

TOM SALOMONE: So first, thank you for coming to the first session. Welcome to AU. I mean, this is a great adventure. How many people are new at AU? Yes, yes, very good. So I think you're going to have a great experience here at AU. I know people are really-- everybody, every company that comes puts on their best effort here. You're going to see things here you won't see other places. So I think you'll really enjoy AU. So before I begin, I also want-- I'm Tom Salomone. I work for Lenovo, and I do this session every year at AU. And some people come every year to learn about what's new in the technology. So I try to talk about the latest and the greatest so you can understand what choices you have and what it really means to you, based upon your application.

And besides me, I also have Denis Sessanna. Excuse me. Easy for you to say-- Sessanna. And Denis is a AE at NVIDIA. So he's a real technical person. So not only can tell you the overview of the NVIDIA stuff, but he can get down into all the details about applications and everything else. So if we don't get to your question during the presentation, make sure to stop by and either see us here afterwards, or see us in our booth afterwards as well, because we're very happy to answer your questions.

So with-- So I'm going to get started now with the regular presentations. So this is the topics I'm going to go over, the basics about how our system works. So some of this may be insulting to some people, because they know that already. But some people don't know it, so I'm going to go over the basics. And that just leads, naturally, into why we certify. So I'll explain that a little bit, then talk about processor selection, talk a lot about Turbo Boost 3.0. It's new this year, and I want to make sure people understand what that means to them. About tuning workstations, we're going to talk about graphics cards, rendering, selection, performance, some technology changes, some uniquenesses we have at Lenovo you should know about, some emerging technologies, like ARVR, point cloud capture, new things. There's some new things there, and there's even more that we could talk about than just some quick summary comments.

So this is the basic outline of a processor from Intel. So you can see the name of this actual processor. So when you go online, and you look at an Intel processor, you're trying to decide, you're going to see all these acronyms like UPI and PCIE, and so this is what this means. So the first one I show is cupie QPI. That's a very simple one, quick path interconnect. And then,

when you look at the Intel chip, it's either going to have one or two, right? It's either going to have that where it communicates to another processor but you don't have to worry about that. It's just information for you. So it tells you whether you can use two processors or one processor.

Then is the memory, right? So the memories-- the memory, here, that we're showing with the Xeon has four memory channels. There's others that have two, but in the I7s, I5, I3 range. They can have less, but they can also have four. And so you want to look for four memory channels, because that's the best way to optimize your memory and your memory performance. And we're going to go over that a little bit later. PCIE, there's a lot of new things going on around PCIE that started maybe two years ago now. PCIE, PCI Expr-- or PCI Express, right? Pcie is-- has always been used for graphics, right? So that's been the main graphics bus. But now there's a new storage device called M.2 in terms of the form factor, and so that storage device now uses a PCIE bus, which is a much faster bus.

And then there is a chipset. So it used to be you always had to choose or look at the chipset that corresponded to the processor, but that's always one for one. You can't really select that chipset. And so that's how the system communicates to the USB, to the SATA and SAS drives, and other things, like ethernet and so. So that's the hardware side. Now there's the software side. There we go. OK. So you start with the system hardware. Of course, that's all run by ones and zeros. And then, there's the bias, right, which is no longer really named the bias. It's really named UEFI, and that changed in Windows 7, but everybody still refers to it as biased, because bias had been used forever, right? It seems like with computers, and UEFI doesn't, kind of, roll off the tongue as well either. I think so.

And then-- and so in there, that's where you find your system drivers, your graphics drivers, everything on the board has a driver, right? The processor has a driver. Everything has a driver, right? So you're probably familiar with graphic drivers, because you have to update those kind of regularly, but everything has a driver. And so there's-- and all that is driven by the Windows operating system. And the graphics, specifically, is driven by either the open GL or the direct exits in the Microsoft Windows, and used to be, as some people may know, open GL was separate from Windows, and over time it got incorporated.

And so now on top of that is your CAD application, AutoCAD, Inventor, Revit. And there's two things that really distinguish CAD applications from most all other applications. And that is they have a kernel, right? And then they have a section that is really what's displayed to you, your

user interface, things like that. What the kernel is, the kernel is the formulas, the mathematical formulas, that keep your design accurate. So when you're designing something, you can measure it, right? You can do a lot of engineering calculations with it. And all that is done with this very complex mathematical calculation.

And so when they came out, these kernels, right? This was a real breakthrough. Thousands upon thousands of man hours were put into developing these kernels, and because they are-- they follow math, right, they're not changed very much. Now you can't display the mathematical formula, right? If you could display it you would-- it'd be useless. So they developed interfaces for that and that's what you get with AutoCAD and what you get with other applications. So AutoCAD used to use the Asus Kernel, if you've ever heard of that phrase, right? So that's the mathematical section. And that's really important to you, because the kernel is very serial.

And so when you make a change, right, it's going to go through all the math in a very serial way. It calculates part of it, then takes the answer from that and plugs it into the next formula, on and on and on. And that's the way it was developed. And people have been trying forever to make it parallel to take advantage of the multiple cores, but it's way, way too complex to do that. And so this is true with all CAD applications. Every CAD application out there has a kernel, and then what the companies, like Autodesk, work on really is the display and the UI. They don't do very much to change the kernel, except maybe to get more access or some other things.

And then another thing you should know about systems is that every manufacturer has unique UEFI, right? So everybody has their own unique one. And so that changes some things about what you can do with the system, how it can perform, things like that. So that is unique.

