranged from medium to high spatial resolution 2 in 3D simulations, various timescales and conducting numerous parameter studies. So how we mix all that up is we basically have two types of applications that's how we think about it. One is a multi-physics application which can involve a combination of things like hydrodynamics, particle transport, and involve the most complex geometries, the models that we use are very approximate. And so when we run very detailed simulations on those, and this is where the parameter studies come in, you can study how well those models are and how well the uncertainties those models arise.

Once you know the uncertainties, you can then use more accurate [science] based applications such as a particular hydrodynamics algorithm that may look at some uncertainty in the hydrodynamics, allow you to figure out what the fundamental physics that is going on there, come up with better physics based approximate models that go into the multi physics simulations. And then repeat the process over and over until you get basically a predictive science capability for these types of simulations. That's the problem we're trying to solve at a high level.

Tim Crawford: At a high level. That's a really complicated problem to have to solve for though. You talked about the computational requirements, and you talk about the technology that kind of fits into solve some of these physics models and the missions that you're going after. Maybe you can spend a few minutes talking about the partners that came into play. And because I think this is an interesting part of this story, it's not just about the technology itself, but it's about how the technologies and your partners came together. Matt, could you spend a few minutes to share with the audience who those folks are and how they played a role?

Matt Leininger: Sure. So this all started at least most recently, back in 2016, when NNSA awarded our commodity technology systems contract which was called CTS1 for short, to Penguin Computing, and that involved Penguin working with Intel, which Intel select as both the processor and the high performance network provider and we've been purchasing systems under this contract for the last several years. And the partnership with Intel and the interconnects is particularly critical in our case, because the physics simulations we're doing are very challenging as I mentioned. You can't just solve them on a single server, you really have to break up the problem distributed across thousands of servers, and then use that high performance interconnect to tie the pieces back together again, to kind of reassemble the problem.

That's turning a problem that was extremely complex to impossible on a single server to one that you can solve in a few hours or a few days. So the partnership with Magma is the latest system is being delivered to live more to this contract. And again it utilizes Penguin as well as Intel's latest generation processor, but we've also added liquid cooling technologies from coolIT. So the partnerships that we've been evolving here under the commodity technology systems contract for the last few years, have all come into play and kind of culminated with this latest machine Magma.

Tim Crawford: I do want to delve into Magma and delve into the cooling technology and how that innovation is really playing a role. Because I think for folks that might have HPC environments, high performance computing environments, or even in the enterprise space. This is a pretty new technology that's coming into the mix. So let's maybe double click on Magma in terms of what is Magma and how does that fit in?

Matt Leininger: Sure. The mission for it it's the latest system that was just delivered to Livermore under our contract with Penguin. All the systems that we've purchased under this contract previously are still in very heavy use, and we have a great deal of demand for more and more computing cycles, strictly directed towards our stockpile stewardship program. And the way Livermore goes about designing our systems in general is we use a concept called the scalable unit, where each scalable unit you can think of it as basically a Lego building blocks that includes not only compute nodes but also some infrastructure nodes.

Then you're allowed... you can build up larger scale systems by SUS together. And so that's how we basically go from a system of maybe a scalable unit of a couple hundred nodes up to something or 3000 nodes or more. And Magma in our terminology, Magma is about a four scalable unit system. Physically, it fits into about the size of half a tennis quarter. So it uses the second generation Intel® Xeon® Platinum 9200 processors, as well as the dual rail Intel® Omni-Path interconnect. And in total, it has 760 compute nodes. About 70... a little under 73,000 cores, theoretical peak for double precision performance of 5.3 peta flops, and about 300 terabytes of memory capacity. And so, as far as our workloads that I mentioned before kind of the first order, although they are intensive on the network, they're most intensive on memory bandwidth. And one of the things we really liked About the [Intel] Xeon Platinum processors as they have a tremendous amount of memory bandwidth per node. And therefore, we can kind of remove that bottleneck at least a bit for our applications and then deliver both in a very capable and economic HPC cycles to our mission critical applications.

