Intel Developer Forum
Tuesday, Sept. 26, 2006

PAUL OTELLINI KEYNOTE

Paul Otellini: Good morning. Good morning. Thank you, and welcome to IDF. Today I'm going to talk to you about what's inside, and more specifically what's inside the silicon chip. Where it's going, how it's changing, why it matters, and what the opportunities those changes present for all of us here this week at the Intel Developer Forum. For computers, much like people, despite our differences on the outside, it's what's inside that counts. Now it wasn't that long ago that many were writing about the fact that you don't need any more processing power. The answer was clearly no, in the PC Magazine article from 2001. They even went as far as saying that the only people who would need more processing power were NASA scientists and designers who needed fast computers. Well the world's changed. If you fast-forward to PC Magazine last month, what they say is speed matters again. And they have a different view now. Now more than ever, we need the power that comes from new silicon. Performance matters again after four or five years.

Why is it? What's changed? Well the simple fact is we have, and what we do with our computers has changed. The things that people want to do require more processing power than ever before. PC games are approaching movie quality, as you'll see a little bit later. Google desktop search is becoming a de facto on the way we find data in our hard drives. Everyday tasks are evolving to include photo editing, antivirus software, some people doing video editing, and so forth. And of course the latest operating systems, like OS X and Windows Vista, are all driving the need for more and more compute power. If you take an example of YouTube, a 60-second, simple clip on YouTube today would have taken almost the full capacity of a microprocessor in 2003. Here we're comparing it to the first Pentium M, the chip that went inside the first Centrino, essentially would have been consumed running that one 60-second clip.

In 2004, about half of a Pentium IV top-of-the-line processor would have been consumed running that same clip. You look today, though, processing performance has given us to the point now where with the new Intel Core 2 Duo, running that clip essentially doesn't consume just a few percent of the microprocessor power. It leaves much more capacity and capability for other things that we want to do while we're enjoying that clip. Video won't just stay there. Video will evolve, and the next threshold that will happen on the Web is high-def video. It's only a matter of time before YouTube starts showing high-def video streams. In fact, this Christmas high-def will start to reach the mainstream price points as HD camcorders drop below $1,000.

When you look at the difference between high-def encoding -- that is the

compression required to put it on the Web and then take it out to your PC later on and decompress it -- you start looking at this, it requires eight times the kind of computer horsepower to do high-def versus regular definition. As high-def becomes more prevalent, processing powers will move back up that far to the left, and we'll have to start all over again.

It's not a stretch to say, though, that at the same time that the world is consuming more and more application capability in our day-to-day lives, we're also becoming more and more mobile. We've seen an increased focus on energy efficiency over the last four or five years. We all, in our notebooks, want to have longer battery life. We want to have smaller machines, sleeker form factors. Intel first talked about this shift to energy-efficient performance at IDF in 2001. We said then our goal was going to be to continue to increase performance while driving power consumption down over time. It's not a stretch, I think, to say that this is the biggest single shift in computer microarchitecture we've seen in well over a decade.

We started delivering this capability in terms of energy-efficient performance where it mattered most -- in the notebook with the Pentium M in 2003 in Centrino. In January of this year we significantly increased the capability of our mobile offerings by introducing the Core Duo Chip. It's the first dual core microprocessor for notebooks. At the same time we did that we also started seeing that same chip evolve and move into smaller form factor desktops, but also start moving into ultra-dense low-power servers.

In July of this year, just a couple of months ago, we introduced the next phase of energy efficient performance with the Core 2 Duo processor. The most exciting thing about this shift to energy-efficient performance is the innovation it's enabled from you in the industry to take advantage of these capabilities. A company that clearly knows about innovation from the inside out is Apple. I thought it would be very interesting to have them join us today. In fact, this marks the first ever appearance at an Intel Developers Forum of an Apple representative, so let me bring out Phil Schiller, senior vice president of marketing for Apple.

[Applause]

Phil Schiller: Good morning. Thank you very much.

Paul Otellini: Take it away.

Phil Schiller: Thank you for inviting me. Good morning, everyone. I am so happy to be here today because just recently Apple and Intel achieved a major milestone in our industry. Just recently we completed a transition of the entire Mac product line to Intel microprocessors. I think when we look back on 2006 this will be seen as one of the top tech stories of the entire year. It is a major accomplishment and something that only could have come about because of the incredible teamwork of Apple and Intel engineers working together and because of the power and low-power consumption of these new microprocessors. For that, we really thank you, Paul, and thank everyone at Intel. It has

enabled some incredible new products.

For example, we started this year with Paul and Steve together at Mac World introducing a new iMac from Apple – an iMac with an Intel Core Duo inside. For those of you looking at the slide who don't know the Mac product line that well, you probably think this is just a monitor. It is not. We have one running right up here. It is an entire computer. It looks like a monitor, but the computer has disappeared. It's all built in -- everything you need for a great personal computer experience. Not only is it a great experience, but a very fast one too because of the Intel Core Duo built inside it.

In February we shipped our first Notebook with an Intel Core Duo in it -- our Professional MacBook Pro. The MacBook Pro is an incredible, high-performance, professional notebook. Unlike everyone else's notebooks, we make them thin and light. It's 1 inch thin, 5.6 pounds, but has all the power customers need to do professional video editing, to do music creation, advanced imaging applications, all in one sleek form factor.

In March we did something that really surprised people. We took our small form factor desktop, the Mac Mini. This is just 6.5 inches square, 2 inches thick, and inside it, we put an Intel Core Duo. So an affordable, entry-level desktop, small form factor, but amazing performance and everything you need for a great desktop computing experience.

Also in March, we introduced our whole new consumer notebook line, the MacBook. The MacBook also has all Core Duo chips, dual-core performance across the consumer notebook line. It as well, is just an inch thick, 5.2 pounds, and has been a big hit with customers. At this point we have all moved all our notebooks in March over to Intel microprocessors. All Core Duo -- faster transition than anyone else in the industry, and it has really shown in our success. We announced not long ago that our market share of notebooks had risen from 6 percent in US retail in January to 12 percent this summer, a huge jump due to these incredible new notebook products that we made together.