Now, so I think that discussion helps people, maybe, understand certification. And when you think about certification, this is an example with Lenovo. It could be similar for HP or Dell or whoever else, but you see our desktop systems, right? You see our mobile systems, and then you see the graphics cards that go with those and the graphics cards that go in our mobiles, and then you see the different operating systems. All those combinations have to be checked and tested, right? So we know they're right. So you can just start thinking about how many combinations they are. And when-- and when you actually figure it out, this is what it looks like. For our desktop think stations, it's 75 different combinations, right? And for I think pads it's 36 combinations.

So those are the combinations we start with, then we have to test it for every application. So for Autodesk, it's Autodesk applications and then versions. So there might be two versions or three versions currently used. So we might have to test it for more. And so then you have to multiply that times the graphics numbers. So you can see, once you start with a 100 or whatever it is and times maybe 20 of those and two or three graphics drivers, right? So pretty soon it gets to be a really big number in the thousands. And so we have an engineering team that works on this. So does an Nvidia. They have an engineering team that works on this. So does an Autodesk, right, that makes sure all these combinations work. And so nobody ever really talks about it, but they just say it's certified or it isn't, but it's important to you, because that's the only way you're going to know it's really been tested by us, and so that it'll work.

OK. So let's talk about processors. But one thing I did want to mention, going to go back here just one quick second. When you see this layout, right, the software components, and the display, you can see that the CAD application is really taking over control of the graphics. It used to be, in the old CAD, you actually had to hit the Escape key to get out of CAD, to get back to the operating system, but that's-- now it's much smoother obviously. But it does, the CAD application is taking over and what you see on your screen. And that's a lot different than other applications, as well. So obviously Word and PowerPoint don't do anything like that.

So in processor, so the first thing you have to choose from processors, because of the way the software is architected, is CAD or CAE or kinds of things. So in the CAD, gigahertz comes for us, because it's generally a single processor, but it can expand a little bit to take advantage of some of a second core, as well. But CAD-- for CAD you're choosing gigahertz first. for simulation, rendering, those types of things, data processing, you can use cores first, so you can get multiple cores. Or, we're going to talk about later-- a little bit later, GPUs and how important that is to consider when you're thinking about those types of applications, because you may need tho-- you might be able use GPUs instead of cores, and so-- and gain some real advantages. So this is the first decision that you make.

And then you go to the Intel web page, and you can look at the different processors that are available. So this is a series of processors that have been out in the market, and starting at the low, going to the newest. And so you can see that the big thing to look at is process in the middle, where it's 22 nanometers versus 14 nanometers, right? That's the size of the-- half the size of the distance between two points on a chip. This is smaller than the wavelength of light. So if you think about that, this is extremely tiny. And as a matter of fact, I think I put up there,

point two nanometers is the size of an atom.

So you can see 14 nanometers. There's not too many atoms distance anymore. That's how tiny these things are, and processors. And that becomes very difficult for Intel to make improvements. And they're looking at 10 nanometers and seven nanometers, now. They're trying to develop production methods for those. Well you can see how it's shrinking, right? And as they get closer and closer, right, the ability to get closer is going to be a real challenge. And so people are looking at things like light in the future and other things. So it's hard for Intel to continue on the great rate that they had maybe five years ago or ten years ago with constant improvement in performance.

So what does that mean to you, right? So on our desktops, and so I think this is similar with others as well, right, there's the core set of Intel processors and then the Xeon E3s. And so those are in the low end, which is our P310 small form factor, and the 3-- P310. And then there's the E5s, which are in the other ones where you can get dual processors or single processors. And so this is the dual processor ones we have, and for those-- for dual processors you would use those for 3ds Max, maya, or any finite element analysis activity.

And for the low end, right, our entry level, that would be great for AutoCAD, but Revit, Inventor, you'd probably choose something at a higher level. So you can just see the differences between our desktop products, and so that's really a good way of thinking of it. And at the end, our 410, that's a new product we introduced. It's a low cost-- the lowest cost entry level of the E5 products. And E5 is going to become more important to you than you think when you look at turbo boost. And we're going to show you that in a minute. You may be at the other end, E3 or I7s, but when we started talking about turbo boost 3.0, I think that'll be more important.

And so I just want to show you an example of benchmarking. So these are benchmarkings we run, and this shows something that is typically two years old versus current. So you can see that the two-year-old one that we rated as a one in the benchmark. And this is the catalyst benchmark. And you can start to see-- see if this works. Oops. Nope. I don't know how I did that. There we go. Sorry. Hit the wrong button. Operator error again. Let's see, how about this? This is review. Everybody remember? OK. Let's see. All right.

So you can start to see the difference. So you can see, if you look at the CPU index, there wasn't much of a jump, right, from two years ago today. But the disk certainly is a big index. The index changed. The graphics, there's a big change. In the total, to you, there's a big

change. But if you look at it by subcomponents, you can see-- you can really see the difficulty Intel's having trying to improve performance version over version. And that's important to remember. So and we're going to go-- we're going to talk about how to deal with that in a minute.

And so then, here's an inventor benchmark. This is a relatively new benchmark, I think. And so you can see they do benchmark differently, with drawing times and graphics and our hard disk drive. And you can see that old to new know there is a significant performance. Now the same is true for mobile workstations, right? So here is the mobile workstation processors. And what was new last year, a year ago this week, I think, was the mobile-- the Xeon mobile processor. And so that's now been put into mobiles. And this is how it looks.

