Intel IDF 9-20-07 Morning Session

[Beginning of recorded material]

[Music and video plays.]

Dadi Perlmutter: Good morning everyone, and welcome to the third day of Intel Developer Forum. When I see a large crowd on the third day, I have three options. By the way, everything I have is three options. It's either you've had a terrible two days, and you say [it is] the third day and I have to make it. Or you have nothing else to do. Or the option I prefer -- by the way, I'm not going to ask what is it really -- that you really had a great two days, and you just hope that the third is going to be at least as good, if not better. I think it will be.

Just to recap what we talked about -- or the highlights of today. I forgot to tell you. Could you rewind the point, please, for the previous one? Okay. A small recap of yesterday. You did hear about our vision for mobility, really making Internet on the go using cool and getting even cooler and better solutions from all the way from notebooks to mobile Internet devices using great CPUs technology, wonderful graphics, and enabling wireless and broadband wireless across the board. You did hear later from Knut about what could be done with nonvolatile memory, way beyond the way nonvolatile memories are used today, going from disk caches all the way to solid state disks to really close the gap and solve the I/O problem, making sure that this part of the system gets way faster. I bet, as I used to be a technical person in the past, that the discoveries of the fellows was way better and way more entertaining than a discussion or hearing the gritty presentation from the executives on stage, and you had a chance to really talk to them, ask tough questions, get through the details, get to the guts of the architecture, the solutions, the everything from our fellows. You got a wonderful opportunity to get training on gaming and showcase and ask about these, the social interaction, the great stuff you are doing.

So what's going to happen today? Immediately after me, Renee James is going to talk about the thing that really interests you, the developers, how she, with the Intel software team, is going to take and work with the developers to really enable all the great technology that we have showed you in the past two days. Then, it's going to be the wave of the future. Justin Rattner, Intel CTO, is going to talk about life, or the virtual life. By the way, people who have the virtual life are going to be way better than the physical life. I'm not so sure, but he's going to show how wonderful that's going to be. And by the way, always the future is better than the present, so we'll wait for him and get that stuff. And then you have the last chance to get the training and the discussions and talk and meet the fellows to the last questions you had in mind and you didn't want to ask.

As I believe you really enjoyed this IDF, those who want to hear it again could fly to Taiwan. It's going to be on October 15 to 16. And then like any good show, it's going to be the sequel. Of course, next year in spring, we are going to bring the next messages, the next wave of products, next wave of technologies. We would love to work with you, the developers, going on to the future.

And a few words about the flowers that we had outside. These flowers are really for our celebrating 10 years of IDF, but I think more than that. They are very much symbolizing blossom, growth for the next wonderful 10 years together in IDF, but more than that, in real work on bringing great technology to the market for the success of all of us. Thank you very much.

[Applause]

[Video]

Renee James: Good morning. Good morning, okay. This morning I first of all want to thank you for following all the signs to day three of IDF. We have a pretty exciting morning for you, starting off with, before I get going, at the end of my presentation, we are going to announce the winner, the lucky winner and 18 of your friends, to a Gwen Stefani concert. And like good IDFers, we're using RFID technology so we'll know if you're in the room, and you need to be in the room to win. So that's exciting incentive number one. Second incentive, we have a lot of new developer information to share with you this morning, and then Justin is going to pick up from there with a fantastic presentation about the future.

So I'm going to switch gears from what you've heard the last couple of days, and I know it will disappoint you that I'm not going to show you more roadmaps and I'm not going to share future compiler projections. And, instead, I'm going to talk about something that Intel has been doing over the course of the last 18 months to revamp our developer and community efforts.

So you heard from Paul and from Dadi and Anand and Pat about all of our technology initiatives. And, today, I'm going to discuss how you, as developers, whether it be software or hardware or in between can get engaged with these different platform initiatives. We've changed, in the last year and a half, significantly our approach, our philosophy, and the tools and contributions that we're making available to the developer community. And I'm going to go through that in some detail.

But before I get into the changes, I want to tell you the one thing that hasn't changed. And that's our strategy, our objective, our goal for software at Intel. It's always been the same. Port of choice. And what that means to us is that all the software across platform[s] in all the different communities or different usages runs best on Intel architecture, period, full stop, it's as simple as that. My entire team, many of whom I can see in here today, who've been doing sessions for you in the last couple of days, is dedicated to making that come true. So we'll talk about how we're going to do that and how we'd like to you to get engaged with us as we go forward.

So the first thing we have to talk about is what's changed, so what caused Intel to consider that we might want to take a different approach. And what we looked at, and I think the best proxy that we could use for what's going on in the developer community, and this is individual developers as well as companies, is where are developers coming from in the future? So what I'm showing behind me is IDC projection data for 2011 of the growth rates of developers worldwide. And I don't think that I have to call out the data for you to notice that more developers -- and the growth is faster -- the growth correlates to number of graduates from computer science programs worldwide. And you can see where they're coming from. It isn't large companies in mature markets, necessarily, it's from new growth markets.

In addition to where these developers are coming from, there's a change in how they're developing. We have seen and undergone one of the biggest changes, I think, in approach to a more collaborative-based development design and iterative approach. So all these new developers are being trained and coming into the community, if you will, with a new approach. Development is no longer the purview of one company with one idea. It's a community of people collaborating.

So as we go forward, I want to kind of compare and contrast what Intel used to do with where we're going. In the past, I don't think I need to tell a lot of the IDFers this, but we used to design the hardware and put out a spec and then we'd work with you through sessions and/or applications engineering to get either your hardware or software or what have you, the system design to work with it. And it was very much Intel information on Intel's time. What we've seen is that that's not how the community or how developers want to really work with technology in the future, or work with Intel, for that matter, and that's the feedback that we've really heard from you. People want to work in communities.

And for the purposes of our presentation for the rest of the morning, I'm not going to define communities; I want to tell you what we think it is. It's a group of passionate people committed to innovation and breakthrough, which is why we're all here, equally motivated not only by community and collaboration but by competition. So what we have moved to, and what I'll talk about for the rest of the morning in specific detail, is an open embracement of a collaborative approach in development. Our developer programs, our new initiatives are all very much embracing the concept of how you can get more engaged with Intel much earlier than ever before.

We've revamped our entire approach. We have developed specific communities of interest, which I'm going to detail out, around each one of the technologies you've heard about from Pat and Dadi and Anand. And most importantly is the philosophical change in approach. We are contributing intellectual property, our tools, our people, our ideas, and sample code to really participating fully with developers going forward. And we think what this signifies is that Intel has moved to Internet time and off of Intel time with development.

So we revamped our site. You can get there from intel.com or inteldeveloper.com or intelsoftwarecommunity.com. Any way you want, they all link together. It is an interactive site that's moved from a repository of information that you may or may not have ever used to an interactive site with over 500 blogs, lists, developer competitions -- one of which I'm going to announce the winner of later, another upcoming treat -- contributions are listed there, projects. We actually have developer and expert sessions that we host in Second Life. You'll see some of that in Justin's presentation later. And the response has been overwhelming. We did this less than a year ago, and we already have 65 percent more developers participating with Intel than we ever have had in the past.

So today I'm very excited about the next step that we're taking. We are announcing a new addition, which is called whatif.intel.com. In an unprecedented move in the history of Intel, we are putting up what is very much like Google Labs or an Alphaworks site where we will share our early ideas, our alpha code, thoughts, lab projects with the development community, get your feedback, allow you to interact with us on where we're going. We've posted three of our first projects up there, so we're ready to go. One of them is some thoughts we have about our performance tools that wouldn't even come out until 2009, but we've put it up there and are wanting to get feedback on that. And we look forward to doing more of that. And most importantly, what we're hoping with this site is that we can get more engagement with you and hear more from you about what you'd like to be seeing from Intel in the area of developer support. I can see you way back there [unintelligible] [Silvers].

All right. So one community that we've been engaged with for a long time, and this is not new, is the research community. And some of you are hear today. I understand that we have several universities that were invited. That's not new news. We have fantastic university relations and will continue to go forward on all of that. What is new in the last year and a half is that we've taken an additional step and focused our efforts on curriculum development. Specifically, I’m going to use the example of parallel programming.

One of the most profound changes coming to developers, and you saw all our roadmaps for the last couple of days, so I'm going to come back to this point because this is key to how we're going to use all of those cores in the future. We developed curriculum, specifically undergraduate curriculum, and we trained professors in over 400 universities to teach parallel programming. This is something that we're spreading; it's worldwide from Shanghai, [Jetun] to, you know, U.S. universities. We hope by the end of the year to have a thousand universities joined up and teaching parallel programming. In addition to that, we have trainers that we'll send out, and we're working with training companies to really forward advance information about where Intel's roadmap is going and how developers can be ready and trained to interact with the technology that we're building.

So, another community that is critical to how we develop and how software will be developed for Intel platforms going forward is the open source community. And Intel has become more of a key contributor to the open source community in the last couple of years than we have been in the past. Part of this is the change to being more open and having information more readily accessible. This behind me is a very small subset of the contributions and projects that we're engaged with across Intel. My group is not the only one. All different groups across Intel contribute and participate in open source. We think that this is an important piece of innovation and how the community can bring new uses to the platform.

We've worked really hard in the last couple of months -- I'll say 18 months if you will -- and couple of years. Going forward, this is part of an ongoing process; we'll be continuing to add to it. And I think, in the words of Alan Cox, "We can see some improvement." In 2001 Alan suggested that Intel was secretive and hard to work with. Today, he would say that our cooperation with the community has not only paid off to making us the platform of choice, but also, we're much easier, much more open and collaborative in our activity.

So that's it on kind of what we've done to revamp the overall program, and now what I'd like to do is talk about the technology initiatives you've heard about and how you can find more developer information about those.
So specifically, around mobile power, power in the data center and virtualization, as well as how we're going to use parallel programming and how that relates to visual computing, I'm going to go through some of our very specific initiatives.

