Intel IDF
9-18-07 Morning session

[Beginning of recorded material]

[Video playing.]

Female Voice: Ladies and gentlemen, please welcome Pat Gelsinger.

Pat Gelsinger: Good morning.

Voices: Good morning.

Pat Gelsinger: Okay, let's try that again. Good morning!

Voices: Good morning!

Pat Gelsinger: Much better. Thank you very much. Welcome to IDF. We're excited about today. In 1997, we launched the first IDF. Intel was at the forefront of silicon, the Pentium II, a whopping 350 nanometers, 150 megahertz. You remember when that was fast, right? Also, 1947, the first transistor was invented, so the 60-year anniversary of the transistor, the 10-year anniversary of IDF. Right? And we think today is going to be the best IDF ever as we kick it off this year. And, of course, what makes a great IDF, great content. We're excited about three days, right, full of content and the classes and the tracks and then the keynotes.

Today, in his Intel keynote, Paul Otellini will give an overview of our strategy. I'll be on later today describing our enterprise strategy. Tomorrow is everything mobile, right, we here we'll have Dadi and Anand presenting. And on day three, right, software and the future worlds that Justin will be describing for us. At lunchtime today, I'm very excited to have Gordon [Moore] presenting. Really, what figure could you look at in the semiconductor industry who has had a greater impact than Gordon Moore? And today you'll have a chance to hear from him. Training, gaming, IDF showcase, just a tremendous amount of wonderful activities we have in the three days of IDF.

And, of course, the ultimate geek fest, right? So here we are, some of the, you know, some of the greatest brains in the industry, in technology gathered together, and you need some of the greatest technologies. And we have some of our Intel fellows, some of the senior fellows who are here, some of the Intel fellows. And, you know, go ahead, stand up, guys. A bunch of them are here, right, our Intel fellows? Stand up in the front row here, right? You know, oh, here we go -- there we go -- sorry. There we go, right? And these are the certified smart guys of the industry. But it's also not just the Intel brains, right, but we also have industry leaders, right, from many of the companies who are listed here, who are participating, training, and presenting as part of the tracks training and showcase today. And over 30 percent of the sessions are led by industry leaders as well as the Intel leaders who we've pointed out.

We also have two new innovations in information sharing. We have Channel Intel, launched by myself and Justin this week, right? I get a little bit frightened when I see myself on YouTube, a little bit -- I'm just not quite sure about that media, yet, right? And a fun and informative way to keep up on what we're doing. Also, we have Open Port, our Web 2.0 community for IT professionals. And we have over 1,700 registered IT professionals as part of that community. And, finally, make sure to visit the Upload Lounge to share and learn as part of this new phenomenon.

And, of course, IDF -- the Intel Developer Forum? Nay. The Industry Developer Forum. And we see the role of our industry partners here as absolutely essential to the success of what we do. And I want to say a personal thank you to our gold as well as our silver sponsors who enable us to bring a tremendous IDF.

In closing, let's celebrate the last decade, the innovation, the things that we have done together to literally change the world. Today, let's kick off the next decade, a decade at least as exciting at the last. Thank you very much for your participation with us at IDF. And now for the rest of the show, Paul Otellini for our initial keynote of the day. Thank you very much.

[Applause]

[Music]

Female Voice: In a world moving at the speed of thought, today's extreme technology is tomorrow's mainstream. In the business world, relentless competition drives the quest for extreme performance at lower power. For consumers, the desire to interact with video, music, and games drives the demand for extreme capabilities. In communication, WiMAX brings extreme connectivity to new users in new places, while breakthrough in new form factors let users take mobility to a whole new level.

[Music]

Female Voice: Ladies and gentlemen, please welcome Paul Otellini.