Well, that left just one more product to complete our hardware transition, and that's our professional desktop. In August, we introduced and shipped the Mac Pro. This is a professional workstation. And so for this product, what we did is we put in two Xeon Dual-Core chips. This delivers an incredible amount of performance; it is the fastest Mac that we have ever made. But not only is it fast, unlike the previous generation of product it replaced, it didn't need liquid cooling inside. That saved a tremendous amount of cost, so we've gotten very aggressive on the pricing, more aggressive than our competitors. But we also provided new features inside because of the space we saved. We now added up to four internal hard drives, two internal optical drives, four full-length PCI Express slots, and an amazing amount of features for customers because even here in the workstation space, performance per watt really mattered to deliver a better product.

So far, I'm talking about hardware, but there's also been a huge transition in software as well. At the same time, we've brought over our entire operating system, Mac OS X, and all the software that runs on it native to these great Intel microprocessors. And this is really where the magic of the relationship can shine. Because Apple can move very

quickly to deliver new software features to take advantage of all this great hardware a lot faster than PC customers are used to. For example, we can take advantage of things like the latest SSC instructions for media applications very quickly in Mac OS X. We can take advantage of the new power-saving modes of these great microprocessors to deliver better battery life and thinner and lighter designs in our notebooks. And we can take advantage of the 64-bit support of these processors, like the Xeon, across our operating system without giving up 32-bit compatibility, something that so many others have struggled with. This is the transition that we've just completed less than a month ago. And with this transition, really it's just the beginning of the amazing things that Apple and Intel can do together to really bring to market the best computers that we, our customers, and our developers have ever seen. Thank you.

[Applause]

Paul Otellini: Thank you, Phil. It's been a real pleasure working with you and your team. Thank you for all your support. Thank you.

Now, Phil talked about converting the Apple product line to the Core Duo chips. The thing that has been really since then -- and I would hope that Apple will embrace -- is the Core 2 Duo chips, which we launched in July. And from our perspective, this is the most significant leap forward in terms of our performance, and particularly energy-efficient performance, and it was introduced just 60 days ago. In an era where 5 percent or 10 percent increments in performance seem normal, the Core 2 Duo was a stunning leap ahead. Let me give you some benchmarks. This happens to be from PC Magazine: 80 percent better in game performance, almost 40 percent better in audio, 35 percent better in video encoding, and most importantly of all, it does all this while consuming 20 percent lower power. It's jaw-dropping fast, in their words. The momentum for this product is really quite, quite good. Demand has been very strong. In fact, Intel has shipped over five million Core 2 Duo processors in the 60 days that the product has now been released. It's the fastest 60-day ramp for both the desktop and notebook processors in Intel's history.

We've had a very busy summer. We've launched the core architecture now across a number of segments, a single microarchitecture that spans everything from the sub-notebook to the DP server today. This architecture delivers the world's best and most energy-efficient performance. And it's not just a little bit better; it's anywhere from 35 to 80 percent better performance with better energy efficiency.

So what could be better than what we've done already? What's next? What's next is after Dual-Core, Quad-Core. In July we announced that we are going to accelerate the schedule of our Quad-Core products from Q1 of '07 into Q4 of '06. And I wanted to give you an update on that product line today.

The initial Quad-Core products will be targeted for the enthusiast or the gamer segments of the marketplace. This becomes the first Quad-Core chip in the industry that

will go into and is designed for a personal computer. The initial desktop version of the Quad microprocessor will be introduced as a member of our Xtreme family so it will be a Core 2 Xtreme; that's the brand name. And the product will ship in November, two months from now. And in Q1 of next year, as we bring on more capacity for these products and different price points, we'll bring it down and we'll bring it into the mainstream under the brand name of Core 2 Quad.

The early reviews of the Quad-Core product have really been very, very good. My favorite is from Tom's hardware page here, "It's like being catapulted a year into the future." Now as I showed you earlier, the Core 2 Duo, which was introduced just 60 days ago, was a big leap forward for us and for the industry. And in fact, the Core 2 Xtreme product today is the fastest, single-socket microprocessor on the planet. The Quad-Core version is 70 percent faster. This is a huge step forward for the second time inside of a four-month window. One area where Quad-Core will be extremely valuable, I think, is gaming. And to tell and show you more about how this works in gaming environment, I'd like to invite out Markus Maki, who is the founder and chairman of the game developer, Remedy. Markus?

[Applause]

Paul Otellini: Can you tell us a little bit about Remedy?

Markus Maki: So Remedy founded in 1995, and we're a small game developer, and we're probably best known for our action blockbuster titles, Max Payne and Max Payne 2. And now our team's been working since then on Alan Wake, our latest game. It's been 18 months since we have showed that to anybody here, and our team's been hard at work making it work and look great on Intel Multi-Core Architecture.

Paul Otellini: Perhaps you can tell us a little about what environment you've got it in here?

Markus Maki: Sure. So we're going to run the game on this beautiful system from Falcon Northwest. It's running on an NVIDIA graphic solution and it has an Intel Core 2 Quad processor that's been [over-clocked] by more than a gigahertz. So it's running at 3.73 gigahertz.

Paul Otellini: Over-clocking is allowed here, huh?

Markus Maki: Yeah, it's allowed here.

Paul Otellini: This must be one of the fastest gaming machines on the planet.

Markus Maki: It sure is, and it's been really great for us. So let's see what we've been up to. What we are building here is a psychological action thriller. Alan Wake is a writer who goes to a small, quiet town to heal his insomnia, but his writing and his nightmares start to turn real. Let's see how that's working out.

The game is set in the Pacific Northwest, and it's a beautiful and diverse location for its nice outdoor environment for our outdoor technology.

[Music]

Markus Maki: So let's take a look at the technology we have here. What we have is incredibly detailed [in-game actors] in a huge, seamless and dynamic game world. This kind of game simply couldn't be done on a single-core processor. We are using the multiple cores here to stream data in seamlessly in the background and also we're using the [CP] of course to prepare data for the [GPU surrender]. And this allows us to deliver a dynamic and versatile environment for the player. Let's see how our world simulation works and [tackle] the simulation [parameters] in real time.

So the game supports everything from this beautiful sunrise to a crisp, clear autumn day. And of course we support the more dark and gloomy, stormy settings as well. What we are building here is a psychological action trailer. And this kind of technology is perfect for us, perfect for building the mood and the atmosphere for a game.