Yesterday, you heard from Anand about MIDs, mobile Internet devices. It's a new category. We're very excited about the opportunity here. One of the things that Anand talked about was the low power needs and the memory footprint needs and the requirements of full Internet resources. And he demoed technology based on a community that we launched in July called Moblin.org. And I wanted to bring the community piece of it back up, one, to clarify that this is not a company. This is a community of interest that anyone can participate in around what we're doing on mobile Internet devices. We donated the initial code and some of the samples and got the community going, but as you heard yesterday, Canonical and Red Flag and others are already engaged, and we expect more engagement from other companies going forward. And we think this is super exciting from the perspective of a new approach, from the beginning, ground up, new category of products that Intel is starting, and a community working on it.

So specifically, the four areas of interest in Moblin.org are, of course, information and technology sharing. That goes without saying. Most interesting, integration and interoperability work, which is critical in these devices, tools, contributions, which Intel will continue to make and others. And most importantly, the power management and memory work. So I encourage you, even if you're not interested, to just go check it out and get involved and see if this is something, if you're interested in this, you want to get part of the community and contribute and hear where we're going.

Okay. Communities have a way of spawning other communities. And the team that's been working on Moblin.org got very excited about the power and performance breakthrough technology that they started to work on. And they really, they came back and said, "Hey, we could do something in the data center very similarly." And out of that grew another idea.

You heard from Paul and from Pat that energy consumption is a huge issue and using U.S. data only from the EPA, data centers and servers are driving a significant increase in the use of energy consumption in the U.S. -- 2X in the last six years -- and represents 1.5 percent of the total energy consumption in the United States today. This is a major environmental issue. We kicked off an initiative about it, and you heard about it earlier this week. It extends globally, of course. So one of the things that the developer team has been looking at is, how could we create a community around these environmental issues, and specifically, to data center power management?

And what I'd like to introduce is the newest community today called lesswatts.org. LessWatts is, under the Climate Savers umbrella, an organization specifically focused on data center, power management around Linux operating systems in the data center. And to help me talk about that, I'd like to introduce Wim Coekaerts from Oracle to come out and talk to us about his work. There he is. Good morning.

Wim Coekaerts: Good morning. Good morning, Renee. Good morning, everyone.

Renee Jones: Come on down. Wim's the vice president of Linux engineering and has been involved with the team on starting up lesswatts.org, and why don't you tell us about what your guys' objective is, and the initiative itself?

Wim Coekaerts: The initiative is actually very exciting for Oracle. There's two reasons for this. So first of all, Oracle is a huge user of Linux in-house. We have thousands, if not tens of thousands, of systems out there, and so we can do some energy saving ourselves. And then secondly, for our customers.

Renee Jones: Right.

Wim Coekaerts: So, let's start with internal usage -- so within Oracle, Linux is our base development platform. So the QA farm, which is thousands of machines, run Oracle regression tests day in, day out, 24x7. And then we have the developer workstations. We clearly have many thousands of developers within the company all over the world. With Linux being our base platform -- its X86, X86-64 platforms -- so it's very important for us to be able to save some money in that space as well, right?

So what we're trying to do together with you on this program is basically we tune our data centers. We make sure that we have the right measurements in place. And then what we try to do is we run Oracle products and we see what the power consumption is. We try to tune the operating system to see where we are using power where we shouldn't be, we're running services where we shouldn't be running them, and so try to come up with ways that we can save energy ourselves within the company.

Renee James: Within your own data centers?

Wim Coekaerts: Within our own data centers. Right. So secondly there are the customers. Oracle traditionally has been working with the DBAs, and DBAs don't necessarily care about power savings. They just care about the database being up and doing their piece.

Renee James: Right.

Wim Coekaerts: But for many years we've been doing Linux OS support -- originally, informally for many years, and then last year we announced the

Unbreakable Linux program. And that sort of gave us a way to actually also talk to system administrators.

So now we have contact from an Oracle point of view directly to the OS people, to the system engineers, to the designers, all the way up to the Oracle [product stack]. So we get a lot of feedback on that. And so one of the things is of course talk about consolidation/power savings.

Renee James: Right.

Wim Coekaerts: Consolidation in many cases is about power savings because they're saying all these systems are doing nothing, but they're using a lot of energy.

Renee James: So you'll be contributing to this organization much of that learning out of Oracle.

Wim Coekaerts: Yes. So we have this program called Oracle Validated Configurations. It's a Linux program where we basically give customers feedback on -- you know, you buy this system from a hardware vendor, you install enterprise Linux and you install the Oracle product stack on top, and this is how [we tune it]. So what we're trying to do now, working with you, is to extend that to say we have validated configuration programs that are also tuned for power savings.

Renee James: Okay.

Wim Coekaerts: So all this stuff we're finding out internally on fine-tuning the OS -- there are the user space changes we have to make; there are the kernel changes we're doing together with [unintelligible]'s team and other folks from Intel and the community -- and so all that together will get a nice bundle that we can offer to the customers.

Renee James: And this is all going to be in the organization lesswatts.org? Will be contributed?

Wim Coekaerts: We'll be contributing all that stuff back. Because everything we do on the OS side -- just like with you, and so forth.

Renee James: Perfect. Anything else you want to comment about how this might be implemented at Oracle or implications for your own data centers?

Wim Coekaerts: Well, you know, within [our own] data center, we use our own best practices. So the validated configuration info is what we also deploy. So the things we're asking our customers -- or we recommend our customers to do -- we're going to recommend our own data center teams to do as well.

Renee James: So the best-known methods of your data centers will also be found in the LessWatts. Great. Thank you.

Wim Coekaerts: Thank you.

Renee James: Data center power management is not the only issue, and Wim started to talk a little bit about consolidation. I should mention -- I'm sorry --

that also Red Hat will be part of this. They weren't able to be here today, but Paul Cormier wanted to make sure that I let everyone know that they would also be supporters of this initiative and very excited about where we're going here.

So transitioning to and staying in the data centers, I want to transition to talking a little bit about virtualization. Virtualization is a key technology. You heard from Pat about what we're doing with the community on virtualization. Our strategy is very similar, which is to be the port of choice and work with all of the different virtualization vendors. He showed a variety of different platforms in his presentation, but the one that he didn't show that is relevant to both data center power management as well as virtualization, is the Microsoft platform. And I think it goes without saying to this audience that that is one of the largest developer communities in the world. Around Windows. And we've had a very long and productive history, greater than 20 years, of working with Microsoft to bring innovations to market, both hardware and software.
So today I would like to have a special presentation from Microsoft about where they're going in power management with Windows Server 2008 and show us their new virtualization technology, codenamed Viridian. And to demonstrate this, I'm very excited to introduce the general manager of virtualization strategy, Mike Neil, who we've worked with for a long time, and Mike's going to finally show us the fruit of all of our efforts.

Michael Neil: Renee, great to see you.

Renee James: Nice to see you, Mike.

Michael Neil: Thank you. It's a genuine pleasure to be here. Thank you.

Renee James: Yeah, I'm happy to see this product that's going to almost launch.

Michael Neil: It's great. Let me take you over and show you what we're doing.

Renee James: Okay, super.

Michael Neil: So this is a machine running Windows Server 2008.

Renee James: First we should say "server demo."

Michael Neil: Yes, exactly. So this will be the sexy and exciting server demo that we all worked on. So, this is a system running Windows Server 2008. This is the most robust and flexible solution that Microsoft's had in this space. And I wanted to kind of walk you through and show you a little bit of the virtualization capabilities that our teams have been working on.

Renee James: Why don't you, for the audience -- I don't know if my microphone is on or not -- can you tell us when we would be able to see some of this as you go through it?

Michael Neil: Yeah.

Renee James: When the audience would be able to interact with the technology?

Michael Neil: Yeah, and that's exciting. We're actually --

Renee James: No futures here.

Michael Neil: No, no, this is -- well, it's just slightly a future. So we expect to release the release candidate next week, and it will contain the Viridian technology so that everybody here will be able to download it and try this out. So we're really excited to have people take a look.

Renee James: Great, that's great news.

Michael Neil: So, one of the things that we did with our virtualization technology is really integrate it into the Windows experience. And so we wanted to make it very easy for customers to be able to do that. So if I bring up the server manager, I can see we have various roles that we can add to a server. And so one of those roles is our virtualization role that we see here at the bottom, and that one's already installed in this system. And you'll notice that this is a system running our hypervisor that we worked with you on. It uses the Intel VT technology to be able to power that hypervisor built on top of that.

Renee James: Thank you, great.

Michael Neil: But the machine looks just like any other Windows systems. There's no real difference for the customer. They really just interact with a system like they would any other Windows box.

So, that's the server manager. And let me just show you our new MMC user interface for our virtualization environment. This is our management console. We have five virtual machines running, and I'm just going to show you a couple of those and talk about some of those capabilities as well.
One of the things that I should point out is that through our virtualization technology, we're really building it in as part of our core platform. And you talked about developers and how important they are to the ecosystem. And so not only is Viridian a technology that people can use to do consolidation and that customers can interact with, but everything you can do through our user interface is done through a set of APIs that we've jointly developed with the industry as part of the DMTF. And working with Intel and others in the industry, we've produced a common set of management APIs so that others can build on top of our platform and create very exciting solutions.

Renee James: For those in the audience that don't know about what you did with the DMTF, can you just take one second from the demo and explain? Because I think it's pretty revolutionary, especially as virtualization is kind of maturing into the market.