Tim Crawford: Okay. And so if we look at the technology that's into Magma, and how coolIT kind of came into this, I want to talk about the innovation specifically around Magma. I mean, Magma itself is a massive footprint. It's amazing technology that you're tapping into. But let's maybe spend a few minutes and spend some time talking about each of those components and breaking that apart, because I think this is the core piece that differentiates this specific technology but also how you've put it together. So can you maybe talk for a few minutes about some of the newer innovations that are introduced as part of the Magma implementation?

Matt Leininger: Yeah, we discussed one that's the memory bandwidth per node is tremendous, which is one huge innovation that we really like to see. Another one that I mentioned briefly was the coolIT liquid cooling technology that was built into these [Intel] Xeon servers. And the CoolIT solution provides directed chip liquid cooling, that brings a coolant right to the components, particularly the CPUs and the memory modules. For the CPUs, this is very critical, we would not be able to achieve the overall performance of the system or the density that we got at Magma without direct liquid cooling. Plus in general, as a production computing resource, one of the errors we see in general across production machines is memory errors are sometimes the biggest cause of issues on machines, basically errors on the memory DIMMs. And so the direct liquid cooling on the memory module as well we're looking forward to reducing the operating temperature of the DIMMs and hopefully therefore reducing the overall number of memory errors we see over the system lifetime of Magma. The other piece we liked about the liquid cooling design from coolIT on Magma was, it's really the ease of maintenance and serviceability of the system. There are a couple of examples of this.

First is the use of the blind-mate and dry-break connectors. So you don't require a manual disconnect. So you can unplug the server and it will automatically unplug the coolant lines without leaking any coolant. Second, is how the coolant is delivered, I mentioned that we just delivered now into CPUs to the memory DIMMs, but the way they have it set up to memory DIMMs, it doesn't interfere with the serviceability of the DIMM so it makes this basically a very easy machine for our admins to administer and maintain and service. It doesn't make it that different from things we've done in the past that have been air cooled, which is a really nice feature. So all this basically adds up to very clean, easy to maintain, easy to service, liquid cooling solution for these commodity systems.

Tim Crawford: Can you maybe talk a little bit about your experience with cooling within the data center, for example, is this your first time working with cooling systems within a system?

Matt Leininger: No, it's definitely not a first we've liquid cooled and air-cooled machines for a long time. We've started over the last about five years to move these commodity systems from being 100% air cooled to now being more and more liquid cooled particularly to the least, the components that are generating most heat like the CPUs DIMMs, or even the future GPUs. So we have a lot of experience in doing both air and liquid cooling. But one of the challenges we found was continuing to do air cooling is a couple. One is just facility wise, If you have a really big machine that's all liquid cooled and then you also put in your data center, you know several thousand node cluster with air cooled becomes an engineering problem of how you optimize it. Get all the cold air flowing in the right way, and that often takes weeks and months for people to figure out how to fine tune that and often just takes a lot of manpower.

We found that in the future, we think it's going to be much easier in terms of data center design and delivering solid systems to our users of moving our commodity systems be more and more liquid cooled. The second aspect of it is more than technology, we're starting to see CPUs, the power per socket is going up. They're starting to crush thresholds, where to get the really good performance per socket and kind of the best bang for your buck, so to speak on HPC workloads and starting to require the CPUs that need to be liquid cooled, in order to get the densities and the performance of the systems and so those two things that we see both are facilities and where technology is going in the market is driving us to bring liquid cooling into these commodity systems.

Tim Crawford: So if I understood you right Matt, what's you're saying is, in order to get the performance that you really need, you have to leverage liquid cooling to take advantage of these technologies at that performance level.