[Applause]

Paul Otellini: Good morning. Good morning, and let me add my welcome to IDF. I hope you enjoyed the video this morning. What it reflects is the theme of my talk. My talk today is all about extremes, extremes in product, extremes in technology, and extremes in usage. And it's also about, though, how we as an industry have to collectively come together to drive new technology into widespread adoption. In short, the theme of today's talk is extreme to mainstream, how it's our job as an industry to drive technology from its inception into widespread adoption.

I'm going to explore three areas today. The first is a bit of a history, how 40 years of Intel technology have helped create the mainstream that we have today and what the impact is. The second is to explore some of the unique capabilities of our company in terms of allowing us to develop and deliver new technology at a rapid pace. And thirdly, most importantly for this forum, where are we going. What are the new technologies, what are the new markets that Intel is reaching out to explore for all of us to be able to participate in as we grow our business?

Let me start by taking a look at our innovations to support how we live today. And to do that, I'm just going to start with a picture. It's a picture that is very common to all of us. Probably sitting out in the Upload Lounge in front of this room many of you looked like this, or down at Starbucks, or on a campus, or at an airport. Essentially, it's a Wi-Fi-enabled Centrino notebook, USB, Bluetooth, MP3 players. It's a picture of the ordinary. But five years ago, this would have been a picture of the leading edge of the extreme. Ten years ago, this level of capability would have been simply impossible. How did it happen? How did we get here?

There are a number of technologies that were put in place to make this happen. First and foremost from Intel's perspective, the microprocessor. Invented almost 40 years ago as well, 4004 all the way up to the Core-based product, Core 2-based products that we have today, a relentless pursuit of Moore's law in bringing new technology to the market year after year. We realized early on that connectivity was a critical element in computing. We were one of the three companies that did the specification and development of Ethernet, and over the intervening multiple decades, we pushed communication technologies to the point where we now have hotspots for Wi-Fi all around the world.

Memory -- Intel was the leader in the early days of semiconductor memory, moving from the old core memory of mainframes to more reliable, lighter, higher-capacity solid-state memories, from SRAMs to DRAMs to the EPROMS to the non-volatile memories that you see in Flash today that are in the midst of all our handheld devices like phones, MP3 players, portable video cameras, and so forth.

And lastly, as things got smaller, as things got portable, power efficiency became more and more important. Beginning with the 386SL, we have a track record of power efficiency in our microprocessors in our platform products, and we've now extended that out to the latest innovations in SpeedStep technology.

Now, Intel didn't invent everything on this page and everything in that picture, but we've sure invented a lot of it. And what you're seeing today is that each of these pieces are needed, had to be in place in order for us to draw that picture. I think going forward that the message to all of us as an industry is very important and very clear -- that the innovations that we as an industry are making today are the basis of the future of the computing environment, and probably for the basis of the digital world. And our goal at Intel has always been to develop cutting-edge technologies and bringing to market very rapidly.

So let me shift and show you how we do that. What makes this all happen? From our perspective, there are really three fundamental capabilities at Intel that we're exceptional at. The first is silicon-process technology, the heart and soul of Intel Corporation, executing, delivering on Moore's Law year after year. The second is our Intel architecture and increasingly, the platforms based around our microprocessors that bring life to computing. And the third is our ability to create markets. That has to do with the scale of the company, the scale of our capacity to deliver and meet your needs as new products come out, and most importantly, a global reach and ability to diffuse new technology into the marketplace at a very, very rapid pace.

Let's start with silicon technology. And forgive the science lesson here, but I thought it was important to get down to the most essential element of Intel, which is our silicon processing technology and at the heart of that, the transistor. There is in every transistor, and there's a picture of them up behind me here, a very, very thin insulator layer. It's a thing we call the gate dielectric. It separates the gate from the rest of the transistor.

Since 1968, when Intel was founded, that insulating layer was made out of silicon dioxide. As late as a decade ago, or even more than that, 15 years ago, our senior technologists started realizing that this was problematic. Why was it problematic? As transistors got smaller and smaller, there was increasing propensity to have leakage, to have power dissipate through that gate, which is not good for transistor performance and certainly not good for the overall power characteristics of a system.