So let's go back to Mr. Wake here and transition to nighttime to see how the lighting in our game works. All of the lights and shadows in the game are completely and absolutely dynamic and volumetric. This level of lighting has never been delivered in a large-scale outdoor game before. But, okay, I'm feeling a bit confined here, so let's go and seek some wide-open spaces to really demonstrate what we can do with four CPU cores. So we're now on the other side of the Cascade Range, and let's unleash the forces of nature, or the forces of the Intel Core 2 Quad processor on this nice and quiet location here. See what happens.

[Game sounds]

So here we're using a whole CPU core just for the [physical] simulation alone. You simply couldn't deliver this level of [physical] simulation along with the rest of the game experience on a single-core processor. And as a game developer, we're really excited on how far we're now able to push the in-game simulation and create more dynamic, versatile experiences for the player.

Paul Otellini: [And we're very excited to get more and more of this.] Wow. Don't let him get hurt.

Markus Maki: All right.

Paul Otellini: Thank you, Markus. It is truly extraordinary work.

Markus Maki: Thank you, Paul.

Paul Otellini: I think it's like watching a movie.

Markus Maki: Thank you.

Paul Otellini: Yeah. We appreciate it very much, thank you. That was really extraordinary. There's lots of excitement from a number of our partners as we're about to bring out these Quad-Core systems. Starting with the PC manufacturers who are going to announce and launch systems around it. There are 13 up there that have indicated so far that they'll launch with us in November. We hope to see more. There are lots of these Quad-Core systems out on the floor at IDF over the next four days for you to be able to play with. I suspect there's not a demo of that game, but there are other games out there. Quad-Core inside, coming in November.

Now, at the most fundamental level, the thing that makes all this work, the thing that makes this come alive is the transistor. And a lot has changed in the world of computing, but there's been one constant, and that's Moore's Law. As many of you may know that Gordon has retired from day-to-day work at Intel, but the responsibility in terms of carrying on his law has been passed on. I'd like you to meet some of the stewards of Moore's Law.

[Video plays.]

Male Voice: We hear a lot about Moore's Law, but I don't think that most people are aware of what doubling every two years really means.

Male Voice: When I first joined Intel, the smallest transistor dimensions were about 3 microns. Today the smallest dimensions are around 30 nanometers, 100 times smaller. Over the years, our process technology has changed dramatically. Our wafer sizes have grown from a 3-inch diameter in 1978 to today's 12-inch diameter, and we have switched from bipolar to NMOS to CMOS technology, and more recent innovations have been [locate dielectrics], copper interconnects, and [strained] silicon transistors. In each case, the goal was to improve performance and to reduce power. So while the concept of performance-per-watt is somewhat new to the industry, power considerations have been a focus for Intel all along.

Male Voice: At Intel we have a huge advantage over our competitors because we design and make our masks in-house. Our mask teams collaborate closely with our project teams to develop resolution, enhanced capability, to deliver higher-density] [patterning earlier and with lower cost].

Male Voice: Moore's Law isn't like a natural law. It doesn’t just happen. We began working on 65-nanometer over five years ago. And are now working on processes for use five years or more from now. Gordon Moore formulated a law. Now it's up to us to keep him honest.

Paul Otellini: Last year, Intel introduced the world's first 65-nanometer transistors. This was a significant improvement in performance-per-watt. These transistors deliver up to 30 percent better performance, a five-fold reduction in leakage at the transistor level from the previous generation of 90 nanometers, and 30 percent smaller. In fact, up to 10 million transistors can now fit on the tip of your ballpoint pen.

So how are we doing with 65 nanometers? This month for the first time Intel's microprocessor shipments on 65 nanometers are great than those on 90. We've crossed over in terms of our shipments from 90 to 65 nanometers. This week we will have shipped our 40 millionth 65-nanometer microprocessor.

For those of you out there who might be keeping score -- and I suspect at least one guy is -- the rest of the industry has shipped a grand total of zero in terms of production microprocessors. For you trivia buffs out there, though, even more important, these 40 million microprocessors represent something in the range of eight quadrillion transistors. That's 8,000 trillion 65-nanometer transistors. I've been assured that Justin Rattner, our CTO, has done the count on this, so it's absolutely verified.

What's next? We don't stop. We're the stewards of Moore's Law. I thought I'd talk a little bit about 45 nanometers. We have a major transition coming this year. Forty-five nanometers will deliver an additional 20 percent performance boost and up to another five-fold reduction in leakage power at the transistor level. Intel was the first company to demonstrate 45-nanometer SRAMs working earlier this year and right now the process yields for this generation of technology are tracking right on par with the 65-nanometer technology at its point in the development cycle.

I thought I'd show you a picture of the world's first and only 45-nanometer production facility. This is called D1D. It's in Oregon. For those of you wanting one, it costs about $3 billion. It's a 212,000 square feet of world-class production facility. This facility is running test wafers now. It runs production wafers scheduled for the second half of 2007.

But we're Intel. We can't rely on a single Fab to meet your needs and the needs of the marketplace. I thought I'd give you a sneak preview of the next two 45-nanometer Fabs that are under construction. These are Webcams of the next two Fabs -- one is Fab 32, which is in Arizona, and the other is Fab 28, which is in Israel. Both are well along. Actually, the clean room is finished on this one over here in Arizona. This is the gown room that's being built on the construction on the outside. We're starting to install equipment on this one. This is a little bit farther behind. We'll have production in the Arizona Fab in the second half of next year and in the first half of '08 in the Israeli Fab. In grand total, we'll have three, very large, world-class Fabs up and running on 45 nanometers. The grand total investment for these three Fabs is $9 billion. The sum total of the clean room space is a half million square feet of clean room.

You might be asking yourself, "What the heck are they going to do with all that space?" Well, we're going to build products. Most importantly, we're going to start with

microprocessors. Our first 45-nanometer microprocessor design is almost complete. In fact, the design will be completed next quarter. Beyond that, we have a total of 15 distinct, discreet microprocessor designs being undertaken today at Intel on 45 nanometers, meeting the needs of the desktop, notebook, and server market segments. We'll roll these products out over the life of this technology.

But it's not just 45 nanometers. We have a model now for sustained technology leadership. We will, of course, continue to run and deploy new silicon technology on the two-year cycles you've come to expect from Intel. Here I've shown the 32-nanometer cycle coming in the 2009 timeframe. But we're augmenting those with a new microarchitecture every two years as well. The Core microarchitecture came out this year. In '08 we have [Nehalem], the next generation microarchitecture. In '10, [Gesher], the next generation after that. What you can get from Intel, what we will commit to you for these three generations is a 300 percent increase in performance-per-watt between now and the end of the decade. We're not going to slow down on Moore's Law and we've got the design capability to make it happen.