Michael Neil: Yeah, absolutely. It's important that, from a management perspective, that you have management solutions that can span the virtual environment as well as the physical environments. And so we really worked with the industry, including competitors like Gamware, partners like Zen Source, Intel, IBM, HP, and others in the, and Dell, in the DMTF to create a common set of management APIs that the industry can adopt and use throughout all of the products out there. So there's quite a bit of traction, but we really built that in our initial offering of our virtualization technology.

Renee James: Into the operating system?

Michael Neil: Correct, yeah.

Renee James: Okay, thank you.

Michael Neil: So this is a virtual machine running Windows Server 2008, which we're also excited about, of course. And one of the things that's exciting here is this is a quad processor virtual machine, so able to take advantage of the Xeon quad processor technology right out of the box, which is good to see.

Renee James: Good job.

Michael Neil: One other virtual machine I'd like to show is our server core environment. And this one isn't as exciting, because it's just a command line environment from that perspective.

Renee James: It's almost like a trick question nowadays. Can you program a command line?

Michael Neil: Yeah. So, this is really an environment that we've developed as part of Windows Server 2008 to be a lot more flexible and reduce the footprints associated with deploying certain roles.

Renee James: Footprint as in power and memory?

Michael Neil: Power, memory, disk footprint, all of the above. And servicing gets easier because there's fewer components on there. You're not going to update this machine because of an update to Windows IE or something like that. So the reduced footprint really provides a lot more flexibility for customers to be able to deploy more roles and to have more systems consolidated as well.

And then the last one I wanted to show was some work that we've been doing with the broader industry. And so this is Novell's SLES running on Viridian. And we've been working with Novell to optimize Linux on our platform. We realize that our customers --

Renee James: I never thought I'd hear that statement.

Michael Neil: I know. It's a little strange, isn't it?

Renee James: Yes, but it goes to our point about community.

Michael Neil: Absolutely, and it's important because our customers want to consolidate their heterogeneous environments.

Renee James: They do.

Michael Neil: And they want to be able to run Linux as well as Windows side by side in their virtual environment. So we've been working with Novell on this, we've also been working with our partners at Zen Source and Citrix on the integration of technology, so that we can have even a better experience with folks running Linux on top of our platform as well.

Renee James: Fantastic.

Michael Neil: So just to end on, I think one of the things that we're really excited about, as we consolidate these systems, and then with Windows Server 2008, we also have power management turned on by default, the millions of systems that we'll support between Intel and Microsoft that are out there, we're going to dramatically reduce the power consumptions across the world. And that's really exciting to see as well, so . . .

Renee James: Thank you.

Michael Neil: Good.

Renee James: That's wonderful.

Michael Neil: Great.

Renee James: And I look forward to this coming out next week.

Michael Neil: Yeah. Keep an eye out for it.

Renee James: We will.

Michael Neil: Great. Thank you very much.

Renee James: All right, thanks, Mike. I know server demos are always rough, but I thought since it's coming out next week, much awaited for, and Mike has been working very hard on the cross-platform elements to ensure that we can manage in a heterogeneous environment and reaching out and working with Zen as well as the Linux crowd. I think that's an important piece of a change in the way that everybody is working today. So we wish Mike luck on his upcoming product launch, and we look forward to the alpha next week.

All right. Now I want to talk about something that's kind of newer and dear to my team's heart, and that's parallelism. Every roadmap that you saw, that I won't repeat, was about the cores and what we're adding to our roadmap over time. And the available cores for software performance, for application utilization. And I've talked about the university program, but what I'd like to talk about now is the very specific developer program around parallel computing and what we're doing.

First of all, we, of course, have our software community site, which has a wealth of information -- tools, sample code. We actually run live Webinars where you can come and [get] free training. As I said, we also host events in -- Second Life developer events. There is a knowledge base. There are live application engineers who can blog with you and answer your questions. You can actually even get a call back from an application engineer if you're having trouble threading your code.

Much of what you're going to see on this site has come from developers like yourself giving us feedback on what you need to be able to take advantage of the roadmap that you've seen from Intel. But, in addition to the site, it's not a roadmap, but I have worked in the Intel -- my product line, the tools product line. We do have an amazing set of what I think are leading tools in the area of threading. And you know us for our compilers and our tuners and VTune, but you may not know about the thread analyzing tools or the thread building blocks or some of the cluster tools, and even some of the performance libraries for threaded code that we have available.

So I think it's worth taking a look if you're struggling, if you're even thinking about threading and how you might get engaged with it. Intel has a full set of tools. As I said, online resources as well as live resources. We are pulling out all the stops from a developer perspective to really get people engaged and moving. Quad-Core is not, you know, our biggest challenge, right? You know, we have eight cores coming and beyond, and that's something that is pretty important work for us to start now.

In the sense that we think this is so monumental, and I told you we're training developers today, we have, in an unprecedented move, decided to open source one of the fundamental building blocks of threading, which is the threading library. We announced about six weeks ago that we open sourced threading building blocks 2.0. This is cross-platform library; it works across Windows, Linux, Mac OS, Solaris, what have you. It has complete support. There's lots of documentation books about how to use it. There's a community around it. But we decided that it was more important to move the development community forward on Intel's roadmap around threading than for us to keep it to ourselves.

And I have to say that in all the years I've worked in software at Intel, almost 20, I didn't think we'd ever open source one of our tools. So this is, as I said, a profound shift in our thinking about how to get people moving and working and engaging with us bi-directionally around Intel platforms. So we ran a contest, because we love to run contests now, that's part of our community activity, for the first person to thread something using the Open Source, what we call TBB, or the Threaded Building Blocks, and I'm happy to announce today that Vincent Tan, a C++ programmer from Sydney, Australia, was the winner. He got a Dual-Core laptop so he can code mobilely anywhere in the world. And what he threaded was a Source Forge project called [P-Archive], although there's been much debate on my team about the pronunciation of P-Archive -- P-Archive or Parchive, depending on who I believe. And he was the first one to turn in his results. So thanks Vincent for that and we'll continue to run developer contests around threading. We have gaming contests going on for threaded gaming. We have a lot of, I would say exciting, maybe not as exciting as Gwen Stephani, but contests going on for developers.

So let me show you what we've been doing in the ecosystem. We've been working for a while on threading. I mean this is not new news right? So this is a snapshot of ISVs who already are taking advantage of Intel Quad-Core through their threading efforts in their commercial products. And it's a short list. We have many, many, many, many more. I wanted to just give you a snapshot of what's already out there.

And one of the areas that you'll see on the -- I guess your right hand side -- is games. Of course, visual computing is one of the areas that takes direct advantage of parallelism, and without going into a lot of detail on that, because Justin is going to wow you with it, I would like to just take a minute to show you one of the things that we've been working on. So Paul showed you on Tuesday the Pandemic game. That was a quad-core game; it had been threaded to use all four cores. It was running rendering on one, physics on one, game play, animation. And it was, I hear by all accounts, a pretty good demo.

But what my team has held out here is an eight-core game. We haven't shown it before. It's running on a Skull Trail motherboard, which is an Intel to-be-released motherboard, with two core 2.0 Quad_Core processors, four discreet graphics cards, eight gigs of memory. They're going to show you our first ever eight-core gaming demo.

[Video demo]

You can see the performance monitor across the bottom. It's going to run for a second. So we worked with the team at Capcom, who developed this game, to take advantage of our multi-core processors. And what they used is really what I just told you about. They used VTune, the thread profiler, the thread checker, the building blocks to allow them to find, detect parallelism, find out where are the right places in their code that would even benefit from parallelism were, and go ahead and thread the application. It's a nice job. Thank you.

So, gaming is a lot of fun. There's a lot more to it than just great graphics. There are significant other pieces, which I won't go into too much, but I do want to say this game uses the physics engine from the company that we acquired last week. And I'm happy to announce this morning the deal closed and Havok is a wholly-owned subsidiary of Intel. They are the world leader in gaming physics, animation, and rendering behavior. And we're very excited about that. It's part of a bigger visual computing strategy that we have, and you're going to hear from us, from our developer team, from my team, more about SDKs around visual computing tools in the next year, so stay tuned for that.

I'm going to leave it at that because Justin's going to pick right up on this thought around parallelism and visual computing. And I’m going to wrap up by talking a little bit about where we're going from here.

So I said that communities are key. It is becoming deeply engrained into my team certainly, and we're spreading it to the rest of Intel. There's no turning back. This is a major change for us philosophically. We really are wanting to encourage you, and I can't say this enough, to really engage with us in a different way. Become part of our developer communities, give us feedback, tell us where we're not hitting the mark or giving you what you need as developers around all of this technology that you've heard and are learning about in these sessions in the last couple of days. That's why we're here. I encourage you to go to the Inteldeveloper.com or Intelsoftwarecommunity.com, and we look forward to hearing from you.

So before I close and I turn it over to Justin, we do have to announce the winner of the drawing for the Gwen Stephani concert. So by attending the keynote and having scanned in, you're eligible for a luxury suite with 18 of your friends on Friday, November 2, at the Oakland Oracle Arena. I didn't know it was named Oracle Arena now. So they're going to tell me -- see, I don't even know who the winner is, so they're going to flash the name up there, I hope.

[Music.]

Renee James: [Unintelligible] are you here? Congratulations. Congratulations. I hope you like Gwen Stefani. Congratulations. All right. Thank you. Intel will contact you with all of the details after the show. All right, I want to thank Wim, and I want to certainly thank Mike for both joining me here this morning. And I want to thank all of you for sticking it out and following the signs to the third day. We are going to work as hard as ever to contribute to the community and help your experience with Intel as developers be better than ever. We look forward to working with you, and together, we're going to create the future. Thank you.

[Virtual presentation by Justin Rattner.]