So this is the mobiles that we have. We have 14 inch P40. We call that a yoga. You can flip it in and make it a tablet or flip it all the way around and make it a workstation. It has [INAUDIBLE] capability. It's really cool. We have the P50s, which has an extra battery life. So that can actually go 12 hours. Our civil engineers love that, because they can take it out and don't have to worry about plugging it in for a whole day. The P50, which is our most popular one for CAD, and our P70, our new 17 inch, right? So that's what we have and others have something-- similar things.

And so the Xeon is really put in the high end, right, because the high end workstations for mobiles are the ones that can handle the power draw that these have. The other ones, the entry level ones, have lower power designed processors and graphics parts so that the batteries can last longer. And so if you're thinking about using a mobile workstation, right, the P70 with GPU acceleration is really good for these applications. And so you want to pay attention. We're going to go over that a little bit more, but-- if you're doing that, but obviously it only has four cores, because that's really the limit on mobile. So core-- extra core count won't help you in the mobiles. Then these are really great for AutoCAD, Civil, Revit, Inventor.

OK. So I want to talk a little bit about turbo boost, and help you understand turbo boost. And so this is the slide I use every year, but it's from an Intel white paper on turbo boost, this chart. And the gigahertz rating is really at-- it's set depending upon the thermo design power limits. So they decide this based upon the distances and everything. So if you hit this speed, it's not going to deteriorate the chip over time. If it goes above it when it's at-- when the processor's hot, it can deteriorate the chip, and make the life of the chip much less.

And so they found out that, hey, when it's cool, right, we don't have to worry about that distraction. So we can boost the speed. So we can boost the speed up, and so they call that turbo boost, and turbo boost one, two, and three. And what it automatically does-- you don't have to pay any attention to any of this stuff, because it does it inside the chip. But it looks at the workload, the active cores, estimated current, all that stuff, to figure out, can it go faster or not. And the only thing you have to do with turbo boost 2.0 is make sure it's on or off, if you don't want it. And you can say, why don't you want it? Well some people, their software-- not Autodesk-- their software is designed to be parallel and meet up, and if some has turbo boost, other isn't, it screws up their timing. But for everybody here, turbo boost, you'd want turbo boost to be on.

And this is a little history of turbo boost. So it's just, sort of, important to remember. So when turbo boost first came out, everything was boost-- excuse me, the cores were boosted, which is the real CPU part of the process. And you can see the layout of the processor. So the cores, many years ago, that was just the CPU, right? They used to have-- for people that have been around, they used to have floating point chips and CPU chips and things. All right, but that's the main cores.

And then-- and they found that, that worked OK, but not really great, because the cores would run faster, but it needed to be fed faster in order to take advantage of it. So they decided to boost everything together. And that worked a little bit better, but still not what they expected, because the chip-- because everything was sped up, the chip would heat up too much and too quickly. So they lost time. And so they came up with this idea of making it energy efficient and have the two different elements, the core and the non-core, separate. And so that was the 2.0. And so that it could then figure out what to do and speed things up when it needed to in certain areas, or not. And then Broadwell came out, and they came out with 3.0. And this requires certain processors. And those processors now have-- they have special turbo cores that can go higher, right. So the other ones, they are the standard cores, and they just overstress them. But this one has special cores just for this. So it can last longer. And so when you're running CAD, the reason this is important is because you're taking advantage of the single processor, right? So you're taking advantage of a single core. And when you make a design change, it goes through all those calculations, it spikes up to the maximum, and comes back down when you're done, right? And you can't work as fast as a processor. These are billions of cycles a second, so it's going to sp-- it's going to spike up and then come back down, right? And then you'll make your next change, it's going to spike up and come back

down. And that's how you-- that's what's happening when you're doing your CAD applications is it's spiking like this all the time. And so we can take advantage of the extra speed and give you extra performance.

And so with turbo boost 3.0, they're supported processors. There's this-- you need a supported processor, a supported operating system. You need the Intel driver. You have to download this driver. It has a user interface, so you have to interact with it, and you have to make sure its enabled in the BIOS and the firmware that supports this feature. And so the processors have at least one core with a maximum turbo boost frequency higher than the other. That's what makes these processors unique. And you can assign key applications to the fastest cores using affinity functionality, if people know that. And what an infinity does, you can look it up on Google, it really lets you assign your application to a specific core.

And you can then, also assign all the other applications to another core. And that's really important for performance, and I'll show you in a minute. And it does require a user set up. So it's not on or off. So this is something you're going-- you'll need to pay attention to now that maybe you didn't pay attention to before. And turbo boost-- the official title is Turbo Boost Max Technology 3.0-- these are the processors, there's Xeon and Core i7s that are available with this functionality. There's operating systems. There's only the 64-bit, right, 7, 8.1, and 10. So these are those. And then if you go to the Intel web site, on their download page, this is right from their download page, you'll just see that each processor, you need to download the information, the drivers, and the program for each processor.

So this is how you get to turbo boost 3.0. So if you're choosing, you want to look for those. And then I'm going to show you what it does. So I'm going to explain it, then I have a quick video I'm going to show you where we explain it even better, but that's the task manager. It's not showing much usage, but down here on the right is the Intel software that you would interface with. And you can see, if you can read it, the lists on the right in the cores that are the highest priority, in terms of turbo boost selection. So the top one has the extra turbo boost, and that's listed in performance.