Performance and performance-per-watt are perhaps most important in the data center. Back in 1996, 10 years ago, this was the fastest supercomputer in the world. It was at Sandia National Labs. It was the world's first teraflop supercomputer. It was called ASCII Red and inside of that product was just about 10,000 Pentium Pro microprocessors in 85 really large cabinets that took up about 2,000 square feet and cost about $50 million. We're 10 years into that and we've been deploying that product and we thought, well, what's better than a teraflop computer in a room? Obviously, a teraflop computer on a chip.

What this is a production prototype of the world's first teraflop on a chip die. Each of these dies has 80 simple floating-point cores on them with stacked SRAM. There are over 80 die on this wafer. Each 80-core die is capable of a teraflop performance with the ability to transfer terabytes-per-second of data between the cores and the stacked SRAM. Now it's a prototype, but we think that the capability embodied by this prototype chip is going to be commercially available in a five-year window. That allows us to do some very amazing things. Justin is going to talk more about this in his speech a little later on, but this kind of performance gives us, for the first time, capability to imagine things like real-time video search or real-time speech translation from one language to another. What's inside matters.

Let me shift gears now and talk a little bit about markets. And let me start first with the enterprise. The dominant trend in data centers today, as you've heard about and read about, is energy efficiency. In fact, it's one of the reasons why you've seen some of the three largest data center operators in the world, Yahoo, Google, and Microsoft, all announce that their next-generation data centers are going to be built in the Pacific Northwest, because that's the cheapest electricity available in the United States today. So continued focus on energy-efficient performance is now driving people to consider sites for locations that they wouldn't have considered before.

In terms of delivering energy-efficient performance, though, Intel had a rather large leap ahead in June when we started shipping the industry's leading performance and performance-per-watt dual-core server chip, the Intel Xeon 5100. It's based upon the new core microarchitecture. This product has also had tremendous momentum, much like the Core 2 Duo did that I talked about a few minutes ago. As of today, three months into the shipment curve, we've now shipped over one million of these products. It's the fastest ramp of any server chip in the history of our corporation. In its first three months of shipment, the Xeon 5100 accounts for half of all of Intel DP shipments and almost 40 percent of the entire DP marketplace -- in its first three months. The product is moving very, very nicely. We're very pleased with it.

But just as with desktops, the next evolutionary step in servers is also going to be quad-core. In November, we will introduce a quad-core Xeon, the first quad-core CPU in the industry for the high-volume market segments. And beginning early next year, we'll introduce a lower-voltage version of that quad-core down at 50 watts, a 40 percent reduction from the standard voltage part that we'll introduce in November.

The quad-core Xeon will deliver 50 percent improvement over today's core microarchitecture -- that's the dual-core Xeon 5100 I talked about -- in the same power envelope. That is, it goes into the same socket, in the same chassis, in the same power envelope, at a 50 percent improvement in performance.

Much has been written in the last year about Intel losing its momentum, losing its leadership in the server market space. I believe very much that with this new set of dual- and quad-core microprocessors we've now regained our leadership. And these products are gaining traction across the customer base and around the world.

To talk to you about why customers are now adopting these products, I'd like to invite out someone who knows a little bit about this area, Tom Barton, who's the CEO of Rackable Systems.

[Applause]

Tom Barton: Hi, Paul.

Paul Otellini: Hi, Tom. Good morning. How are you?

Tom Barton: Good, how are you?

Paul Otellini: Good. Can you tell us a little bit about Rackable?

Tom Barton: Sure. Rackable Systems is the industry's fastest growing x86 server and storage vendor, and we're very focused on large-scale deployments. To give you a sense for what I mean by that, for a single customer in a single data center, we've actually deployed 10,000 servers based on Intel's Xeon Sossaman platform. And when you're focused on big deployments like that, customers really care about power consumption,

they care about performance-per-watt, and we as a company, are always looking to deliver more performance at a lower power consumption in a smaller footprint.

Paul Otellini: Intel launched the dual-core, lower-power microprocessor about three months ago. How has this impacted your business?

Tom Barton: Well we really appreciate Intel's renewed commitment on performance-per-watt. Like I said, that's very important to our customers. And based on the improvements that we're seeing, we're actually expecting that the sales of our Intel-based servers are going to grow by a factor of five in the first quarter of '07 versus the first quarter of '06. It's clear to us that Intel's got leadership now in performance-per-watt for the Xeon Woodcrest Series 5100. But performance is great, and equally stunning are the improvements in energy efficiency.

Paul Otellini: That's great, Tom. I see you brought a little something with you.

Tom Barton: We did.

Paul Otellini: Maybe you can show it to us?

Tom Barton: Yeah, we brought one of our racks here. This is a 22U rack, and using our patented approach to density, we actually mount 40 servers back-to-back here and we've still got room left over for networking gear. We're also showcasing our DC powered option that leads to greater energy efficiency as well.

Paul Otellini: So if I'm doing my math right, 40 servers, dual-core, dual processor, that's about 160 cores at work here?

Tom Barton: Well you're half right, Paul, because we actually built this with the Clovertown Quad-Core, so what you're looking at here is 320 processing cores, which is a world's record for X86 server density. There's actually enough processing power here to generate three teraflops and put it in the world's top 300 supercomputers all in this 22U rack.

Paul Otellini: Wow.

[Applause]

Paul Otellini: But as they say on TV, wait, there's more, right?

Tom Barton: Right.

Paul Otellini: I understand that this is also available in a full configuration which is roughly twice this size, over 600 core teraflops, and that would put it in the top 80 supercomputers of the world.

Tom Barton: That's right, in a single rack.

Paul Otellini: Fantastic, unbelievable. Thank you so much for sharing it with us today.

Tom Barton: Thanks, Paul.

[Applause]

Paul Otellini: Now let me shift gears and tell you a little bit about the home now. The home's seen a tremendous amount of change over the last few years as well. In the last 12 months, perhaps the biggest change is the explosion of web-based video coming into the home. I plotted here YouTube. YouTube has now grown from almost nothing to over one billion page views per day. And YouTube, as big as it is, still only represents less than 10 percent of the total video streams that are on the web today.