Matt Leininger: Yeah, it's not only just to get the highest performance, it's to get whatever the best performance per dollar is, which is really what we're after in these commodity systems. It used to be we could kind of get away with really good air cooling and still get that best bang for our buck so to speak. But with the technology advancing and things becoming more dense, the power per socket going up in order to stay in the loop, it's driving us to move even these commodity systems to be liquid cooled as much as possible.

Tim Crawford: How does these systems kind of fit in with regards to the existing rack? When you think about the PDUs you think about all of the networking equipment that goes into these racks where these cooling components kind of fit in, how does that affect to airflow? Because that's a core piece to the story is understanding how to most effectively design these systems in such a way that you're taking advantage of the technology, but the cooling doesn't impact airflow. Can you maybe talk for a few minutes about that?

Matt Leininger: Yeah, sure, we weren't going for 100% liquid cooled, like I said before, we really just want the components that generate the most heat like the CPUs, the DIMMs, perhaps in the future GPUs to be direct liquid cooled. So you're talking about something that you may be doing 70 to 80% liquid cooled then the rest, air cooled a little bit. So one of the nice things with Penguin and Intel and coolIT working together on this was that they could engineer a system that could deliver all the liquid cooling, but at the node level and the rack level and the system level. You don't have components interfering with that little bit of airflow that you do need through the system. And so overall, it's a very well-designed system from that perspective of getting us this gets back to kind of getting the best bang for your buck on these systems.

If you really have to do 100% liquid cooling, you have to get into chilled doors and other things that start to add complexity to the system, which although engineering wise is possible, it's certainly not always the most cost effective way to go about doing this. So sticking to try to get about 70 or 80% of the heat captured through direct liquid cooling, and then the rest, look a little bit of airflow is really what we're going after here. And Magma in particular, from our three partners is very well designed to achieve those goals.

Tim Crawford: Okay, great. We kind of think about where we go from here. You're throwing a lot of information to us, a lot of great information, a lot of details. But what's next, where are you going from this point forward?

Matt Leininger: All the CTS one systems that we procured over the last four years now, including Magma will continue to deliver HPC cycles to our users over the next several years. We'll continue to track technology roadmaps I'd mentioned as part of the group that I'm in at Livermore. We are preparing for our next round of commodity technology systems procurements could occur starting in late 21. And that will be under kind of a second round of the CTS procurements called CTS2, pretty simple. We expect to have an RFP out for that later this summer and then use the contract award and hopefully late this calendar year, and then start delivering the next round of next generation systems so to speak to NNSA starting in the second half of calendar year 2021 and then going through 2024. So that's what we're really looking at right now, we still have a few systems that we need to deliver between now and when CTS two starts. So, CTS one will be ongoing, all the systems we've deployed, they're still running and delivering cycles to our users, but we always have a growing need for the more and more people demand for these types of systems.

So we basically have a continuous ramp and procurements in order to provide those to our users. And so CTS2 will be the next big round here in the next few years.

Tim Crawford: So this is not a one off situation, this is a continual cycle that you go through?

Matt Leininger: That's correct. And even though it's we call them CTS1 and 2, they have a lineage that goes back about 15 years. So really CTS2 to be our fourth round and doing this, where we do have large beginner like this, and I didn't mention before but it delivers systems not only to Lawrence Livermore, but to our two sister laboratories, Sandia National Laboratories and Los Alamos National Laboratories as well. So the for these types of systems were all three labs using the same type of hardware and supporting the same types of systems for our users.

Tim Crawford: Great, Matt. Thank you. So at this point, I want to shift gears a little bit. We've got a number of questions that have come from the audience. So the first question it was mentioned, you hope to get improved error rates on memory, is there any data correlating memory error rates to temperature?

Matt Leininger: We have some data matches at Livermore but between Livermore San Diego and Los Alamos where we do seem to have some errors that are correlated to that, there's memory errors of very complex things that can come about in a lot of variety of ways, but temperatures at least one factor in it. And so we are hoping that by liquid cooling the DIMMs on Magma that we will see a reduced rate of the memory errors, but this is also bit of an equivalent experiment as well. We will be gathering data on Magma over its lifetime to see if that really holds true.