And as they tunnel through there, we had to find a new way. We had to find a new material, a new technology to deal with the gate dielectric. We started doing research in this in the 1990s. We looked at thousands of possible permutations. We cracked the code a few years ago in 2003. The breakthrough that solved this problem is a new materials-based insulator that's based upon a material called hafnium. Hafnium allows us to produce a very high-K gate dielectric, or high-resistance gate dielectric, along with metal gates, to be able to essentially address that leakage problem and deliver high performance.

We've replaced silicon dioxide with this new gate, and the results are pretty stunning. On one axis, we were able to deliver on what is traditional for Intel with silicon generations, which is another 20 percent faster performance at the transistor level. But the real breakthrough is that we took the power, in terms of gate leakage, down by a factor of 10. This is enabling us to provide a whole array of new products with high-performance, low-power characteristics and ultimately address new markets.

I call this the magic of 45 nanometers, which is the latest silicon technology that Intel is going into production on as we speak. The beauty of this technology is not in just what it could do that's predictable, and that's the area in the middle of this curve of this chart. We've always addressed a given level of markets, in terms of notebooks, desktops, and servers, with our standard silicon technology. With 45 nanometers, though, we get two things that are new to us. We get a threshold level of density in terms of the amount of transistors we can cost effectively put onto a die for other features, but also the power characteristics of that.

So, for example, you'll see Intel do integration of things like graphics onto the microprocessor in this generation of technology, something we've not had the die budget to do before, with prior generations. At the low end of the market, where power and form factor matter, you'll see us build very, very tiny die based upon a [Silverthorne] core that are system-on-chip implementations to address the needs of ultra-low-power portable computing. At the high end of the market, you'll see us use that transistor budget in a very aggressive fashion to build an array of X86 devices or Intel architecture cores in a many-core environment that I'll talk about a little bit later for a product called [Larrabee]. The overall capability here I think is pretty unique. Other companies have announced 45 nanometers. No one has announced the inventions in terms of the hafnium-based high-k dielectrics that we've talked about today. We will be, I think, unique in terms of the power, the performance, the transistor density, and the yields in this generation of technology, and we're well ahead of the industry.

You've come to expect us bringing out a new silicon technology generation every two years, and in fact, we've done so for 14 years, from 1999 through today, 180 nanometers down to 45 nanometers today, going into production. We're making great headway on the next generation, though. We have to be ready for what's next. The next generation is 32 nanometers. It uses a second generation of high-k/metal gate transistors, and I'm happy to say that we have our first -- the world's first product fully functional on 32 nanometers. This happens to be an SRAM. There's a 291-megabit SRAM array of die on here. Each die is 1.9 billion -- that's with a B -- billion transistors. Pretty amazing capability. What this does is starts to give us the know-how, the confidence to build mainstream microprocessors on this technology a short two years from today.

Let me shift to platforms. Platforms have been the major thrust from Intel's development perspective the last few years. And in our definition, our parlance, a platform really is, very simplistically, an Intel architecture microprocessor tailored for the task at hand, a collection of chipsets, logic chips, graphic chips that sits around it, again, tailored for the task at hand. We look at market segments and identify user value features for those markets, like WiFi was for mobile or security manageability are for the enterprise, add those to the platform, and wrap them up into one architecture. The benefits, I think, for all of you as developers are pretty widespread. It's a quick time to market. You can ramp into new technologies, new market segments very rapidly by taking advantage of this platform approach. There's a very, very low cost associated with this because much of the work is done upfront and there's very high performance.

I think, though, increasingly, the long pole in the tent, in terms of getting new products out, is software. And what this platform approach does for software a developers, particularly in new and adjacent markets, is it makes your job a lot easier. We are providing a wide availability of tools for all these products, and the software reuse -- so the nature of the compatibility of the architecture and the nature of the compatibility of the Internet being written around Intel architecture is really available to you as you design new devices that have Internet capabilities built in. And, of course, we have a very consistent roadmap nowadays focusing on where we're going, so you can have high predictability on your development around Intel's development.