As this continues to grow and as more and more content moves onto the web, I think it's giving rise to a new class of devices. And the one thing that this new class of devices has in common is that they connect the Internet to the television. And there are a number of them up here today. Over here you've got the IP set-top boxes, HD DVD players, digital media recorders are all here today and have one thing in common; they have Intel architecture inside them. In fact, the Toshiba HD DVD, which is on the first machine to the right of the Rackable, the big rack over there, the black machine, is the second generation of the Toshiba HD DVD player. This is based upon Intel architecture inside it. It ships in October, and it hits a new price point for HD DVD players, sub-$500.

So you have to ask yourself, why are people in consumer electronics putting Intel architecture inside this wide array of devices? There are really two reasons: First of all, it turns out that many of the Internet services and access capabilities that we expect to have on the PC is most easily available and most easily ported quickly to CE space by using Intel architecture. It's the availability of the common software stacks that are written around the PC to make this a very easy design paradigm. But equally important is the need for performance inside these devices. And the ever-increasing need for performance naturally leads them to products that are made for the PC, like Intel architecture.

Let me move beyond the CE devices and talk a little bit about the evolution of the PC. In January, we launched a product called Viiv, and this is Intel's first thrust around the media PC. Viiv is a collection of technology designed to offer the best media experience ever. It's a place where you aggregate all different forms of your content that's out there from the Internet to your personal content. It's also a device that allows you to connect to other devices around the home to view that content on a larger screen in a more friendly, most-often, couch-oriented fashion.

I wanted to give you an update on Viiv. Since January we've done two big things: We've incorporated the Core 2 Duo into the product line. And just recently, in the last month or so, we've released a new software stack that enables the Intel media server, which allows consumers to be able to access Viiv content from the comfort of their TVs

in their living rooms or other rooms around the house. To make this work though, we've had to work with the industry to enable a series of devices that allows for this to happen in legacy televisions. People aren't going to go buy a new television just to take advantage of this. So we needed a device that interfaces between the Viiv PC and those TVs. That device is called a DMA or a Digital Media Adaptor, and we're making some very good progress on here. We've had specifications out for well over a year. A number of manufacturers are now moving into production. The first out there today is one from a company called NETGEAR. Their DMA retails for around $250; it's available starting this week.

But you can see as we move into the peak selling season here, there are a number of other companies that will have devices on the shelf from Acer, Linksys, Creative, D-Link, ICube and Buffalo, all supporting Viiv, all verified to work with Viiv and taking advantage of this 10-foot user model that's coming out today.

Now there's one other device that is very common and it sits next to most of our televisions, which is a set-top box. Wouldn't it be great if the set-top box that we already have acted as a DMA as well? Well earlier this year, at CES, we talked about a relationship between Intel and DIRECTV. And today I'm pleased to show you the first DIRECTV product that's verified to work with Viiv. It's DIRECTV Plus HD DVR and it's verified. It has incorporated DMA functionality, so it's that silver box over there. It has standard set-top box functionality with the conditional access but also able to now take content from the Viiv PC, either from the Internet or your personal content, and show it on that same television that the set-top box is connected to.

The DMA functionality inside that box is in its final stages of verification, and it will be delivered to customers in the fourth quarter. DIRECTV expects to sell millions of these devices over the next year. In addition to this device, DIRECTV is also developing an add-in tuner card that will plug into Viiv PCs to allow for conditional access-based content to be viewed on the Viiv PC. So it works in the other direction so you can turn you PC or even your notebook into a set-top box receiver using conditional access with the addition of this add-in card -- bi-directional, very powerful usage model.

Now all of this is meaningless without content. Content is making very, very good progress. In January when we launched Viiv, we had 50 verified services. Each service had a multitude of offerings underneath it. And it had things from sports to TV to games. As of today, we have over 150 verified services, many of which are unique to Viiv. We have over three million music tracks, over 100,000 music videos, over 12,000 television shows, 1,000 hi-def movies and hundreds of games.

But in the last three weeks there have been four announcements that I wanted to call to your attention. The first is with Yahoo -- Yahoo Fantasy Sports. I know many of you may be fantasy football players. There are over six million fantasy sports players in the world today. Today to be able to do this, you need to have a notebook or a desktop computer connected to the Internet and a TV to watch the broadcast. This new service from Yahoo, Fantasy Sports, allows the user to be able to integrate those two and to be

able to have that information associated with the fantasy sport of your choice up and interactive, available to you from your couch while you're watching the game in the other screen on your television. It's available only for Viiv for some time.

The next one is from AOL. AOL has started shipping and made live their 10-foot video interface, their online library. It's the world's largest library of online premium movies, full-length TV shows, music videos, and music concerts. The way you access this database is, again, from your couch with your remote control. This is the industry's first 10-foot video search capability. It was created and verified on Viiv.

Later this week, NBC will launch a new service. It includes free downloads of several new NBC shows that are exclusive to Viiv and Centrino users. In fact two of the new shows that they've committed to are Heroes and Studio 60 on Sunset Strip. We expect this to be a very popular service as well.

And lastly ClickStar. We talked about ClickStar at CES. This is a company which is going to bring to market the first ever co-release of a Hollywood movie, a new premiere movie, and the Internet release. The movie's called *10 Items or Less*. It stars Morgan Freeman; there's a picture there. They've committed to release this onto the Internet within two weeks of theatre release, so it's now compressing the time between theatre and the time it's available on the Internet. This couldn't happen without this open-standards approach that's creating a very large amount of content from a variety of suppliers using different media sources, different DRMs, different data types, all coming together under one industry standard around the Viiv business model.

So content's changing dramatically, but the desktop is not evolving nearly as rapidly. There are some examples shown here where OEMs are incorporating new Core 2 Duo technology into ever-smaller form factors. If you think about the way we live, though, in our living rooms, in our dens, in our homes, we care a lot about how things look and about how they sound. We don't want them noisy. So today we think that we need to accelerate the pace of adoption and technology creativity around these small-form-factor, fanless, noiseless devices.

So today we're announcing the Core 2 challenge. This is a $1 million prize or set of prizes. It's a design competition, and it's intended to create the sleekest, smallest PCs based on Viiv and Core 2 Duo. It's open to all OEMs and ODMs. There'll be objective judging. I'm one of the judges. I'm reasonably objective. And the winner will be announced at next year's IDF. The details on this challenge are available out in the hall or on Intel.com.