So they'll list the fastest one at the top-- as the top one and then down. And so we were running this. We took the snapshot, and you can see this as an Intel tool that we use to look at each core and know how it's performing. And you can see the pinkish color, right? That core is maxed out. That's the one we get assigned to the process. So now I'm going to run a video and you're going to see-- and hopefully the sound is going to work. No. Nope. Didn't-- didn't

work. OK. I guess I'm going to do my sound. OK. All right. So I'm just going to point to this and - so I don't know why it's not running, but it's not.

So here we have system information. Again, we have the Intel application where you'll see the cores. You can actually take AutoCAD or Revit, and put it into this white section that has nothing, and you can assign your applications to that. And that's what I would recommend you do is take what application you use, and assign it so it looks at that core. If you don't, what Intel will do, they'll look at the applications that are using the most performance and they'll make that assignment for you, if you use the default. So if you have a browser that is out doing a download, right, and-- it might fall into that category. And you don't want that in that category. You want to reserve that for the applications that really impact your productivity. And then we monitored all the cores on the bottom. And you would have seen if we had the video working right. It would show you that as well.

So then in addition to that, Lenovo has introduced this-- our own performance tuner that does a lot of things similar to the app, but not just for those special proces-- processors, but for any processor you have. And it works on old systems and new systems, as well. But it does more. It looks at things like global settings, affinity, power management, visual effects, graphics, our bios, and so forth. So it looks at all these different things and gives you the best settings for your application. So if you're running AutoCAD, it's like the AutoCAD except the default settings, and then you get the best performance you could get for AutoCAD. So here's an example, and we show this in our booth. So if you're interested in learning how to do this, you can stop by and see it, but you can see the tremendous improvement by some of these settings.

And so this is really streamlining your whole computer to run your application as best as it can. And this is more impactful than, I think, a single CPU decision or a graphics card decision or anything else, right? Because you can see-- you can see how much a performance it really makes, because when we make stuff, it's really-- we have to make it more general for everybody. But you can make it more specific for what you use it for. OK. So now we're going to talk about GPUs and Dennis is going to take over and speak for a minute.

DENNIS

SESSANNA:

So good morning ever-- to everyone. Thanks for getting up bright and early to come out and hear some of our tips and tricks. Just curious, how many of you are using high end 3D rendering in your workflow? So that's quite a few people. So NVIDIA, of course, is a graphics company. But, traditionally, rendering has been CPU based. So there's quite a few different

rendering packages out there. There's many that are incorporated into your application. So you have something called mental ray that has been in [INAUDIBLE] for a long time. NVIDIA has something called Iray that they distribute. You have other packages from chaos and there Vray. Octane has a really-- [INAUDIBLE] has a really good renderer called the octane. I do spend most my time working with Iray. So we'll talk a little bit about how GPUs can accelerate Iray. Iray is a photo realistic render. So it's a physically based render.

I work with mostly the automotive industry, so pretty cars. They want to actually take and render their cars before they actually have to go and build other cars. So I've got a couple of people that-- in the automotive industry that are looking at Iray. Iray is integrated into 3G Studio Max, but also, we have released plug-ins that you can install into Maya and 3D Studio Max that has the latest version of Iray. So Iray will run on CPU cores, of course, but they run a lot faster on GPU cores, our graphics processor. So the higher end graphics that you have, the better rendering performance that you're going to get. So I think the next slide is actually talking a little bit about that rendering performance.

So what-- this is a benchmark that was done by Lenovo. It's not something that NVIDIA did. So it's a third party, so it's not skewed in any way. And the-- this is rendering time and the number of seconds, so to take and render a scene. The lowest one is just two, I think it's two, CPU sockets working together and rendered the scene, and that took 452 seconds to render that scene. If you move up and-- to the next level, let's get rid of the CPU sockets and just use a single M4000. That single M4000 will outperform two full Intel sockets in this rendering. So this is all Iray, using Iray for 3D Studio Max. Go up to the next level, hopefully people are familiar with our product line. M6000 is our highest end graphics product. Actually, we have a P6000 and P5000 are just announced by Lenovo this week.

TOM SALOMONE: This week, yeah.

DENNIS

SESSANNA:

So that's our next generation technology. So just with using a single M6000, you jump up to the render time done in 172 seconds. And then, up the line the cool thing is that as you add extra graphics horsepower, you can actually really increase your render performance. So just from going from a single M6000 and adding two CPU sockets, you really don't increase your performance that much. Today the CPUs do actually do-- they do a pretty good job at adding performance, but not as much as adding a second M6000. So going from 172 seconds, adding a second M6000 will get you up to 98 seconds render time.

And then, again, adding more CPU sockets, so having three M6000, so going from one M6000 172 seconds to three M6000, it takes all the way down to 71 seconds. And the CPU cores really-- they add a little bit of boost, but not a whole lot. So the cool thing about Iray is that you can run your rendering on a Iray on your local system. You can do interactive rendering or batch rendering. But we also have another product called Iray server. So Iray server is something that you install on any system, windows or Linux. The great thing about Iray server is that you don't need a full license of your application on that system. So you don't need a license of 3D Studio Max to use Iray server for rendering.