Justin Rattner: Hey, everybody. Great to be here and I'm really pleased to see the house full this morning. That's wonderful. I hope I was able to stimulate your attendance and those of you who were thinking of leaving early have stayed put. So we'll try to have some fun. I really appreciate the fact that Renee finished a little early because I'm going to need all the time I can get. We've got lots of videos and some interesting demos and things like that.

Anyway, first of all, welcome to the research and technology keynote. I've kind of gotten out of practice with these things. I've been doing opening keynotes. I did the opening in Beijing, and that's kind of a different thing, more of a corporate thing, but I really like talking about the technology, so it's great to be back doing the research and technology keynote. I have to say it was quite a trip in from Second Life, all that rezzing and derezzing and everything, all my atoms are scrambled up now, you know? I don't know how those Star Trek people did that, but it's quite an experience.

This morning I want to talk about virtual worlds and what they may mean to the future of the Internet. You know, in many ways, this feels like 1993, those of you who were born in 1993. These IDF audiences get younger and younger and I feel older and older. You know, at the time, everyone was talking about the online services. I had an AOL account and a CompuServe account and that was really cool. My best friend was on Prodigy. I never understood that, you know? Prodigy seemed like, you know, kind of TinkerToy computing. But he thought it was cool -- Ph.D. out of Stanford, no less.
As a DARPA PI, of course, I had access to the Internet and, you know, was doing email. And the IT folks at Intel said, "Email? We're not sure, especially email outside of the company. Woo, what a concept."

So, you know, we were hearing about the worldwide Web, and we were downloading the NCSA browser and things like that. But, you know, we didn't really feel the technology was anywhere close to mainstream. We sort of looked at it and went, "Oh, well, that's pretty interesting. Maybe the scientific community will figure out what to do with this stuff sometime." But few of us, literally I think very few of us, imagined that in just a few short years the Web, the 2D Web as we know it today, would be a global phenomenon.

But this is not 1993, this is 2007, and the world is a different place. Not only has Moore's Law been at work, as we heard on Tuesday as Gordon described the evolution of Moore's Law, but other laws have been charging ahead at the same time. And, you know, we've graphed just a few of them here. And some of these laws didn't even exist in 1993. I mean, you know, there weren't wireless networks, at least ones that people, consumers could get to. So we didn't have any laws predicting the growth of wireless network.

You know, most importantly, though, I think, in the time since 1993, our use of the Web has really changed the way we think about interacting as human beings. I think everybody can agree with that. The Web has just made the world a different place. So the question is, the question I'm trying to pose this morning is are we at a similar point with respect to the 3D Web that we were in 1993 with respect to the 2D Web? Is all of the activity we see going on, and we'll talk more about it here as we go along, just a precursor to another one of these huge leaps?

If we are at the threshold of the 3D Web, then some of the biggest market opportunities are just ahead of us. That's true, as Steve Kleynhans of Gartner, he's a VP of research at Gartner -- in fact, Steve may be, I thought I saw him earlier this week. Steve, are you out here? I don't know. Raise your hand if you're out here. I don't see him. Maybe he headed back. He gave me the quote and, you know, I got approval, so we're good to go on that. But, you know, he made this observation that the biggest market opportunities happen when we change the way people interact with computers. That's a pretty simple idea, but quite true, and, I think, one of the factors that makes this notion of moving from the 2D Web to the 3D Web, or, if you prefer, 3D Internet, so exciting and so full of hope and potential.

So the question is, is another revolution about to happen? And let's look at some of the evidence. No, this is not a legal case here this morning, but let's look at some of the evidence. There is certainly a tremendous amount of activity going on. I don't think anyone can doubt that there's tremendous experimentation taking place on the Net today. We see this huge growth in digital communities. And to sort of get a better feeling for this, a little better understanding of what's going on, we've broken it up into four quadrants here.

There's the creative sector where people can contribute ideas, thoughts, you know, express themselves, things like Wikipedia and, of course, the blogs, and express themselves in other ways through content creation, certainly YouTube and Flickr and things of that sort. Then we have the commerce sector which has grown tremendously over the years, you know, Amazon, eBay, all of the shopping services. And then we have the social sector, which has really grown tremendously in just the last few years, things like MySpace and, of course, Second Life, as we saw in the opening. And finally, we have the play sector, where we put the online games, World of Warcraft and Xbox Live and some of the new Sony services.

Some of these are tantalizingly, tantalizingly close to what we think of when we say 3D Web. But others, quite frankly, seem well-entrenched in the 2D experience. So the question is, how are we going to get from point A to point B? How are we going to get from the 2D Web to the 3D Web? Well, let's look at this in five key dimensions. Nothing too complicated, here. We're going to look at presentation, sort of the way things look. We're going to look at user content -- how much content can users contribute? Is that even an option in these various venues? This notion of persistence -- and if you haven't heard of that, it's kind of a new thing for me -- persistence is this notion that you can leave evidence of your travels through this particular part of the Internet. So leaving tracks is the concept here. You know, in Second Life if you carve your initials on the tree, when you come back days, weeks, months later the initials are still on the tree. That's persistence. And then social interaction, and finally modeling -- or if you prefer, physics. And physics is going to underlie the modeling.

So starting with the static Web, we sort of baseline that. Relatively limited in terms of presentation, user-generated content persistence, social interaction and modeling. Not much going on.

Now if we move to the right a little bit here, towards a richer experience, we place things like MySpace where, in fact, the presentation has taken a step forward relative to the static Web, as has content, persistence -- because you can put things up and they stay there, and you can come back and find them. Social interaction certainly is a big part of MySpace. Not much going on in terms of physics, but to be expected -- this is largely a 2D experience.

Then we move to something like Google Earth, and Google Earth is really kind of interesting. As we were mapping this out, we had big arguments of where to put the dots. Clearly, presentation has taken a big step forward with Google Earth; there's no question about that. I mean the first time you experience Google Earth there's always the "wow" as you're flying around the planet and zooming in on various portions of the globe. There's not much of an opportunity for user content, although Google is starting to address that, and you can now go to the 3D warehouse and stick the Eiffel Tower on Paris if you want to do that.

But really, Google Earth is not much for social interaction -- not that it couldn't be done, but it hasn't been done yet in Google Earth, so we won't mark it very high there. The orange dot is still at the low end of the left end of the scale. And finally there's not really much physics going on. We're not modeling global climate change or anything like that in Google Earth.

Next is Second Life. Now Second Life is particularly interesting in that it tends to move all of these factors significantly to the right, moving towards this richer experience that we're talking about. The richer experience that's really going to change the way users feel about the time they spend online. And it's unique in that it consistently moves. It doesn't have the sort of scatter analysis that some of these other venues had, like Google Earth.

It's still not, however, the 3D Internet. And I want to be clear on that point. As much as Second Life and the other virtual worlds like it are pretty cool, they're not what we would consider or define to be the 3D Internet or the 3D Web. What we want to see in the 3D Web is exceptional levels of presentation, as we'll talk about later. We want cinematic quality in the graphics. We don't want to look at the relatively primitive images we see today. We want to make user content creation much easier. Of course we want to retain the persistence, but we want that persistence to apply across virtual worlds, not to be contained in a single virtual world. And higher levels of social interaction, better human-computer interfaces to enable that, and we'd like the behavior -- the way things act in the virtual world -- to be much more natural.

Well, I said Second Life is only an example. We often use it as a proxy because people are familiar with it. But there are literally dozens and dozens of these virtual worlds available on the Internet today. Club Penguin, which is focused on children, Webkinz, Kaneva, EVE Online, World of Warcraft -- if you're a game player. So clearly, it's not a niche. In fact, Wired Magazine says this year, virtual worlds will attract some 60 million users, which is really an incredible number and that's why you start to think, "Well, gee, you know, maybe this isn't as niche-y as I thought, or kinda flash in the pan. Maybe there's something really going on here, you know." Where's the knee of the curve? You can get to 60 million users, maybe 100 million users in a year or two in not inconceivable.

Okay, well, let's break it down a little more. You know, it's not monolithic, even though these are all virtual worlds of some kind. There are a number of flavors. First of all, we have, if I can get this out, the massively multiplayer online role-playing games. Who came up with that? It sounds like it came from the Intel marketing department, doesn't it? It's just what Intel would call it, you know, and everybody would go, "What the heck? Can't you just call it an online game?"

Anyway, that seems to be the handle for this kind of virtual world, and you know, these are focused on achieving the ultimate, winning all the goals, saving the fair maiden. You know, there's a lot of cooperation and strategy. You have the guilds in World of Warcraft. Very significantly, this is a rich, 3D environment, and the games have really pushed the 3D experience, and I think to their credit. If you want to look to the future and how the general experience might be a few years down the road, you look at the games, because they tend to be doing the best job here.

Then we have the metaverses, you Neal Stephenson fans, you know, all read Snow Crash some years ago and learned about the metaverse and moving back and forth between first life and second life, if I can use that metaphor. Now there are different kinds of metaverses. You sort of have the client-server variety -- that's Second Life. We're all familiar with that, but there are also peer-to-peer types, such as Croquet and Qwaq. In fact, we're going to look at that in just a minute.

The important thing, however, is that all of these, to a degree, feature user-generated content, persistence, as we were just talking about. Although, you know, the metaverses could be a little bit more exciting in terms of the visual experience. You know, that's going to be really important to them. Although the computational challenges and the communication challenges of doing that are formidable, which is why we all get up and go to work in the morning.

There are a couple more types of virtual worlds I just want to mention. And we'll actually look at one here a little bit later. The next one up is paraverses, and this is something of a mash-up of virtual worlds and physical. So we're combining the physical with the virtual to form what's called a paraverse. A good example, as we'll see, is this idea of virtual surgery, where doctors can experiment, practice, learn complex surgical procedures, sort of operating in a virtual environment that is projected onto the real environment, what the actual patient looks like. So that's a particularly interesting kind of virtual world, and there's lots of potential there for training and education and maintenance and so on.