At the heart of the platform is Penryn, it's the microprocessor, and our next microprocessor is Penryn. This is taking the same Core 2 microarchitecture, which is shipping today on 65 nanometers and moving into the next generation of silicon technology. Penryn, in the large Quad-Core configuration is a 410-million-transistor die. We will launch these products on November 12th, about two months from now. And it is the first of 20 unique microprocessors that are being developed on 45 nanometers -- that scale of products I showed a few minutes ago -- that's coming out. There'll be a large number of SKUs that are launched on November 12th for servers and for high-end desktops. And then, in the first quarter of 2008, you'll see us bring out another 15 or 20 SKUs or so to address not just those markets, but also the mobile markets as we move this technology into the mainstream across all of our product lines.

We're very, very excited about our customers' adoption of this technology. There are over 750 design wins around the Penryn family as we approach a launch. And as I said a few minutes ago, Intel is in production on 45 nanometers. And Penryn today -- and in fact, as we are speaking, wafers are moving through the FAB in anticipation of shipping to customers for that November 12th launch.

We feel, though, that it's not just packaging -- not just silicon technology that has to be pushed into the envelope here. It has to be packaging technology, as well. We are designing a new generation of packages for our microprocessors that is 60 percent smaller than the prior configurations, allowing you not just better flexibility in your small form factor designs, but also lower costs for those that are focusing on low MPCs for emerging markets and other needs. Smaller is better, smaller is cheaper.

We're also focusing on building cleaner products. We announced last April that we were 100 percent lead-free in terms of our products shipping out of Intel in our microprocessor product line. Today, I'm happy to announce in addition to lead-free, as we move to 45-nanometer generation of microprocessors, and the 65-nanometer generation of chipsets that surround those microprocessors, Intel will also be halogen-free by the end of 2008. Lead-free, halogen-free, smaller, cheaper, cleaner for the environment.

So, what's next? What's next is Nehalem. The tick-tock model that Pat will talk about later on today has implicit in it a number of things. We have to move to generation-to-generation of silicon technology, and in alternate years we'll move microarchitecture-to-microarchitecture generations to keep the pace of innovation going.

So, this year, we have 45-nanometer Core 2 with Penryn. Next year, we'll have 45 nanometers with a new microarchitecture called Nehalem. Nehalem is a very dynamic design from a number of perspectives. From an Intel perspective, it's a very modular design. We have the ability to change the configuration of cores, to change the configuration of cache size, to change the configuration of I/O, power envelopes and so forth to be able to meet the needs of an increasingly diverse segment of our product needs.

From a developer's standpoint, though, of systems, it's also very dynamic. You'll be able to configure the real-time needs of the system in disenabling power states and so forth, to be able to optimize the performance for the given task at hand. At the largest configuration that we'll ship in 2008, they'll be an eight-core product. Eight core on one die, and each core will have two threads. So, each eight-core die will be supporting 16 threads. Think about what the performance could be in a dual-core or Quad-Core and beyond configuration for simultaneous multitasking.

Now, I said that we are on track to second half '08. What I wanted to say is that the design is complete. Nehalem was finished about -- I guess about a month ago. We have wafers running in FAB. This is one of the first Nehalem wafers that's come out of FAB. Each die on here has 731 million transistors on it, and it represents the next generation of Intel microarchitecture.

To tell you a little bit more about this stunning product, I'd like to bring out Glen [Hilton]. Glen is one of the Nehalem chief architects and an Intel Fellow. Glenn?

Glenn Hilton: Hi, Paul.

Paul Otellini: Hi, how are you?

Glenn Hilton: Great.

[Applause]

Paul Otellini: Glenn, I'm so proud of the accomplishment that you and the team have done here. Can you maybe tell the audience some of the neat things you built into it?