So Iray server has two modes. It has a streaming interactive mode. So you take your data load it up in 3D Studio Max, you connect to the remote server. It does all the rendering on that remote server and then streams the rendered pixels back to your workstation. So let's say you normally do your CAD on a mobile workstation, you need some more rendering horsepower, you could have a three M6000 or P6000 base system sitting on someone's desk, and that resource can be shared across multiple users. So I think the next-- so with-- this is batch rendering with I don't know. Is this--

TOM SALOMONE: [INAUDIBLE]

DENNIS Yeah so this is batch rendering with two piece 10-- [? 910s ?] with a total of 6 M6000s doing
SESSANNA: rendering, goes down to 35 seconds to render your scene. So pretty good scaling. And again, that resource can be shared across multiple users. So I think that's about it, right?

TOM SALOMONE: Yeah. You want to do this?

DENNIS Oh, oh yeah. The other thing is that the P910 is the highest end system. You could also get a
SESSANNA: lower cost system, the P510 and use-- put a single M6000 in that system and get similar rendering performance. So I think that's what you wanted to talk about.

TOM SALOMONE: Yeah.

DENNIS OK.

SESSANNA:

TOM SALOMONE: OK, good. Yeah, so one thing I do want to point out is that, if I just go back here, is that these aren't any old Xeons, right? These are next to the highest core counts, 16 cores each. So when he's talking about comparing CPUs, this is about as fast as you can get, in terms of cores, right? so it is pretty good, but you can also see where the 172 is on this chart, and then,

where the 171 is on that chart, but--

OK, so let's move on and talk a little bit about, well how do you select graphics cards? So that's rendering and it's an awesome thing to look at. You can-- like DENNIS SESSANNA talked about, you can cue up your renderers and get very quick response times. And so there's this new idea out there that says, rather than put expensive machines under everybody's desk, maybe take that money and buy a much more expensive machine, but have everybody access it. And so you can get things done a lot faster. So that means, to me, if you do a lot of rendering, you want to start thinking about at least exploring that idea.

And so what do you consider when you think about graphics cards? Well there's really three things, processor, memory, and driver. And if you have an old graphics card, you've all had this experience, you update the driver and things are faster. And that's true. When NVIDIA comes out with their very first driver, like they're just coming out with pascal, their main interest is to making sure everything works properly, and not just your applications, but all the applications that are out there. And that takes a lot of work, right, to do that, and a lot of time. And then, as they progress over time, they also find ways to make it faster. And so if you haven't upgraded your graphics driver in say two months or a year or something, right, you want to look at that, because you can get just some extra performance just by upgrading your driver.

OK. Now, how do you size graphics cards? Well that really depends on how you're going to use it. And so the things to consider, 3D model size, right, so number surfaces. And so if you talk to NVIDIA, they talk about number of polygons and textures and how big things are, but what does that mean to you? It means that you've got a lot more surfaces, you've got bigger model sizes, you've got different materials, maybe, that reflect differently. And so those are the things that you want to think about.

The next is the number of pixels, the monitor size, and screen resolution. So 4K takes more power than 2K, for instance, or a bigger screen, a 30 inch, takes more power than a 24 inch. There's more pixels, right? And so monitor size, all that stuff, is important, because you're driving that with your graphics card. The need for movement, some people in the AC market might never move their things, right? They might take like four or five different views and just work on those views, never do a lot of movement, but if you're in CAD, every command you're almost moving something, right? So movement is important.

And so high frame rates. So virtual reality is one that requires a lot of movement, right, to keep things accurate. It is like 90 frames a second, right? So it's very, very fast. And so there's a lot of need for movement for virtual reality. Ray tracing, lighting effects, that also affects the graphics card. And so advanced lighting techniques, you're going to want more power in your graphics card. So if you're doing these things, the more you do them, the higher end graphics card you're going to want to use. And, of course, I just want to remind people, certified drivers, because that's how we test and we make sure all these things work together.

So now I talk a little bit about displays and monitors. And so most people today are probably using a 2K display, right, full HD, it's 2.8 megapixels. Well 4K, right, it's ultra HD, 8.1 megapixels, and it's two times the resolution of 2K. It's a huge improvement. And one of the biggest productivity improvements that has had in the CAD industry is when people have gone from one screen to two screens, right? Now some people are using three screens. Anyone using three screens out here? Yeah. See? If I told you 10 years ago people would be used three screens, people would've thought I had a hole in my head, but that's really-- people are using it, because it's important. It's a productivity boost.

And so resolution is also a good productivity boost, and the prices have come down. So it's much more clear. You can't really show it on a projector, right? It wouldn't show up good. But crisp, and the prices are down. So you can see a 27 inch price now. So the prices are there, they're also available on mobiles. And this is an example of a 4K drawing, right? So you can see the 2K versus 4K. Lines are a little crisper, you just get better definition, you can see things better, and so anyway. So that's a little bit about the 4K versus 2k.

And then another thing that's happened a couple of years ago now is that the mobile graphics, right, is very similar in performance, now, to desktop graphics. So there's virtually almost no change in terms of performance if you're running plugged in so you're mobile doesn't power something down. And so that is really cool, but the CPUs are still vastly different, because your higher end CPUs they take a lot of wattage. So on your desktop, your CPUs are still different, but your graphics is almost the same.

And so this is what NVIDIA recommends for graphics card positioning, and so these were only guidelines. They aren't firm. But if you want it-- if you had-- if your typical AutoCAD Revit Inventor, you'd use something like an M2000 or an M2000M, which this is showing mobile graphics, right? But if you're going high end, you definitely want the higher end graphics, so like 3Ds Max and things like that. And so like DENNIS SESSANNA said, if you have Iray's

server, you can take advantage of those GPUs. So you can do that rendering on your mobile, right, if you have the high end mobile graphics. And so that's something to really consider.