And finally, one that's very interesting, and maybe I'm stretching the definition here a little bit, but it's so neat, I think it's worth talking about it. And that's something called. Machinima. It sound like an old Yiddish expression. It's not. Machinima refers to movies or television programs that are actually made entirely inside of the virtual world. You know, no cast, no crew. You don't have to lug the lights or anything like that. You create the characters and the environment entirely within the virtual space and make a movie of it. And if you think, "Well, that sounds kinda hokey, you know, movies in Second Life, I don't know." But in fact, a recent series of shorts in Second Life featuring someone called Molotov Alva, someone who's dropped out of first life and now exists only in Second Life -- that series of shorts was purchased by Home Box Office, by HBO, for a six-figure sum, and will actually premiere on HBO I think sometime this fall, as I recall.

So, you know, it's not just for fun and games. People are actually getting serious about creating this kind of content. Well, okay, enough, you know, with the taxonomies, the virtual worlds, and all of that. Let's take a look at one. And the one I've chosen for this morning is Qwaq -- we've already seen a little bit of second life. But Qwaq is particularly interesting because it's aimed at creating virtual worlds that are useful in the enterprise setting, which is another important aspect of this move to the 3D web. It can't just be a consumer phenomenon; it's got to apply in the workplace as well. So if you will, join me in welcoming Greg Nuyens from Qwaq. I don't know if they call you Mr. Qwaq.

Greg Nuyens: Yeah, you could.

Justin Rattner: But it's great to have you here today.

Greg Nuyens: Good to be here, Justin.

Justin Rattner: Thanks for coming. Okay. Well, it's not Second Life, that's for sure.

Greg Nuyens: No, that's right, that's right, absolutely.

Justin Rattner: So tell us about -- these are called Qwaq Forums, right?

Greg Nuyens: Yes. Absolutely. So, as you mentioned, at Qwaq, we're focused on virtual spaces for enterprise collaboration, and so here I've taken your dress style today -- we're looking over my shoulder at you.

Justin Rattner: I'm kind of boring.

Greg Nuyens: Well, okay. One of the great things about virtual spaces --

Justin Rattner: [I wore this shirt and everything.]

Greg Nuyens: Yeah, we can have a place for every purpose, so if you want to dress yourself up like a 2D browser, we can just press the arrow and change to another site. Here we're looking at you, and you're going to join me in here.

Justin Rattner: Okay.

Greg Nuyens: And we'll be able to see that whenever we choose a new avatar there, you get your new impression. Your shirt nicely matches. And we think that people are going to want to do this more and more in virtual spaces, as they get more prevalent.

Justin Rattner: Now, who else is here?

Greg Nuyens: Oh, well JDM joined us from Portland here, one of our colleagues from Intel.

Justin Rattner: Hey, JDM.

Greg Nuyens: Wave JDM.

Justin Rattner: Yeah.

Greg Nuyens: And in addition to having the correct appearance, which is the natty shirt we see before us, it's very important --

Justin Rattner: Oh, this thing? [Yeah] [unintelligible.]

Greg Nuyens: No, no, please. People are going to want to be able to choose the other elements of their behavior, like their gait. So here, for instance, we see an abstract representation of the walk in the balloon man here.

Justin Rattner: Yeah, Michelin Man --

Greg Nuyens: Exactly. The [unintelligible] [site] guy. Yeah. So here you've chosen your avatar, and now that you're dressed for success, we're going to return to the boardroom and take a quick look at -- actually, I'll head back, sorry. Here we're going to take a look -- you're going to be doing some virtual surgery later on.

Justin Rattner: Yeah, that's right. That's right. Don't give it all away.

Greg Nuyens: I wish you well -- I wish you well.

Justin Rattner: Oh, I'm not going to do the surgery. Please.

Greg Nuyens: Oh, I understood you were the patient.

Justin Rattner: Oh, god help me, no.

Greg Nuyens: So in here we can look in our boardroom environment -- of course, if you're going to be able to work with a distributed team, what you need to be able to do is use your documents and applications just like you do when you're working together in a physical space. So when I reach to this Web site here and I pick this, what I can see is that this is a live document, and we go to the other -- next page, and we see some nice coverage on Justin in Business Week. Of course, you can do the same thing, and when you pick a link, I see the change that you did no matter where we both are. So distributed teams can now collaborate regardless of where they are.

Justin Rattner: Okay, but, I mean, just looking at, you know, productivity apps I guess blown up here wall-size, it's not quite the 3D Internet experience I'm thinking about.

Greg Nuyens: Well, we understand Intel loves numbers, but we'll go beyond that. Of course at Intel and other places, people want access to their enterprise applications. You're quite right.

Justin Rattner: Right.

Greg Nuyens: Beyond the productivity apps, they also want to be able to consult things that are behind their firewall, so the enterprise server allows you to do exactly that. And here we're going to look at a couple things that we're doing. In fact, today, we announced -- Qwaq and Intel -- that we're collaborating, bringing Intel's Miramar technology into Qwaq Forums.

Justin Rattner: Oh, yeah, yeah, oh, that's fantastic.

Greg Nuyens: Yeah. So if we look at this representation of some of the Intel IT infrastructure, we can decide that we want to fade that live application that is the distributed virtual space into a representation of all of the Intel IT. Here we're looking at three different data centers. In fact, you can probably tell me what those are.

Justin Rattner: I think this is Folsom.

Greg Nuyens: Okay.

Justin Rattner: I thought I saw the Israeli flag somewhere over there, so that's probably Israel over there. It looks like we've got a problem going on here somehow.

Greg Nuyens: It's always something, isn't it?

Justin Rattner: It's always something. If it's not one thing, it's another.

Greg Nuyens: That's right. It looks like we've got a thermal alert here, and we can see, through some of the modeling of the data center, whether --

Justin Rattner: Okay, so you're moving literally between enterprise app to enterprise app here, as we're trying to diagnose the problem.

Greg Nuyens: That's right. And each one of these logical elements in the information space, just like in the virtual workplace, is linked together, and we can traverse those links as well. So if we had more time, we'd go look at the incident manager, make sure that the right team's on it, but I think --

Justin Rattner: We pay people to take care of these things, so we won't go --

Greg Nuyens: Exactly. Plus, it might be hard.

Justin Rattner: Yeah, that's right. Probably beyond our skill.

Greg Nuyens: Yep, that's right. So we can bring back the forum -- this is, again, a live app -- and we can transition, then, between our 2D enterprise apps and our 3D immersive space that we access them in, sharing all those in a single context.

Justin Rattner: By the way, I really like this conference room. I'm going to talk to my site services people about duplicating the lighting.

Greg Nuyens: Absolutely. We'll have it drop-shipped to you forthwith.

Justin Rattner: Okay. You can have the 3D print button?

Greg Nuyens: Absolutely, that's stereo lithography of a whole room. That's going to be great.

Justin Rattner: Okay.

Greg Nuyens: So looking at the time, it looks like, actually, you're next session is going to start here, so why don't we go on in and take a look for a second. Here's the auditorium. It looks like people have started to show up in here. We have the films and the models. One of the nice things about virtual space is not only can a team who's not physically together work together, but you can do things that are impossible in the real world like growing a scale model that's actually going to be featured. And you're going to learn a little bit about this model shortly.

Justin Rattner: That's right. We're going to go to work on a patient here in just a minute.

Greg Nuyens: Well, since they're prepping the patient, it probably means that it's time to go. But thanks very much.

Justin Rattner: Okay, thanks, Greg.

Greg Nuyens: It was great being here today.

Justin Rattner: That was Greg Nuyens, everybody.

Greg Nuyens: Thank you.

Justin Rattner: Thanks, Greg.

[Applause]

Justin Rattner: Pretty cool stuff. Okay, so it's usually at this point in the conversation, at least as I've discussed it with my colleagues at Intel, that people say, "Yeah, we were here five years ago, you know? And it didn't happen then, and it's not going to happen now." So while you're trying to get everybody excited, you know, the chorus of "Nos" begins. And, in fact, you have to respect that. There was a lot of work done five, even 10 years ago, on the Internet to try to create these 3D spaces, and it didn't really go very far.

So why do we think it's going to happen now if it didn't happen then? Well, we've already mentioned early on about, you know, all of these trends and their own laws that govern their growth and advancement, but there are other trends, megatrends as we'll call them, that I think are equally important in the evolution to the 3D Internet.

The first of these is, of course, social networking, which has grown tremendously over the last few years as people find these environments a way to meet people, communicate with people, friends, family, new acquaintances, what have you. User-generated content is incredibly important. Who would have thought that people would rather watch sort of crude videos being uploaded on the Net than primetime television programs? But, in fact, primetime viewing has dropped by half. What are people doing? They're on the Net watching TV, little TV, not big TV, but they're watching that kind of TV. For whatever reason, they find it pretty exciting.

There's also broadband connectivity. Even if you go back five years, and certainly if you go back 10 years, you were doing great if you had a 28.8 or 56K modem. There really wasn't the bandwidth. And maybe you had ISDN or something like that, a few more bits, but broadband, in terms of megabytes per second of bandwidth into the workplace or into your home, just wasn't there. So that's another big trend. And, of course, we continue to build out that infrastructure. And now we're talking about wireless broadband connectivity.

I think another important trend is high definition television, high definition displays and the content to feed those displays. Users are getting increasingly accustomed to very high quality imagery as it's provided on formats, either by satellite, or increasingly by HD-DVD and the hi-def disk formats, and now available, you know, through the players and through the game consoles and all sorts of different sources. So people have this expectation now that whatever they're going to watch, they're going to watch it in high definition.