Tim Crawford: And you had mentioned before that one of the biggest issues is the memory reliability due to heat.

Matt Leininger: It's definitely the memory reliability. We expected heat at least one component that can trigger the errors and over the lifetime of any machine, we typically see [swapping out] there's one thing that's just kind of continuously done, even when the machine is very stable DIMMs can eventually go bad, you have to swap them out. So we're hoping by bringing the temperatures down inside the node, directly cooling the CPU as well the DIMMs at least reduce that. How much it reduces it, we'll just have to wait and see.

Tim Crawford: Got it. So we have another question that's kind of along the same lines. Is there a specific heat of the coolant itself?

Matt Leininger: That's a good question and the answer is yes and I don't know it off the top of my head. So we'd have to get back to them on that one.

Tim Crawford: Okay. All right. So kind of along the same lines is memory. Memory bandwidth per node was mentioned as one of the primary innovations of Magma. But I didn't catch the specific number two total memory capacity was mentioned, but I missed the total storage capacity of the system. At this stage, is that something that you can get into or? I know you're still working out some of the details there too.

Matt Leininger: Yeah, the total amount of memory on Magma is about 300 terabytes. So in terms of DRAM capacity, that's the right number. In terms of storage, such as a parallel file system. There's no one number it actually...Magma was purchased as an individual unit. We purchased some additional storage through other procurement mechanisms as well. But it's tied into several different monster file systems that will service it. So there's not one firm answer to that but it can basically connect into many, many terabytes of storage.

Tim Crawford: Okay, great. So question about Magma itself. So the Magma systems that you have in place today, are they currently in production or are you in the process of bringing them into production? What's the current state of Magma?

Matt Leininger: So Magma has been delivered to Livermore. It's undergoing final stages of installation before undergoing acceptance testing, and then it'll be moved into full production here in next month. So that's where it's at right now.

Tim Crawford: Great. Our next question, how long have you used the prior generation Xeon processor-based systems [Intel® Xeon® processor-based] and how long do you expect to use the second generation [Intel] Xeon Scalable [processor-]based systems?

Matt Leininger: So our previous [Intel] Xeon [processor] we started pushing back in 2016 those are the E-5 2600 in the four series if I remember correctly, those have been running for almost four years now we expect them to run through for at least a few more years. Magma, we expect once that goes into production that will probably serve our users for about five years, maybe even six years is kind of the average we're seeing on a lot of these clusters today. So that's kind of what we expect these new things to survive. Then the same thing is true, once we get the CTS2 systems later, they're kind of we're averaging about five or six years per machine.

Tim Crawford: And you mentioned, the CTS program goes back about 15 years, and if you're getting five to six years out, you're currently in your... I think you said third or fourth generation that the math kind of plays out right?

Matt Leininger: Right.

Tim Crawford: This next question, I think you addressed it kind of to my earlier point, but is this your first use of cooling technology in the data center?

Matt Leininger: No, we've used lots of different cooling technologies. Historically... most recently over the last maybe five years started to use it more in these commodity systems, and this is about the third solution I think that we've used in that space. But we're still learning a bit we have a lot of experience of delivering these very highly integrated systems that are almost over engineered, so to speak. They're very expensive, but they do a very good job as well. But with these commodity systems, where these are the workhorse systems, these are the everyday systems that all of our users get on to and use for code development, testing, running jobs, all that stuff, and we're trying to deliver the most cycles we can with the given dollar amount. So we're trying to find that balance to the best cost performance. And like I said, we're trying to bring in liquid cooling technology into these commodity systems, but do it in a way where it's still cost effective, and allows us to get as many cycles as we can and deliver the most high performance as we can.