Let me shift next to mobility. We've talked a lot about how energy-efficient performance is changing the enterprise and the home, but the largest benefit for energy-efficient performance is in the things that we carry around with us, things that are mobile. We've seen a big shift from the luggable computer of the '80s to the sleek notebooks of today. You saw Phil's very attractive notebooks in his slides. The result of this, though, is that notebook growth is accelerating. There are more notebooks sold per year in Japan

than desktops. In the United States and in Europe there are more consumer notebooks sold than consumer desktops every year. This is up from less than 20 percent of the total market in the pre-Centrino days.

Centrino was an inflection point in driving notebook growth. In 2003, we introduced the first version of this product. It was a low-power microprocessor. It was very high performance, chipset, Wi-Fi; it ushered in this era of sleeker, more connectable notebooks. In 2004, we made another revision to that product line; we increased the second generation of Centrino in terms of microprocessor performance, better wireless, longer battery life. And then earlier this year we announced a significantly increased performance version with the Core Duo product and now the Core 2 Duo product, bringing dual-core performance into Centrino.

I'd like to talk to you about what we'll ship next year. Next year we release the next iteration of the Centrino platform. We've codenamed it Santa Rosa, and it has a number of significant improvements. It will use the recently released Core 2 Duo microprocessor, but to that we add Vista Arrow capable next-generation integrated graphics. We add a faster wireless with 80211.n. We add NAND Flash Memory onto the motherboard. And we have new manageability and security features for the enterprise. Let me spend a few minutes talking about these features.

The platform that launches early next year will have NAND on the motherboard for the first time that we know of out there. This NAND on the motherboard significantly improves Notebook performance and battery life. Essentially what the NAND cache does is that frequently accessed data will be stored on the NAND Flash rather than the hard drive. This means it's more available to the main memory and to the microprocessor. And since it doesn't have to go out to the drive as often, it consumes less battery life under notebook mode. More importantly, because it's closer and faster in semiconductor memory, you have faster boot times, 2x faster application load times, and a 2x decrease in the time you resume from hibernation. All this comes along with reduced power consumption.

80211.n is the next generation of the Wi-Fi 802 standard. Today, most notebooks ship with an AG wireless card or chip in them. The data rate on that AG card is 54 megabits per second. 11.n has more than five times the data rate and more than five times the throughput rate of AG. This gives you stunning new capabilities. You can have multiple video streams; you can actually have multiple high-def video streams wirelessly now being enabled on this new class of notebooks. You'll see more of this tomorrow in Dadi Perlmutter's speech, and he'll actually do a demo of these capabilities.

The next inflection point, though, for mobility is really what I call "Broadband to Go." What do I mean by that? I mean by that that we have personal, portable broadband Internet available to us as individuals wherever we're going. To make it a reality we need two things: We need pervasive wireless broadband, and we need a new class of ultra-mobile computers that we can literally carry around in our hand.

Let me start with WiMAX. We've talked about WiMAX a number of times in this forum. In the last year we've seen a huge leap forward in the deployment and the commitment to WiMAX. Two big things happened. First of all, the 80216.e spec was approved by the IEEE. 16.e is the mobile version of the WiMAX standard. Secondly, we've now seen major telecom operators step forward and commit to deploy WiMAX networks in many countries around the world. I've show two here in the United States, which are Sprint and Clearwire, who have announced in the last 60 days.

I like Sprint's view of the value proposition of WiMAX. They say, "Why are we deploying WiMAX? It's four times faster and one-tenth the cost of EVDO." That's a pretty good value proposition for them and for us as users. The combined coverage between these two companies that's been committed is almost nationwide coverage -- over 100 million users covered by 2008. It's not just the United States; WiMAX momentum is very strong, with over 200 WiMAX trials around the world and a large number of deployments.

How are we going to bring it to computers? Mobile WiMAX will first be available in the form of add-in cards for notebooks later this year and throughout '07. To date, eight telecom operators around the world have committed to incorporate the capability of these cards into their networks in 2007. In '08 WiMAX is integrated onto a single module with Wi-Fi and becomes essentially part of the Centrino platform. If you think back, it was integration into the platform of Wi-Fi that drove the ubiquity of Wi-Fi. We believe the same thing is going to happen as we integrate WiMAX into the platform and take advantage of that ever-expanding WiMAX coverage that will be deployed around the world.

So that was the vector of broadband wireless. What about the devices? How are we going to access this data as we're on the go? Well, on the computing front I talked last year about a new class of ultra-small, ultra-portable devices. I said that in order to create them, Intel had to create a new product line. And that we would create silicon devices that would take the power down by a factor of 10 in order to enable very long battery life. We said we would enable this by the end of this decade. I wanted to give you an update on our plans now. I'm happy to say we're accelerating our plans.

If you take 2006 as the baseline, this year the baseline is a five-watt, ultra-low-voltage microprocessor. Next year, we'll ship one that has one-half the power and takes one-fourth the surface space in terms of driving the form factor density down. In 2008, we'll meet our decade goal of a 10x reduction in power and we'll do this at one-seventh the size of today's offerings. And we have plans to move beyond that to single-chip integration beyond 2008.

If you compare the product we're shipping this year, in '06, to the product we'll ship in '08, the battery life triples, from three hours to greater than seven. So the idea of getting essentially all-day life out of a very highly-portable, highly-functional, highly-featured device is now very much upon us. All of these devices will offer full web capability and have the capability of running multiple OSs, including full versions of

Vista. We think that this offers a new opportunity for all of you, as developers, in the audience to take advantage of this silicon development and find new opportunities for products that come out there. To talk more about this, I'd like to bring out Anand Chandrasekher, who's the general manager of the new Ultra Mobile Group at Intel. Anand?

[Applause]

Anand Chandrasekher: Hi, Paul.

Paul Otellini: Hi, how are you?

Anand Chandrasekher: Good to see you.

Paul Otellini: So what have you got to show us today?

Anand Chandrasekher: So Paul, this is a brand new piece of a prototype. It's based on our ultra-mobile silicon that is going to come out in the first half of next year. And this is basically using Intel's lowest power consuming processors to date. This uses the [Steely] processor, which we're sampling to customers today and will be out in the marketplace next year. This particular prototype is built by Quanta. It has a full PC operating system that it's running. The user interface is a user interface delivered by Street Tech. It's got a built-in camera. It's got a pretty cool user interface, as you can see, a keyboard, if I just rotate this around. It's got a full keyboard, so you can actually use that. It's got Wi-Fi and WiMAX built in, and I'll show you the WiMAX piece in a little bit. But first would you like to see what this is capable of doing?