And I think another significant trend, a megatrend if you will, is this notion of the virtual economy that can exist in the virtual world and be connected to the real economies around the world. That's one of the really intriguing aspects of Second Life, that you can actually build things and sell things inside of Second Life and can convert your Linden dollars into real dollars, or whatever the currency happens to be. So through these trends is really emerging an entirely new set of interactions.

The key difference, and this, you know, is really what I want you to walk away with, is that it is now, perhaps for the first time, technologically possible to do this. It's also socially acceptable to be a part of these virtual worlds, and it's even culturally appropriate. In fact as we've studied this inside of Intel, we're amazed how different cultures around the world have embraced the technology. If you look at Asia in particular, their notion of virtual world is very much this fantasy environment, those sorts of things. You look at what's going on in other parts of the world and there are reflections of the cultural values within that part of society. It's that change -- not just the technology, but the social and cultural changes -- that will really be the foundation for the growth and development of the 3D Internet.

So we've seen a metaverse. Now let's take a look at a paraverse and how that might have a fundamental impact on medical care. Now we'd invited one of the legends in this field, Dr. [Court Cutting] to be with us today and actually walk us through virtual surgery. Dr. Cutting is really an amazing person. Sometimes, as I've heard the story, he stops in the middle of a surgical procedure and writes some code. Believe it or not. And then goes back to the surgery. So he's a real advocate of virtual surgery, and even though his schedule didn't permit him to be here today he did provide us with this video. Let's take a look.

Now, I have to warn you that if you're a little bit squeamish you may want to turn away from some of the scenes. If you elect to stay in your seats, we did not provide motion discomfort bags for anyone. But if you feel the need, get the hell out of here, all right? Let's play the tape.

[video]

Justin Rattner: That was great. Let's bring out Aaron Oliker from BioDigital Systems. Aaron is going to take us in for a closer look at what you just saw in the video. Come on out, Aaron.

Aaron Oliker: Hi, Justin.

Justin Rattner: Glad you're here. We just saw Dr. Cutting talking about some of these cleft palate surgeries. Why don't you give us a closer look?

Aaron Oliker: Right now we're looking at an image of a little girl with a cleft, and we're going to go into an animation and show actually how that surgery is performed. And right now we're making some incisions into the skin, and these incisions create peninsula of tissue called flaps. I see some people holding their hand over their mouths. It's okay.

Justin Rattner: You deal with this every day, right?

Aaron Oliker: You get used to it. These create little peninsulas of tissue called flaps, and you can see here these flaps actually have names to them. Like, one's a C-flap, the other one's an M-flap. And they're like pieces of a jigsaw puzzle. And right now what we're going to do is we're going to raise the cartilage, overcompensate a little bit, and we're going to raise the L-flap --

Justin Rattner: Okay, so these are the nostrils, here, right?

Aaron Oliker: Yeah. So right now we're putting this little girl's face back together. It's a geometric puzzle. You're actually figuring out how to take these pieces of tissue and put them back together as a puzzle to recreate this girl's face. That's what this is all about.

Justin Rattner: Now, we're looking at an animation, right? This is not a simulation.

Aaron Oliker: Yes. This is not a simulation. But that's where we're moving with this.

Justin Rattner: Right.

Aaron Oliker: And so here the little excess tissue is flying away, and we're going to close the nostrils. And when we did this animation, we had to actually have Dr. Cutting look at every individual step of what we were doing because we had no idea how a tissue piece moves. You know? And now we close the entire face, and you'll notice, at the end of the close, that the animation very closely matches her postoperative results. And the wonderful, neat thing about this is that people actually use this to learn how to do surgery, and we've gotten wonderful results from all over the world from -- we have people coming in from China, India, everywhere, saying that this helps.

Justin Rattner: Wow. And BioDigital is developing this kind of technology?

Aaron Oliker: Absolutely.

Justin Rattner: For the medical community.

Aaron Oliker: That's what we're all about. We're about developing technology that will help people learn surgery.

Justin Rattner: Fantastic. Thank you, Aaron. Thanks for being here today. I know you came a long way.

Aaron Oliker: Thank you so much.

Justin Rattner: Let's bring out Joey Teran. Professor Teran, come on out, Joey. All right. So what you was an animation, right? There was no physical simulation involved in that, but that's clearly where we want to go, and Joey, you've been working on that, so you want to give us kind of a preview, a peek, into the next generation of technology?

Joseph Teran: Yeah. So the movies that Aaron showed are pretty spectacular and are really realistic depictions visually of what happens in surgery, but it's created by an animator and it's not very scientific. What we'd like to do is solve the governing physical equations and create a videogame-like environment that residents can be trained in the procedures with in real time or actually come up with new procedures.

Justin Rattner: I can just see the TV ads, "Learn surgery at home."

Joseph Teran: There's actually a videogame that doesn't use simulation but it's based on surgery right now. But to give a quick example of what's actually possible right now with simulation, I'm going to show a simulation of a Z plasty. And so here you see a rectangular slab of tissue that we're going to stretch by making an incision in a Z-like shape and then pulling it in the horizontal direction here.

Justin Rattner: Dr. Cutting talked about finite element methods. Is that what you're using here?

Joseph Teran: Yeah, we're going to use -- finite element method solves for the tissue elasticity here.

Justin Rattner: Okay.

Joseph Teran: You see, after we make this incision, we have this elastic governing PDE for the deformation, where everything should actually go, and we use that to realistically solve for [where] these tissue flaps should go. And tissue flaps, as you saw in Aaron's animations, are sort of the fundamental building blocks of every surgical procedure.

Justin Rattner: Right.

Joseph Teran: Here we're just showing like a proof of concept that we can -- with the finite element method and simulated governing physical equations, we can get an actual realistic result.

Justin Rattner: It looks like somebody's suturing this back together.

Joseph Teran: Yeah, so here these little blue lines are virtual sutures for sewing the tissue back together after you've done the procedure. And it's used to elongate tissue, and it's sort of a ubiquitous surgical procedure. It shows up all over the place.

Justin Rattner: So this is the basic technology for doing these simulations, sort of -- you know, what's the next step and how long will it take?

Joseph Teran: The obvious thing about these is that it's on a rectangular, simple geometric piece of tissue. What we'd like to do is to do a cleft lip and pallet, but you see the geometry of the face for a cleft lip and pallet is so complex that it requires way more degrees of freedom than we can handle in real time right now.

Justin Rattner: So you need more computing power?

Joseph Teran: Much more computing power.

Justin Rattner: I know a company that can solve that problem.

Joseph Teran: We hope so.

Justin Rattner: All right. Thanks, Joey.

Joseph Teran: Yeah, thanks a lot.

Justin Rattner: That was great. I appreciate it. Joey Teran, everyone. Okay, well, you've heard about, you know, metaverses and paraverses and, you know, as we were just learning here in the virtual surgery example, there's, you know, an incredible need for more computing power. And to sort of quantify that, I just kind of want to walk through sort of, you know, the basic loops that are present in the 3D environment.

Okay, so first let's take a look at the 2D web, the one we all know and love and use every day. Some of you are using it right now, I'm sure. You know, here's the basic loop. User clicks on a URL, that goes upstream using HTTP and arrives at the server. The server translates that request to a particular page. That page may be a static page, it may be an active page. But a relatively light load on the server. The page information is assembled, sent back down the wire, back down to something running on the client. There's a 2D browser, as we call it, as we know it, that renders that content, brings up the new page for the user and then the user can iterate on that, click another URL or browse further on that page. So that's a relatively simple loop, relatively light load on the server.

Now, the 3D version of this loop gets much more computationally intensive. So now, we don't talk about clicking on URLs. We talk about users taking actions in this 3D space. Right now, we're talking about spaces and places. We're not talking about pages any longer, right? And by the way, those spaces and places may be largely, if not entirely, user-generated content. So this is not something that could be carefully crafted by a programmer or developer. You know, users are just sort of putting together primitives.

In fact, Second Life has told us that, in fact, Second Life began with essentially zero content. It was just the structure that you see here, and then users have built terabytes, petabytes of content from that basic capability. So there's some action. The server is now collecting all of the actions from all the users that are present in that local region, or the visible region of space, and then resolving all of the interactions, right? So if objects are moving, if avatars are moving, and one is friend of another, you have to make that visibility calculation. When all that's done, then you have to generate a new per-user model, right? So you have to create a unique view, you know. It's not just send the page down there; every user has a unique view of the place, the space, and you have to generate a unique view for every user that's in the visible area of that place. And then that model gets sent down the wire, and now the browser really has a lot of work to do. This is a 3D browser because now it has to interpret the model and construct a 3D visualization. It has to render the model for viewing by the user, and then, of course, there's the iteration.

The interesting thing is, while there's really no real-time demand in the 2D web, there's an essential real-time demand in the 3D web in that you want to be doing this iteration at close to frame rate. You'd like to be able to do it at 30 frames per second. You might not be able to do that given today's hardware, but in the not-too-distant future, frame rate will be possible. So we're updating all of these views, constructing these models, sending them down the wire, having them rendered 30 times a second, truly computationally intense. And let me qualify that for you. We'll look at servers, clients, and we'll look at the network.

So here's some analysis we did. We ran the numbers in the lab. We're looking at the number of users that a single server can support here. If you look at something like EVE Online, one of the massively multiplayer games, you see tens of thousands of users on a single server -- about 35,000 here. World of Warcraft, a little more demanding environment than EVE Online, that number drops to about 25,000 users. But a virtual world fully supporting user-generated content as exemplified by Second Life is now down to about 160 users, and if you've spent time on Second Life, you know you'd really like that down by another factor of two or three to get a good experience.