So now I'm going to talk about some of the other technology changes that have happened over the last couple of-- last year and a half or so. And so one thing that's happened is this new M.2 SSDs. Remember I told you these now interface with the PCIE. They're no longer a SATA bus or a SAS bus, this is a new bus, right? And so they look like a memory stick, right? And so they call them SSD solid state disks, right? And I think the only reason they call them disks is because everybody's used to the word, right, because they really should be SSS, right, solid state storage, but they're solid state disks.

And so in our think stations we took advantage of this, and we put in special PCIE buses just for these devices. So in our case, you can take the device, put two of them in this carrier, all right. We call it flex. And then you can fit two of those in there. So you can actually fit four of those into our higher end workstations. And so it's really something unique we did. And this is what we looked at, and we saw, we did all these studies about drives. And you can see SATA, SAS 1500, RPM used to be the fastest disk drive, right, then they went to SSDs. And you can see that. And now the PCIE M.2 SSDs, you can see what the performance delta is. And so the M.2a, if you have SSDs today standard, M.2s are 4.5 times faster. So it's really-- and they aren't that expensive. So I put the web price, and so that's an discounted price, up there. And so you can see that they aren't that expensive. And and so if you need this performance, you should really go for it.

And the other thing that's been emerging is sketching. And I know Microsoft has great ads about this, with people sketching. But we have a Wacom technology, they have their own unique Microsoft technology. And so this is the same Wacom pen that you get on their big tablets, with 2048 points of pressure. So it's a real industry standard. And a lot of people are finding they can sketch things quicker, get ideas down quicker, show people quicker, and then they go to develop then into models.

So the other thing that's changed is USB 3.1. And you'll recognize that by the blue, that's the thing. It's 10 gigabits per second, and the lower b, the bits. And it's backward compatible. And there's devices out there that are coming out to interface with 3.0. So Samsung has one that, if you read the articles, the title of the article is, "This Is Stupid Fast". That's the title of the article. So if you do backups or transfer data.

And this is the benchmark chart. So you can see at the beginning of the benchmark chart-- so, first thing I had to do before I started this chart was I had to convert gigabits into gigabytes, or megabytes, so these aren't-- you have to make that conversion versus for specs.

And so 60 is the USB 2.0, which is probably what you have today. 245, that's the this USB stick drive. That's about the fastest that's out there. This Samsung drive is 450. And then 3.0 is 640. So you can see that these are available now in workstations. But the drive speeds are just starting to improve to really capture the maximum.

It used to be when we went from 1 to 2, it was the opposite. The drives were ahead, you've got 2.0, you plug in your drive, you've got instant speed. This is the reverse now from that. So and then we're actually at 3.1, which you can see-- it's much, much faster. And so it's going to be a while before people actually have drives that utilize this. But you need to think about that hand in hand if you use that kind of technology.

And then there's also Thunderbolt. Does anybody here use Thunderbolt? Yeah, there's a few people that use Thunderbolt. So Thunderbolt came out, and if you look at the chart, you can see USB 3.0, 3.1, then you see Thunderbolt 2, which has been out for a while. The new Thunderbolt 3, you can see how much faster Thunderbolt is.

And so over time, people think things are going to go to Thunderbolt, rather than USB, so this is the strategy. But you go from what they refer to as passive cables to active cables with electronics in them in order to gain these speeds. And so it's not going to be the same as just plugging things in and using any old cable anymore. But but that's coming.

And then the Wi-Fi. So here is where we are with Wi-Fi. So 802.11 ac is probably the one that you have already, that's the most common. It's the highest performance, and-- I'm sorry, n. 802.11 n. So ac which has just come out as of last year at this time. And so that was the performance of wave one.

Wave two is even faster. So you can see over here, it's 3-point-something gigabits per second. Your standard ethernet cable that you plug-in on your computer has been one gigabit per second for quite a while, unless you-- there are some options to go to 10, but you need to work that with the server, so that you can take advantage of it.

But so what this is telling you is that wireless is going to be faster than wired. Right today. And so you want to be thinking about your wireless connections, your routers, your speed, your

speed back to the computer room to those wireless routers, and see what you can get to really boost your performance. So if you do a lot of stuff over the net, now you can really think about how to get that speed up in better ways, and you don't have to rewire entire buildings to do it. You can do it wirelessly.

Another thing we did is we changed the idea of cooling and laptops. Now I know everybody's had this experience, taking on high performance laptops and trying to put it on their lap, and finding it's too hot. Or put it on a table and think it's going to burn a hole in the table, and things like that. And that's because these processors consume a lot of power.

And so those squares in the middle-- there's a bigger one and the smaller one-- so that's the CPU and the GPU. And then these circle things on the end are fans. And so what we did is we put in a copper bus, and just like pans have copper bottoms, because they distribute the heat equally, or if you've ever soldered a little plumbing on the side, you know if you heat up the pipe here, it's going to get hot a little further away. You can't hold it with your hand. Because it distributes the heat very rapidly.

And so we used these heat pipes and we actually coat the center section black now so that it doesn't give off any heat inside the system. Goes right to the fans, and the fans distribute it. So you can put your hand under ours now, and it's barely warm. It's amazing. So we like to say we put the lap back in laptop. But that the heat does come out in the side, you'll definitely feel it.

And we did the same thing a few years ago with our desktops. So we have front to back cooling. We're the only ones that do that. We don't recirculate the air. And so it keeps things very cool.