Paul Otellini: Well it looks great. I'd like to see if it does anything.

Anand Chandrasekher: [Laughs] Ah, ye of little faith. All right, let's see what this can do. Want to see a video?

Paul Otellini: Sure.

Anand Chandrasekher: All right, let me click on this. And I'll pick *Short Circuit*. I've got three videos loaded here. I'll pick *Short Circuit*, which is one my kids like. So . . . so hopefully the camera's picking that up, right. Now the Street Tech user interface is also pretty cool. If I want to pause this, all I do is run a line through it, and it pauses it. If I want to skip to the next video, I run a line that way, and the next video's playing, right? Now let me go back to the homepage and see if we can find some coffee places to go to. So I write H here. I'm back to the homepage. I'll pause the video here. Okay? All right. Shall we go find a Starbucks Café?

Paul Otellini: Sounds good.

Anand Chandrasekher: All right, let's see if we can find one. Coffee icon shows up

here. So this is actually searching the web live, and it's given me a couple of options. Right?

Paul Otellini: Okay.

Anand Chandrasekher: So I'm going to pick the Starbucks in San Francisco. It says it's 0.2 miles from here. Let's go to that. Let's view the location. Okay, so it tells me where it's at. Ready to go Paul?

Paul Otellini: Yeah, sure.

Anand Chandrasekher: I can get my car, we can ready to go.

Paul Otellini: Car, what car?

Anand Chandrasekher: Funny you should ask. Let's take a look.

[music/video plays]

Anand Chandrasekher: Hi, [Farad]. Good to see you.

Farad James: Good morning.

Paul Otellini: Good morning.

Anand Chandrasekher: Paul, this is [Farad James]. He's the project manager at Volkswagen, and Intel and Volkswagen have been working effectively on standards-based and communications, right? That allows us to let devices like this interact seamlessly with this car and in his in-car entertainment system. So let me show you what I mean by that, right? So let me go right back -- you saw this earlier -- we're on the location. Go back to the homepage, and whoops -- go back to the homepage. And what I'm going to do is I'm going to connect the UMPC to the car, okay? So let's do that first.

Mechanical Voice: UMPC connected.

Anand Chandrasekher: So they're now connected.

Paul Otellini: Wi-Fi?

Anand Chandrasekher: That's right. Over Wi-Fi they're now connected. Now what I'll do is I'll send the point of interest information details over to his navigation system, so he picks up Starbucks here. Let's send that over.

Mechanical Voice: A point of interest called Starbucks Coffee has been sent to you. Would you like to route to it now?

Anand Chandrasekher: So he has the same information that we just picked now up there, right?

Farad James: Yeah.

Anand Chandrasekher: Okay, now this has got to be able to do some more stuff for us.

Farad James: Of course. One of the features I want to show you is online music access in the vehicle.

Anand Chandrasekher: Okay.

Farad James: So thanks to the broadband access I now have from the UMPC, I can listen to thousands of Internet radio stations. So let me show that to you.

Anand Chandrasekher: So before you pick the music itself -- what he's doing now is using the WiMAX capabilities that this has built into it. We're actually using a clear WiMAX CPE to service this entire demo. And we're going to stream music from this, using Yahoo Music, over to his system.

Farad James: Yahoo Music is one of my favorites, because it actually allows me to customize my own music stations. It also allows me to rate music, so if I don't like something, I can just tell the system not to play that song again.

Anand Chandrasekher: Perfect.

Farad James: I will just play a song for you.

Anand Chandrasekher: Okay, go ahead. You should be able to hear that.

Farad James: It's kind of a slow song. It's a symphony. I really like it, though.

Anand Chandrasekher: So this one you can program up.

Farad James: Right. So if I don't like it, I can just tell the system not to play it again and it will not play again.

Anand Chandrasekher: Okay.

Farad James: And, of course, even when I don't have Internet access, I still have access to all the music content on my UMP, so I can stream content from that. So as you can see here, I have access to all my music collection and I can just play that back in the car.

[Music]

Anand Chandrasekher: So that, instead of going to the web, is picking it up off of here?

Farad James: Yes. And of course, I am not the only one to be traveling in the car; I will certainly have people in the back. I cannot allow myself to watch videos, but I can definitely stream it to them.

Anand Chandrasekher: So if you have kids in the back, they could be watching videos?

Farad James: Yes, let me show that to you.

Anand Chandrasekher: All right.

Farad James: Here I will stream the *Short Circuit* video to the right screen, as you will see in a second, and I will stream a commercial to the left screen [here]. Do you see that?

Anand Chandrasekher: Yes, I see an ad here, and I see something else running there but I can't quite make it out. Yes, it's a movie.

Farad James: And this is all coming off from the UMPC and being streamed to the vehicle.

Anand Chandrasekher: Fantastic.

Farad James: I can control it from the front.

Anand Chandrasekher: And we have infrared headsets, too, so each child could be watching something completely different, no interference, and basically interoperability between the UMPC and the in-car entertainment.

Farad James: Right.

[Applause]

Paul Otellini: Thank you.

Farad James: Thank you.

Anand Chandrasekher: So Paul, today what you see is an announcement from VW and Intel which allows this kind of seamless interoperability between these devices and cars to be facilitated so that people who buy cars like this can enjoy this experience on the go and basically bring all of this to the masses. So what you see in total here is PC capabilities, full, rich power of the PC, and full connectivity using broadband wireless capability, effectively in your pocket.

Now today, arguably, your pocket has to be a little large, right? As we shrink our devices, as I showed you on the road map coming out in 2008, these devices are going to shrink. And this 5-inch screen here has 1K by 768 capability, so you're going to be able

to enjoy the full web experience without having to get glasses like the ones I'm wearing. So overall, great capability coming to a pocket near you soon and in the car near you.

Paul Otellini: Great. Thank you. We're really happy to be working with Volkswagen.

Farad James: Yeah, me, too. It's a pleasure.

Anand Chandrasekher: Thank you.

[Applause]

Paul Otellini: Pretty cool. So, we've talked a lot about the importance of power today, and at the end of the day, the real power of technology is its ability to transform lives.