Tim Crawford: Sure. It sounds like to some degree that this is also an experiment too. You're learning about what the capabilities of what you can do for these different projects using these different technologies, both in your work there at Lawrence lab as well as Sandia, etcetera.

Matt Leininger: Absolutely every system we deploy kind of fits into that category. We learn a lot from it. We provide that feedback to our system integrators not the component suppliers. Magma will no different, will provide a lot of feedback to Penguin and Intel and coolIT on how all this technology is working, and work with them over the lifetime of the system to improve anything that we can as well as feed that into the next generation of products.

Tim Crawford: Cool. So I have our next question. Do applications span both the existing private Generation [Intel] Xeon processor-based systems and the newer second generation [Intel] Xeon Scalable [processor-]based systems?

Matt Leininger: No single job will span between the two different systems, but we do have programs that have a huge demand for computing cycles, and they'll kind of spread their workload over both Magma and the previous generation of CTS systems to get all their work done.

Tim Crawford: I would assume that at some level, the more systems you have, and the more capability you have, the better off you'd be, you'd be able to process faster or be able to handle more work coming through the door.

Matt Leininger: Absolutely.

Tim Crawford: Okay, next question that we have. Will you be retiring the prior generation [Intel] Xeon processor-based systems once Magma comes online?

Matt Leininger: No, as I mentioned before, we're definitely not retiring the old systems, even though there are a few years old, they're still stable, they're delivering cycles, we can keep them running. So we expect that they'll probably be around another two to three years before they'll be retired.

Tim Crawford: I would assume, as much as you can get out of them partly because of the expense partly because of, you just need the capacity you would want to use them as long as you could.

Matt Leininger: Yeah, and once they're up and running and stable for us these types of systems you just run them until either they're dead or you really need the space and the power for the next generation system.

Tim Crawford: Yeah, I know that some organizations are thinking about that power density compute algorithm trying to figure out what's the right time to start to retire equipment and bring in new equipment. So it sounds like you're thinking along the same lines.

Matt Leininger: Yeah.

Tim Crawford: Our next question. So this is specific to Intel® Optane™ technology. Are you using Intel Optane technology as part of Magma?

Matt Leininger: Not at this time, we have been investigating Optane [Intel Optane technology] for a variety of workloads, and other non-volatile memory storage technologies. But when we designed Magma, the expected workload did not really require any additional storage printed. So we didn't put anything on Magma at this time.

Tim Crawford: Okay. Our next question, and I think we've covered this in a number of different ways, but I'll ask it anyway. Why did Lawrence Livermore National Laboratory need a system like Magma?

Matt Leininger: So our main users dealing with the stockpile stewardship program had... they always have a huge demand, but that demand through new hires and increased workload and programs coming online, the demand went up enough that the program decided to buy a new machine. So it really traced back to our users needed more cycles to get to work done that they need to get done.

Tim Crawford: Great. Next question that I have here is kind of reading through it here. Can you talk about the CTS1 contract and why these cluster systems are so important to the NNFA?

Matt Leininger: Yeah, I've touched on some of that, but I think I mentioned this before. The CTS systems they're really the everyday workhorses for our NNSA simulation workloads. These are the first system users get on to when they're doing code development, application, integration, testing, debugging, new user training amongst other things, as well as heavy workloads that they need to do to get their production work done. These machines are utilized by nearly every scientist engineer doing simulation work at the lab. And then it's a subset of our users and those workloads that move from the CTS systems to needing even more cycles and that's when they move up to chain and go from our commodity technology systems up to our advanced technology systems. So these ATS systems are the bigger brother to CTS. They're the ones that are kind of pushing the state of the art forward in terms of system size. But that's how it all kind of fits together at Livermore.

Tim Crawford: Great, Matt, thank you so much. And with that, I'm going to turn it back over to Mary to close this out.

Mary Killelea: Thanks again everyone for joining us today. Please look for other exciting customer spotlights that highlight data centric innovations coming soon.