Glenn Hilton: Sure. I'm excited to be here to talk a little bit about Nehalem today and to represent our outstanding Nehalem team. Our team in Oregon's been working on Nehalem for a long time as you might expect. And it's great to see it finally here. When we started Nehalem, our goal was to come out with a new design that was higher performance than the existing industry-leading Core 2 microarchitecture, and that was already very fast. Our goal, we wanted to build the highest performance core that could be used in a variety of systems from the best laptops all the way to high-performance servers.

Paul Otellini: Well, you know, maybe you could show us what else this product could do or what some of the other features were?

Glenn Hilton: Okay. Nehalem was designed to be, as Paul mentioned a little bit earlier, to be a scalable multicore system. But we also spent a lot of effort improving the single threaded performance of Nehalem since many applications don't have a lot of threads. So Nehalem is supposed to cover a wide range of systems and configurations. And to enable this, we had a very strong focus on power efficiency. For example, every feature we added to Nehalem had to pass very stringent power efficiency guidelines, even more than we have done in the past.

Paul Otellini: What about things like interconnect and memory bandwidth and so forth?

Glenn Hilton: Right. Nehalem implements several new technologies, several new important capabilities. One of them is a new high-performance integrated memory controller, DRAM controller. This new DRAM controller provides very low latent -- called Quick Path memory controller. This DRAM controller provides very low latency to DRAM and provides very scalable bandwidth for memory-intensive applications.

Nehalem also implements a simultaneous multithreading that Paul alluded to a little bit earlier. That's the capability of each core to have two logical processors. And each of these logical processors can be working on a different thread or application at the same time. And this SMP is very power efficient, and it really improves the performance of threaded applications.

One other thing I'll talk about is Nehalem implements a new system interconnect called Quick Path Interconnect. And it's the new link that enables multiple-socket systems to hook together. So this Quick Path Interconnect allows us to make very scalable high-performance servers.

Now, Paul, we have a very special guest we'd like to show you today.

Paul Otellini: Okay.

Glenn Hilton: And here it is.

Paul Otellini: What is it?

Male Voice: Hello, Paul. I am Nehalem.

Paul Otellini: Wow, you've got this thing booting already.

Male Voice: I am only three weeks old, and I am already talking.

[laughter]

[Applause]

Glenn Hilton: Now, Nehalem's looking great for having silicon for only about three weeks. And it's great to see it running major operating systems and applications. And we're looking forward to turning it into products in 2008.

Paul Otellini: Oh, so am I. And before I add my thanks to the team, I just want to see what else this can do. You know, this is running Windows today. It looks like XP. I got an email this morning that said you guys had now booted OS X. So I understand that Pat's coming on in about four hours, and that you guys are continuing to work on the product. So could you tell Pat that I would expect the product to do a little bit more than talk to us four hours from now?

Glenn Hilton: Okay, I'll tell him that.

Paul Otellini: Okay, and thanks to you and the team, Glenn.

Glenn Hilton: Okay.

[Applause]

Paul Otellini: So that was our processors -- Penryn, the next generation; Nehalem, the next generation microarchitecture coming on about a year after that one.

Let me shift and talk a little bit about platforms. We've talked about platforms here in the past, Centrino being the first one in terms of integrating wireless, extreme battery efficiency and high performance. The second platform was Pro. We introduced this platform focusing on the IT market, on the need to reduce operating expenses in the enterprise and in small business.

We launched vPro just about a year ago, in September of last year, focusing on hardware manageability and security for desktops. We launched [Intel®] Centrino® Pro [processor technology] in May of '07, combining the goodness of Centrino [processor technology] and the capabilities and features of [Intel] vPro™ [processor technology] into a mobile environment. So far, we've shipped five million units, in the year that we've been in the marketplace. There is very broad OEM support, essentially all the OEMs who produce business SKUs are now shipping pro versions of those SKUs. Very large ISV support in terms of applications, manageability, dashboard, security applications, and so forth. And a very large support from the service providers who work in IT shops, to help lower their costs. A few weeks ago we launched our second generation of this product. Pat will talk more about the future of [Intel] vPro [processor technology] in his keynote later on this afternoon.