And another thing that's that we introduced we're surprised nobody else introduce this-- is diagnostics on your workstations. It's just like your car. To troubleshoot your car, used to be you had to take it to the mechanic. They'd have to figure out what's wrong, and then call you after they try a few things. Some things might be wrong.

Well now, we give you a code, and so all you do is report that code. We know what's wrong, we replace it, and it's that simple. And so there's no troubleshooting. So it's very easy. So I call this read and call diagnostics. You read it, you call up your IT person, or you call us up, they know what's wrong, they fix it.

So there's some other cool technologies we want to talk about, too. AR and VR. You're going to see a lot of that at the show here. We have ours, [INAUDIBLE]. HP has one out there. I think Intel has one. So there's people showing AR and VR. Autodesk has a huge section on it.

And these are evolving. So to show different ways of looking at designs, making design decisions, collaborating, looking at things at scale, as you can see in the upper right. So here's a guy looking at the design while they're building to make sure everything is correct. So they can overlay that with some AR.

And this design experience here, sitting in a real chair, but trying from a AR perspective, how handles look and feel, and so forth. Or a walk through. Or a customer's review. You can look at that on a-- create it and then look at it on a tablet, even. To see what it would look like.

So there's all these new techniques that are out there that are really going to change the way you do your jobs, which you create for people to use. And it's going to be very significant. So Dennis, do you want to add a few things on this?

DENNIS

Yeah. So in the last year and a half, virtual reality has really become mainstream. And a lot of that is because of the cost of the devices. So just curious, how many people are actually looking at integrating VR or AR head-mounted displays into their workflow?

SESSANNA:

So that's quite a few people. So again, I work with an automotive company, and and they've been using virtual reality for a long time. But a lot of times the devices were millions of dollars. So lots of very expensive technology. And all the automotive companies are looking at these low cost devices, and even to the point that they'd love to take an HTC Vive or an Oculus and putting it on every designer's desktop that they could pick up so.

So Nvidia is doing a lot of work to accelerate that technology. When you're using a head-mounted display, you have a left eye and a right eye that you have to render. So we've created a technology called VRSLI, that actually distributes that workload across two graphics processors. That's currently integrated into application from Autodesk called VRED, or VRED.

We're also working with the Stingray team to accelerate their VR renderings. And then of course we have really good relationships with the folks from the Create The Unity and the Unreal Engines. So they're integrating that technology.

The main thing that I will want to point out is that there's so many people that come up to me and they say, I have to use my G-Force GTX for VR. That's the only thing that that's

supported. And that is completely not true. So the P-6000, the P-5000 are actually faster than the GTX 1080 and the Titan X version.

So we're working with all these companies to make sure that our devices are certified. We have great relationships with the folks at Microsoft and HTC and the Oculus team to make sure everything works properly.

So I've been working with a woman, Elizabeth Baron, at Ford, and she's been doing head-mounted displays and VR for a long time. And they're jumping in with both feet with HTC Vives because of the full track space. So very cool technology. Again, I'll be in the Lenovo booth this week. So if you want to talk about VR or high-end rendering, make sure you come back and talk to me.

TOM SALOMONE: Yeah, thanks. Thanks, that's good. So I'm going to talk a little bit more about this too.

And so one thing is to create it. So using VR and AR to create it. When you're going to just look at it, it's even easier and lower cost. So here's an example of tablets and other ways to look at it. Some people still are going to want to use headsets, like maybe in a customer review, where you see this person walking in a blank room, but what he's looking at is this new design that somebody did.

And so there's all these different ways. For maintenance or building, I think there's huge possibilities. Training. It's amazing. And so we have some of that here today. And we're going to show that.

And then Nvidia video came up with this VR ready designation. And so we have recommended configurations that fit the Nvidia VR ready, either AR or VR type of thing. So you want to look at that.

So it could be our high end, where you get the best performance, or it could be our low end desktop, and still get good performance. So depending on your budget, you know you might want to try the low end. Or if you're just starting out or experimenting. That's a good way to start. And then when you realize what model sizes you have, you might want to scale up. Or what you're trying to really achieve with your customers and view.

And so this is our section for-- we are we have actually have five headsets. And one thing we're going to do-- and you should think about this as a way you can use VR too-- is we're

going to instruct you in here on which headsets, what are the limitations, what are the graphics cards, what are the limitations-- that type of thing-- through a VR experience.

So nobody is going to talk to you. You're just going to see that in the VR experience. And you're going to experience what VR is, what it can do, and how it can be used to either sell your products or review your products or modify them or collaborate. So is this is a pretty good.

And then there's another spot that's there. And this is The Future of Making Things from Autodesk. So you'll see this in the exhibit hall, this big red area. And they have 10 headsets, to look at different things. And they're all about manufacturing, so you're going to look at assembly and things like that in there. And so it's another place to look at some VR. And where we are located is the red spot.

Then another technology I was going to tell you about it today is Phab 2, is this device. And this device has just come on the market. You can just start buying this device. And it looks like a huge phone. But it's called a phablet, because it's a phone tablet, if you would. Outside the US, these are much more popular, because people don't buy tablets or computers, they just use this.

And so-- but this one has the capability of capturing point clouds. So you could sit in this room and you could capture the point cloud. And then with this company we have, DotProducts, you can convert it to a 3D image, and then you could take that image and run it through Autodesk ReCap, and you can develop it into a 3D file to use in Revit or AutoCAD.