Four months ago, Intel launched a program called "The World Ahead." This was a comprehensive and holistic program created to help developing nations around the world accelerate the deployment of technology and cross the digital divide. Intel committed a $1 billion investment into this program over a five-year period. We're focused on three things around PCs: Accessibility, connectivity, and integration into the educational systems around the world.

I thought I'd give you a status of where we are today. We've launched the program now in 30 countries around the world. We've worked deeply with governments to create government-assisted programs to make PC access to citizens, either through purchases or through kiosk capability, available to people who wouldn't ordinarily have PC access. We have trained over 500,000 teachers this year and a total of 3.6 million to date at no cost to the schools or the teachers, in terms of how to incorporate technology into their classroom curricula and make it a more powerful learning tool.

A lot of people out there talk about bringing technology to developing countries. At Intel, we're focused on action. When we launched the "World Ahead" program four months ago, I held up a prototype of this device. It's called the Classmate PC. It was designed for education work, classroom use. It's a full-function PC. It's got an ultra-mobile microprocessor inside it. It has a NAND hard drive for durability and for cost. It's small, it's low-cost, it's fully-featured, and it's going to be out in the first quarter of next year.

We've recently put that to a test, to a real-world test, halfway around the world in a classroom in Nigeria. In this classroom, there are 36 students. They had no previous Internet access in the classroom, 80 percent of the students had never touched a PC before. And Intel provided several weeks of teacher training. We installed a WiMAX network into the classroom and the city area, and we deployed dozens of Classmate PCs to have these students take advantage of them. Let's see how it worked.

[Video plays.]

Male Voice: I want to be a medical doctor.

Male Voice: I would like to be an engineer, a computer engineer.

Male Voice: My favorite subject is [science].

Male Voice: I want to become a medical doctor.

Male Voice: Some topics, it becomes very, you know, difficult for you to really figure out what you are going to give to the students so that they can learn. We can do better.

Male Voice: Hi, students.

Voices: Hi.

Male Voice: Good. She's your classmate, right? This is your classmate too. The Classmate PC is a PC designed by Intel to extend education to students via the computer.

Female Voice: Today I performed my first operation and it was successful. [So I'm happy.]

Male Voice: Something as small as that antenna can give the entire school coverage, because it will connect the school to the base station and provide Internet to the entire school.

Female Voice: Now we can talk to somebody that is living in England, here, standing in Nigeria, and get the feedback immediately. So I want to believe that they will see that as an opportunity and learn fervently with the Internet.

Male Voice: It will help us to send messages to other people in other countries.

Female Voice: I wrote my name in the computer, and just beautiful.

Female Voice: I did with the computer.

Male Voice: I want to believe that computer will seriously enlighten [them].

Male Voice: I wish for them to be better, you know, people in society. I wish for them better [health].

[Video ends.]

Paul Otellini: If we can inspire children in Abuja, we can inspire children all around the world. There's clearly a new world ahead, but it won't just happen. We have to come together collectively to make it happen. Technology is improving every element of our lives, in the enterprise, in the home, on the go, and we're bringing the benefits of this

technology to new groups of users a billion-plus strong around the world. In this quest, you can count on Intel doing its part, with the relentless pursuit and deployment of advanced technology. But that's not enough. This forum this week is all about us coming together as an industry to be able to make these ideas a reality. Thank you very much and enjoy the rest of the conference.

This presentation contains forward-looking statements that involve a number of risks and uncertainties. These statements do not reflect the potential impact of any mergers, acquisitions, divestitures, investments or other similar transactions that may be completed in the future. The information presented is accurate only as of today's date and will not be updated. In addition to any factors discussed in the presentation, the important factors that could cause actual results to differ materially include the following: Intel operates in intensely competitive industries that are characterized by a high percentage of costs that are fixed or difficult to reduce in the short term, significant pricing pressures, and product demand that is highly variable and difficult to forecast. Additionally, Intel is in the midst of a crossover to a new microarchitecture on 65nm process technology in all major product segments, and there could be execution issues associated with these changes, including product defects and errata along with lower than anticipated manufacturing yields. Product roadmap projections are affected by the timing and execution of the manufacturing ramp, product defects and errata (deviations from published specifications), changes in the timing of qualifying products for sale and changes in customer demand. For a variety of reasons, we may terminate product development before completion, or delay or decide not to manufacture and sell a developed product. For example, we may decide that the product might not be sufficiently competitive in the relevant market segment, or for technological or marketing reasons, we may decide to offer a different product instead. Revenue and the gross margin percentage are affected by the timing of new Intel product introductions and the demand for and market acceptance of Intel's products; actions taken by Intel's competitors, including product offerings, marketing programs and pricing pressures and Intel's response to such actions; Intel's ability to respond quickly to technological developments and to incorporate new features into its products; and the availability of sufficient inventory of Intel products and related components from other suppliers to meet demand. Factors that could cause demand to be different from Intel's expectations include customer acceptance of Intel and competitors' products; changes in customer order patterns, including order cancellations; changes in the level of inventory at customers; and changes in business and economic conditions. The gross margin percentage could vary significantly from expectations based on changes in revenue levels; product mix and pricing; variations in inventory valuation, including variations related to the timing of qualifying products for sale; excess or obsolete inventory; manufacturing yields; changes in unit costs; capacity utilization; impairments of long-lived assets, including manufacturing, assembly/test and intangible assets; and the timing and execution of the manufacturing ramp and associated costs, including start-up costs. Expenses, particularly certain marketing and compensation expenses, vary depending on the level of demand for Intel's products and the level of revenue and profits. Intel is in the midst of a structure and efficiency review which may result in several actions that could have an impact on expense levels. Intel's results could be affected by the amount, type, and valuation of share-based awards granted as well as the amount of awards cancelled due to employee turnover and the timing of award exercises by employees. Intel's results could be impacted by unexpected economic, social, political and physical/infrastructure conditions in the countries in which Intel, its customers or its suppliers operate, including military conflict and other security risks, natural disasters, infrastructure disruptions, health concerns and fluctuations in currency exchange rates. Intel's results could be affected by adverse effects associated with product defects and errata (deviations from published specifications), and by litigation or regulatory matters involving intellectual property, stockholder, consumer, antitrust and other issues, such as the litigation and regulatory matters described in Intel's SEC reports. Please refer to Intel's most recent Earnings Release and most recent Form 10-K or 10-Q filing for more information on the risk factors that could cause actual results to differ.