As we look at platforms, we look really in two different vectors. One is focused on performance, energy efficiency, and user-valued features that you've now seen in notebooks, desktops, and increasingly in servers. The other vector, though, taking advantage of the capability of the silicon technology we have, is to focus on new markets for ultra low power, where performance is defined as being sufficient for the task at hand, and extreme integration, making them ultra portable and ultra cost effective over time. This brings us into new markets for low-cost PCs, for consumer electronics equipment, and ultimately for mobile Internet devices. I think if you look at the spectrum, this represents, from our view, a growth opportunity for the industry, not just in the markets that we serve today but in markets that are very near to us that have the same characteristics and are waiting for this industry to be able to serve new and exciting products.

One of the things I said we do very well is we help create new markets. To do this, we focus three things on these markets. First of all is scale. Second is capacity. And the third is this diffusion technology, the ability to move new technology rapidly into the marketplace. Scale is important because it allows us to ramp very rapidly. What I've plotted here are the defect density curves for four generations of successive silicon technology at Intel. The curve on the right is our latest technology, 45 nanometers, which is moving into production now, as I said earlier. On this curve, lower is better and steeper is better, so the faster you are down to the bottom of the curve, in terms of minimizing the defect densities, the more predictable your products are.

This allows us to be able to have the confidence to ramp new products very quickly knowing that we'll have high yields and full functionality across the entire array of the die. Without this confidence, we couldn’t possibly afford to invest in technology the way we do today. Our current capacity picture for 45 nanometers is on the slide behind me. We have two factories up and running on 45 nanometers today, one in Oregon -- D1D -- and one in Arizona -- FAB32. We have two more factories nearing completion that will come online in Israel and in Albuquerque in '08 to fulfill the demand that we see for 45 nanometers over the lifetime of this technology. Each of these factories is in excess of $4 billion. We couldn't possibly commit this kind of capital if we didn't have confidence in the silicon technology underlying it, otherwise we'd have a high risk in bringing it up, and in the microarchitecture in the products that are going to ramp very quickly on this technology. You can count on us to deliver the products that you need for this generation of technology.

And the third one I mentioned was diffusion. And I've shown diffusion curves at prior IDFs to show how fast we get to ubiquity in various new technologies. I thought we'd take a different slice today and look back at four areas that I think were incredibly important to the industry in the last five, six years. The first is something I first talked about in IDF 2001. And it was really about moving away from gigahertz and moving towards energy efficient products. The first one of those was Banias, which was at the heart of the first Centrino [platform]. The metric on this was that Banias crossed over its predecessor in seven months. We were able to take new thresholds of power and performance and bring it to the market very rapidly.

In '02, we talked about hyper-threading, and the purpose of hyper-threading, where we put multiple threads in a given single-core microprocessor, was to enable the world of software to get ready for the world of multicore and many-core. And it worked. We had 100 percent of servers enabled with hyper-threading within two years of the launch of this technology [in] over half the clients.

In 2003, we first talked about integrating wireless, Wi-Fi, into the Centrino[-based] notebooks. Today, or actually two years later, 90 percent of all notebooks shipping with Wi-Fi enabled.

And 2004, three years ago, we talked about the move to parallelism, the shift to multicore in all of our products, and our answer to the power wall that we had faced. Essentially today all dual-core processors -- we're shipping mostly dual-core processors. They've crossed over in 2006. And we've shipped over a million in the first four months of this product line.

So when Intel talks about a new technology in this room, I think you can count on us to deliver it to you in the marketplace a short time thereafter and make it a worldwide phenomenon.

So we've talked about the capabilities, we've talked about the history. I think now I'd like to shift and talk a little bit about where we're going next. What's going to be the next mainstream? And today I'll talk about four areas. In the extreme -- extreme mobility, extreme entertainment, extreme problem solving, and extreme inclusion.