When we break that down, if I can break it down, we find that on the server, 75 percent of the time is actually spent in what we'll call computationally intensive components. This is where we're doing intensive floating-point arithmetic using matrix mathematics. Maybe we're doing finite element modeling, as we've heard from Dr. Cutting. We're doing some serious computing on the server in order to construct these per-user models. And when we compare that to let's say the massively multiplayer games, we realize that we're talking about an order of magnitude more computing power, and that's sort of the bottom end. Realistically we're probably talking about a

hundredfold increase in server computation to support the 3D environment from the server.
It's similar -- not surprisingly -- on the client side. And here we're looking at client loads, looking at percent of CPU utilization on the client. The top row in the table is a typical 2D Web site. Processor utilization relatively low, which suggests how far we've come, maybe 20 percent, and there's virtually no load on the GPU. GPUs don't really have much to do in the 2D environment.

Things kick up a little bit when we move to something like Google Maps and we actually have to do some rendering. More load on a CPU, GPU gets involved, starting to pick up part of the task. When we move to something like Google Earth, now the processor load is up and the GPU load is beginning to climb. Now we're talking about 10 to 15 percent load on a typical GPU. When we move to a true virtual world environment, the 3D, the full 3D experience, CPU utilization is now up to 70 percent and the GPU is perhaps at half power. And again, as we'll see in a minute, the visual result is not nearly what we would like it to be.

In this case, 65 percent of the time on the client is actually spent in the compute-intensive portions of the calculations, and it's very similar, very floating-point intensive. Matrix methods, all the stuff you learned in college and hoped you'd never see again; that's what makes this stuff go. And the growth in computational intensity is actually quadratic in nature, right? It's a function of how many things are interacting with how many other things. So you're getting that cross-product. So we project at least 3X the CPU and perhaps 20X the GPU load compared to a 2D environment. So very computationally intensive on the client as well as the server.

Now we can sort of stop there, but it actually affects the network behavior and I want to talk just a little bit about that. What you see plotted here is the bandwidth demand in kilobytes per second for a massively multiplayer game. And as you noticed, it's relatively peaky, but not all that intense. Why is that? Well in the games, when you install the client, most of the content actually moves down to the client. It's no longer on the server. So we don't need the network to do much other than communicate the movements of the avatars and other objects within the space.

If we look at something like Second Life, where you're using user-generated content which, well, maybe you can pre-cache if you're taking somebody on a tour. You can visit the different locations in Second Life and you can sort of move that information in the cache. But if you're moving around fairly freely in such an environment, lots of information is being pushed across the network and now you can see how much more intensive the bandwidth demand is. So when you can't cache the content you really put stress on the network and the numbers we've run in the lab suggest probably 100 times the bandwidth for a full 3D virtual world experience, a la Second Life or any of the others, compared to a massively multiplayer game which is already a 3D environment.

Okay. Well, we've kind of hinted at this. I've made several disparaging remarks, if you will, about the quality of the imagery in the virtual worlds today compared, let's say, to a PC game or a console game. If you look at the video loop that's running now, you notice that a lot of things we'd like to be a part of the experience are missing. Notice there are no shadows or hardly any shadows. Little in the way of reflections. And quite frankly it just doesn't look very natural. So what's it going to take, computationally, to get this experience to the level of reality we'd like to see? And to tell us how we're going to do that, I'd like to introduce Daniel [Pol]. Daniel is a research scientist in the microprocessor technology lab at Intel. Come on out, Daniel.

Daniel Pol: Hi, Justin.

Justin Rattner: Thanks for being here. You heard me talk about how flat these virtual worlds seem to be looking these days. What are we going to do about it? And maybe you can start off by sort of telling us what the problem is.

Daniel Pol: Well, if you look at games or virtual worlds in today's graphics, you often notice there's some problems. For example, with shadows, if you take a look at the left image on the top, it's a nice-looking image here of those bowling pins. But if you look closer at the right image, there's something wrong. And those bowling pins have been knocked over, but the shadow still remains on the floor.

Justin Rattner: Yeah, the shadow didn't move, okay.

Daniel Pol: And the reason why this is, is because an artist manually painted the shadow on the texture of the floor. And when it comes to user-generated content, you can really not expect that the average guy draws shadows somewhere in the scene.

Justin Rattner: Yeah, that's a really important point. Let's pause for a second on that. You know, user content is really different than the kind of content that a graphics artist, let's say doing a game, would create. Because, you know, you might spend hours, even days on a particular scene. But a user-generated scene, I mean, nobody's going to sit there for days and paint the shadows in.

Daniel Pol: Exactly. We're expecting video shadows [where one could] place an object into it. Another example I brought along is the image at the bottom where somebody went to the pain of implementing shadow maps, and they had lots of [unintelligible] [aliasing] effects with those things.

Justin Rattner: Right.

Daniel Pol: But what we want to do is high quality, cinematic-like rendering. And to do so we're using a tool called ray tracing. And ray tracing enables us to make more realistic effects like physically correct glass with reflections, with refractions on it. We can simulate water in a very good-looking way. When it comes to nature, we can increase the number of triangles. We can up the poly count. It looks much better. And, of course, lights and shadows, we can do those things with ray tracing very good.

Justin Rattner: Wow, that's stunning, isn't it?

Daniel Pol: Yeah.

Justin Rattner: Yeah, really amazing. Your reflections are -- and the water effect is fantastic. Okay, so I know you've been working on other things in the lab. Maybe we want to fade out of this and move into the next segment.

Daniel Pol: So at IDF 2004 we showed one of the first real-time, ray-trace demos in the world. And it was running on a really huge cluster with 50 Xeon server processors. And at a resolution of 640x480, we were able to show off about four frames per second.

Justin Rattner: I remember Pat had that huge rack of servers on stage. Notice we've gotten rid of the rack today.

Daniel Pol: No, we don't need a rack today. So last month I was at Leipzig Games Convention and showed off Quake 4 Ray Traced, the map running out on its upcoming quad-core CPU. We were running at a resolution of 768x768 with 90 frames per second.

Justin Rattner: Wow, that's amazing. All right, and where are we today?

Daniel Pol: So today here at Fall IDF, we are going to show a live demo now about real-time ray tracing in Quake. We're having HD resolution running here. We're using an eight-core system, a dual quad-core system. As you can see in the task manager, CPU utilization is almost at 100 percent, so we're making really good use of those multicore architectures.

Justin Rattner: And this is all of Quake 4. I noticed when you went back it's --

Daniel Pol: Yeah, it's the complete map. We don't hide anything to cheat. So it's really the whole thing.

Justin Rattner: And the frame rate?

Daniel Pol: Yeah, about 100 even.

Justin Rattner: Wow, you keep tweaking this. Every time you run this demo, it gets faster.

Daniel Pol: Yeah.

Justin Rattner: Great. We'll have to schedule another IDF so you'll crank the number up.

Daniel Pol: Yeah, so to show you even more of those special effects I mentioned earlier, and even some additional special effects, I brought along a video. So let's please have a look at the video.

Justin Rattner: Okay. Watch closely. This is pretty cool.

[Video]

Justin Rattner: Tell us a little bit about what's going on, what we're seeing.

Daniel Pol: Yeah, so now we're seeing the great reflections for pixel effect reflections. We're running user-generated content. You could just define your material as reflective, and that's it.

Justin Rattner: Yeah.

Daniel Pol: Also, we can enhance game play. For example, we've put a mirror here on the edge. And through the mirror we can look into it, and we can see an incoming enemy, and therefore we can shoot a rocket at the place where he will be, something you couldn't do with the current technology.

Justin Rattner: And what's this?

Daniel Pol: And a game called Prey, they heavily use those so-called camera portals where you could look into a different area where you're going to be, once you walk into it, [or you're beamed to it]. And those kinds of effects are also very easily implemented with ray tracing and could be used in user-generated content as well. And the nice thing is you could put a camera portal in front of another camera portal, and then you have those multiple bounces of camera portals.

Justin Rattner: Right. Going to infinity there.

Daniel Pol: Yeah.

Justin Rattner: Right. Okay. Well, that's what we're looking for. That's the kind of realism that we want to see. Thanks Daniel.

[applause]

Daniel Pol: Thanks Justin.

Justin Rattner: So we really think ray tracing is going to be an essential technology in this move to the 3D Internet. You know, particularly when you're in the presence of user-generated content you don't expect users to make everything look pretty the way the professional graphics artists do in the games. You just want them to build the scene and render the scene, and get that cinematic quality that you saw in the Quake 4 demo. And you know, if the user doesn't like the way the scene looks, just hang up the lights, move the lights, get rid of some lights -- whatever you want to do. Add a mirror, as Daniel was demonstrating, or open a window, as the case may be. But the physical correctness of ray tracing gives you a natural looking scene with a minimum amount of effort.

Well okay, so Daniel showed us a way to get the look we want. We're not going to be satisfied with just good-looking imagery. We want it to behave. We want these worlds to behave in a very natural and realistic fashion. And to do that I've asked someone who, you know, is really known for getting that realistic look in the popular movies you see today. That's Professor Ron Fedkiw from Stanford University, a good friend of mine. Ron, come on out.

[applause]

Justin Rattner: Now this guy is responsible -- this guy -- hi Ron.

Ron Fedkiw: Hi Justin.

Justin Rattner: Ron is responsible for some of the greatest movie special effects that you've seen I think. You know, if you saw Poseidon Adventure -- now, you told me this. There's no real water in Poseidon Adventure isn't that true?

Ron Fedkiw: And in fact the ship itself is not real. I had to put fake water around the fake ship.

Justin Rattner: So no Poseidon and fake water, right. In fact, George Lukas referred to you Ron, I think in SIGGRAPH 2005, as "Mr. Water" for that effect. So congratulations for that, and how many people saw Pirates of the Caribbean, Dead Man's Chest? What was the great effect there? Davy Jones right? All CG right?