So this is a new technology. This is \$499. So this is a new way of thinking about quickly capturing images that you're going to then use. You can also take the captured image and put on a VR headset and view it as well. So we'll have this, you can see this in our booth as well.

So this is the image, once you do it, and then you run it through the DotProducts. And so this is cabinetry, and you can see, he took this image, and he put it in. And then he took the image of the top cabinets, and he connected them. And so that's what that is-- the lower cabinets and the top cabinets. And the other one.

So the whole thing I was trying to get across is, this is really a dawn of a new era. So beyond processors and graphics, but all these other things that are going on. That are changing the way we think about design. And it is going to have a very significant impact on design and

people's can competitive advantage.

And so I just have these. If people want, we can go over these and in our booth. But we can have very specific examples of what you might want to think about buying, in terms of our 15 inch and our six-- and the configurations we would recommend. And then we have our 17 inch as well. So this is another example. And so the 17 inch is really good for those people that want to do rendering using GPUs.

And then our Yoga. We have that down. So if you like to draw, definitely stop down and give us a-- do some drawing.

And then lastly, we're giving one of these away in our booth. So you might win. And win big. And so I would recommend you stop and take a look. And with that, we'll look for questions. Yeah.

AUDIENCE:

Do either of you have any experience with laser scanning, or products, so single-core processors with Xeon, or which level graphics card would be useful? Or large laser scanning data sets?

Yeah, so the lasers scanners that you use are developing point clouds. So point clouds can be small or large. But in terms of your question about processors, so if they're just the dataset, you're not using any cores. You're just using single cores for that. But if you add it into CAD, which is probably your goal, then single cores is still going to be best. So that's what I would recommend. Because I don't know of any parallel processing for the large data sets that exist. So if you're doing rendering, then that would-- that lends itself to parallel process.

DENNIS

SESSANNA:

Yeah, to me, definitely, it depends on your models. You How big your point cloud is. so we were showing a little demonstration application at Siggraph this past summer that was a visualization of a huge point cloud of Nvidia's new headquarters that we're building in Santa Clara. So I think it was around on the order of like 200 million points that we were able to visualize in an HTC Vive in real time.

So that was using too high end graphics cards and special custom Nvidia software that hasn't been released publicly. I know there's other applications out there, but it is-- especially if you're talking about in virtual reality-- you need those 90 frames a second frame rates. So if there have big data sets, I recommend the highest end that we have.

TOM SALOMONE: And we can show you data sets in our lab or in our VR area, depending whether you're

interested in VR or desktop or mobile. Yep.

AUDIENCE: In the beginning of your presentation, you talked about all these applications that you're doing, and then it looked like there was a massive [INAUDIBLE] that you're doing.

TOM SALOMONE: Yeah.

AUDIENCE: Are you doing that independently of the Autodesk certified hardware page? Do you have a location that you're storing information?

TOM SALOMONE: Yeah, so what's happened with Autodesk certification. That's a very good question. So over the last couple years, Autodesk has taken the certification process and pushed it out to OEMs. So we do a lot of testing. We submit the results to Autodesk, they bless it, and put it up on their web.

And so she said, well what about all that massive data? So some-- it's possible to look at some things and say, if I do this one test that really covers these configurations, so we don't have to test everything because that would be really impossible. An impossible task. So we look at what's the same processor, but maybe a different speed, or something. And we look at that stuff and we say, OK, well that's the same as this, so we can we can certify it.

So thank you for that. So we do we do our own testing. We then submit it to Autodesk. and. They audit us, and they audit others. But that's how it works.

AUDIENCE: So you don't have a independent location that you're storing [INAUDIBLE]?

TOM SALOMONE: Say what?

AUDIENCE: You don't have an independent location that you're publishing that information? Like a place we can go--

TOM SALOMONE: We have a website where we publish our certification lists on our web page, yes.

So, another.

AUDIENCE: Yeah, you mentioned that you were doing eye-to-eye with GPUs left-right. That's how SLI [INAUDIBLE] set up, or--

DENNIS: Yeah, so a lot of these new display devices are-- there's a single output. So you have one

SESSANNA: cable. So that is an SLI setting. So the ability to, as the data goes over to the display-- it's a single display but you need a left eye and a right eye-- so it is an SLI set up. Again, there's not too many applications that currently support that technology. So it has to be integrated into the application. It's not just a switch on the Nvidia control panel that enables that.

AUDIENCE: Are you going out of your GPU into your monitor, and out of your monitor into the [INAUDIBLE], or--

DENNIS So there's a technology, I believe it's called direct connect, so it's the ability-- so when you plug
SESSANNA: in HTC Vive, even though it is a display, it doesn't appear in the Nvidia control panel. So we've worked with them to make sure that you don't have to take your VR application and drag it over to your head-mounted display device. In the beginning you had to do that.

Now we have this thing called direct connect that directly connects to the GPU. It doesn't look like a display device to the operating system. HTC and Oculus and Microsoft all support that technology. Makes it easier for users to be able to use these devices, and just kind of plug and play. It all goes to the-- there's a framebuffer in the GPU, it renders to that framebuffer, and then it just scans out directly to the HTC Vive, and not to your regular display.

TOM SALOMONE: OK, so we we've run out of time here. Again,

DENNIS We can we can talk about that--

SESSANNA:

TOM SALOMONE: Yeah, so anybody that has any other questions, just come on forward, we'll happy to answer them. And thank you very much.