Ron Fedkiw: Yeah. It was nice. In fact, one of the big things about Davy Jones was that they motion-captured the actor, and then they layered all of the flesh and articulated [unintelligible] bodies on top of it so that you could read his performance through the CG character that we made.

Justin Rattner: Okay. Well that's the kind of technology we see in the 3D Web. Maybe you can kind of walk us through it and, you know, show us the different kinds of effects that we're going to want.

Ron Fedkiw: Okay. So we just have a few sample movies here. On the upper left corner we're going to give an example of how we can slice a block of tetrahedron up into different pieces. It's sort of prototype for fracture or virtual surgery. In the upper middle, we have a virtual human face. We've actually built a model of this from one of my graduate students at Stanford and then put him in a motion capture lab and recording how his face moves and speaks. And then we built the internal muscles, the skull. He had MRIs done and masks -- and you'll see it move a little bit according to one of the sentences.

In the upper right and lower left are some of the direct results of collaborating with Justin and his team at Intel. And this is some [MPI cloth] simulation that we've done using a lot of processors.

Justin Rattner: You know, when I talked to the virtual world folks, cloth is right at the top of their list. That's going to be an incredibly important effect. And what's in the lower right?

Ron Fedkiw: In the lower right is a bunch of [deformable finer] element models of donuts. We picked the colors so they would match those chocolate and white powder donuts in the end. Just so you'd see the difference.

Justin Rattner: I can see a Dunkin' Donuts ad in your future Ron. [laughs] You know, another special effect. Okay. You know, there are some other things that we hear about a lot, water effects and so forth. And I know you've been working on those as well. Tell us about this.

Ron Fedkiw: The upper left was a very early simulation we did before making Poseidon and the two Pirates movies, to show we had the capabilities of doing this sort of thing. And the rest of them show some of our examples of doing fire. We can burn cloth. A lot of the work we do at Stanford is on interactions between cloth and rigid bodies and fire and fluids. The goal is to get everything to interact together in one world. And then in the fire in the lower left you can see a really full smoke. In the one to the right of that you can see the cellular patterns in fire. It turns out catching the cellular patterns is really important if you want that mesmerizing look where you can watch fire -- you know, you can sit at a campfire or a fireplace and look at it for a long time and get sort of emotionally caught up in it. And having that cellular pattern actually helps. So this is a nice video for proof of concept to show that we can start to capture that mesmerizing effect.

Justin Rattner: Okay, let's put it into motion because it's much more fun to see it moving than just in these still pictures. And you can kind of clue us in on what we're going to see here. So what are we doing here?

Ron Fedkiw: Here we're showing how we can model a mesh, so we can actually cut a slinky out of a block of tetrahedral. It's usually hard to construct these 3D models, but here, in almost real time, we can actually construct a new 3D model out of the given ones. You can imagine --

Justin Rattner: Slinkies, now.

Ron Fedkiw: Yeah, you can imagine making all kinds of complex things out of simple stuff in your games. This is coupling between a curtain and a bunch of rigid rings on a rod, so it's mixing rigid bodies and cloth simulation.

Justin Rattner: So I can hang my laundry out in Second Life, huh?

Ron Fedkiw: Yeah. You see all the contact and collision here that we're capturing, so that things interact well with each other.

Justin Rattner: Oh, and here's your student.

Ron Fedkiw: Actually, this [year we're putting hair on him], so we're getting our hair simulator up and running, so we can have him move around and have some hair.

Justin Rattner: Okay. And what's this?

Ron Fedkiw: This is an armadillo mesh. It was scanned it at Cal Tech, and we made it deformable, so that it models a biological type tissue with sort of a [sacks] full of water.

Justin Rattner: [And helps cloth.]

Ron Fedkiw: And here's the cloth. Again, this is, you know --

Justin Rattner: Yeah, so this idea of the cloth interacting with a rigid body, I mean, that seems pretty fundamental. I mean, if we're going to put the shirt on the avatar, right, that's the kind of effect that we want.

Ron Fedkiw: Yeah, having very realistic clothes helps. Most of what you see when you look at virtual human is the clothes, so it helps a lot.

Justin Rattner: Right. And here are the fire effects, our fire-breathing dragon.

Ron Fedkiw: Again, notice the evolving cellular patterns in the fire, which is something you usually don't see.

Justin Rattner: Right. Now what's this?

Ron Fedkiw: Here we injected some oil and water, and then we drop a flammable liquid on that, and it catches the oil on fire. It shows some of our ability to model all different fluids packed together.

Justin Rattner: Amazing. All right. Fantastic. Thanks, Ron. We'll look forward to it. I know you're going to need a lot more computing power, and we're working on that.

Ron Fedkiw: Thanks, Justin.

Justin Rattner: All right, we'll see you later. Ron Fedkiw, everybody. Okay. So we've talked about getting virtual worlds to look good, Ron showed us the kind of technology it's going to take to make them behave in very natural, very realistic ways. But I don't want you to leave this morning, walk away from the keynote thinking, well, if we just make it look good and act good, we're done -- that's the direct path to the 3D Internet. There's much more involved here. We don't have time to, of course, go through it in an exhaustive way, but I want to talk about a couple of things before we close.

The first of these is, of course, we, you know, can create a great 3D look, but that's not sufficient to get the user really engaged in the 3D experience. And I think as we've seen with the Wii game, you know, a tiny improvement in user interactivity really can bring the user into that experience in a very unique way.

So part of getting to the 3D web is really creating the human computer interfaces that will let us be an active, integral part of the 3D environment. This kind of unusual-looking device here is the Novint Falcon. This is a force feedback device that can be used both to provide input and output in the sense, you know, to give you various forms of mechanical resistance, sort of let you feel your way through one of these 3D environments. I understand they're working on gender-specific models of that. I'll just leave that alone.

And this is actually a 3D mouse. This comes from 3D Connection. And this has six degrees of freedom so you can move very easily. You don't have to pick it up. I'm just showing you this, you know. But you can move this thing back and forth and up and down and twist it and get all the six degrees of freedom. There are many other pieces, I think, of bringing that 3D experience home. You know, we've shown haptics here with the Falcon. Facial tracking is something we've done a lot of work on at Intel, you know, gesture recognition, full-body motion capture. Work's been done in these 3D printers, these OpenGL printers, that let you take something that you've created in a virtual world and actually print it out, actually have the physical representation of that. And of course, the 3D display technologies are going to be a big part of it. Hopefully, we won't have to have everybody with goggles on. We'll have 3D displays and you'll be able to sit comfortably at your desk or in your living room and enjoy the full 3D experience.

Another dimension of the 3D problem is the notion of identities. Avatars are a wonderful thing, but they create problems of their own. It's very important, you know, it's almost critically important in these virtual worlds that avatars preserve this notion of fungibility, as it's called. Fungibility is what allows for role experimentation. You know, you can choose varying degrees of anonymity. You can use techniques like role isolation, where there are many avatars for a single user, or role delegation, where many users share one avatar. The trick is making sure that that's done in a trusted way, so when you're dealing, when your avatar is dealing with another avatar, you can trust that avatar is an accurate representation of a member of the virtual space.

Another aspect of identity in the virtual environment is behaving correctly with respect to accepted norms of social interactions. So we don't want our avatars crashing into scenes and damaging content and doing things that would be considered socially unacceptable.

So identity is going to be very important in the future of the 3D web. Much as we've been developing technologies for trusted identities in today's web environment, we need to adapt that technology, move it forward into this world of avatars and virtual environments, and I think you'll see people sort of overloading their current trusted identity technology as we move forward with this notion of verifiable representations.

Well, of course, you know, the big question is, "How are we going to get there?" Some of this stuff looks pretty formidable, and maybe too big of a mountain to climb. We sort of broke it down here as we see it, much as in the 2D space where we've developed standards and browsers and web services and a variety of programming tools and environments and of course, the underlying run-time technology. We need to do the same thing for the 3D environment. If you look at what's available today, it's almost all proprietary technology. It's like those old wall gardens of AOL, CompuServe, and Prodigy in the early '90s.

There are a few exceptions -- Second Life has done things like open-sourcing the client. A number of virtual worlds have actually done that, made the client open-source. The servers, however, are still proprietary. There's virtually no interoperability. You can't move an avatar in one environment to another, although people are working on that. There are tools, particularly tools that have come from the CAD space, the 3D space of computer-aided design. They need to be adapted, improved, simplified for general use in the 3D Internet. And of course, we have a ton of work to do in the runtimes, because we really haven't been focused on supporting a 3D computational environment as we've talked about.

What can you do to help? Well, we're going to need all the help we can get, so you know, I ask you, I empower all of you to become members in good standing in the team that's going to drive to the 3D Internet. We need to develop open environments, standard environments that are going to let us have these virtual worlds where we can move freely amongst them and enjoy the benefits and capabilities of the technology. We need to create better user experiences, whether it's a better human computer interface, better looking virtual worlds, more natural virtual worlds, what have you. We really need to move forward on the user experience. We certainly need to deliver much more rich content. That's going to require better tools fundamentally, and of course, more opportunities to apply those tools across the virtual worlds. And as I was just saying we need to create a trusted environment in which to do all of this -- one that's safe and secure, and overall reliable for a platform perspective.

Well, that's our keynote this morning. We hope we've given you an introduction to this notion of virtual worlds. I think it's about time for me to return to Second Life, so I'm going to walk over here to the Stargate. I'm not going to de-res this time. I'm going to do a little video projection.

I hope you've all had a great time in Second Life, and -- I'm sorry, a great time here at IDF. As we go back to Second Life, where hopefully we'll have a great time. Hey, bud! How you doing? Good to see you. I thought I ordered a couple of pina coladas. Have you got them ready? Okay, I'll be right there.

Thanks everybody. Enjoy the rest of the IDF. We'll see you next time.

[applause]

[End of recorded material]