Let's start with mobility. As you know, this is where the market is hottest today. It's focused on all of us, people that are on the go, all of our kids, all of our family. All of us want to stay connected, as much as possible, real-time, wherever we're at. You've seen this demonstrated in notebook growth very rapidly over the last few years. This chart plots the percent of our microprocessor PC chips that are shipped every year as notebooks as a percent of desktop. And you can see prior to Centrino [processor technology] that notebooks represented about 20 percent of the market. Post-Centrino, they've grown very rapidly to almost 40 percent of the market, 120 million Centrino [processor technology] have shipped to date, and we are now projecting that the notebooks will cross over the desktop in unit volume two years from now, 2009. Think about this from where we were just a few short years ago.

We, as an industry, identified a need, we marshaled our resources, we marshaled our marketing, we marshaled our technology, and we've brought the price points down to where now you can see a point in time where the notebook is the dominant PC built by the industry year after year after year. Mobility's become mainstream. No one would argue that. But I think we all have problems. We all expect to be connected everywhere. In fact, it's a right, from many perspectives. When my Wi-Fi network is down or when DSL is down, I have a problem at home. I yell at people. So do all of you. When you can't get connected in the airport, you probably get very frustrated. Some things are just still too hard to do or darn right impossible. Mobility has a lot of room to improve, despite all the progress we've made the last few years. We need new devices, we need new services, we need new products.

Interestingly, though, I think that the future of ultra mobility is a lot closer than we think. And it's really going to be enabled by a combination of two things. These new ultra small, ultra low-power microprocessors and chipsets that we'll build, that will enable you to build very, very small form factor devices that have the ability to talk to the Internet anywhere at any time. But that needs a network. WiMAX is, from our perspective, the network that most of these things will be connected on because of its ubiquity, its' potential for ultra low cost, and its capacity for number of clients and the overall broadband capability.

So let's talk about WiMAX. A year ago, we talked about the trials in WiMAX moving from fixed WiMAX to mobile. And, at the time, there were 10 mobile WiMAX trials going on in the world at IDF '06. Today, there's 120. A year ago, there were two large vendors that had committed to WiMAX for North America. Clearwire and Sprint. They're still committed. Today, we're happy to announce that Intel is investing, along with KDDI in Japan in a bid to enable WiMAX in Japan on a widespread basis.

Now, to make WiMAX ubiquitous, we don't just need the networks, we need all the PCs and ultramobile devices to be able to talk to them. And to enable this we are developing an integrated Wi-Fi/WiMAX module, single module that has both capabilities, that will be available for notebooks and ultramobile devices next year. This product is called Echo Peak in its code name. And today I'm very happy to announce that Lenovo, Acer, ASUS, Panasonic and Toshiba are all committing to integrate WiMAX into their notebooks with the Montevina platform launch next year when Centrino [Duo and Centrino Pro processor technology] refreshes the middle of next year.

WiMAX, from our perspective, is moving to the mainstream. If you look at the numbers, they're starting to get to be very large. The contractual commitments today suggest that 150 million people will be covered with WiMAX in 2008, principally in North America; 750 million two years later in 2010; and 1.3 billion two years later than that in 2012. It's just the stuff that we see happening in the marketplace today. So we are on the cusp of a new global network, seamlessly integrated, hopefully the roaming is seamlessly integrated, the same technology around the globe to be able to go into these ultramobile devices and notebooks as they roll out.

On the microprocessor side, the thing that we need to enable this is low power. This is a clip from a slide I showed at IDF 2005. And at the time, I said we were pushing Intel® architecture in two directions. We were going to increase by a factor of 10x in performance by 2010, five years from when I talked about this, on the right side of that scale. And we were going to decrease in terms of lower power on the left-hand side of that scale for a class of new devices. We are tracking to the right side. We are on track to deliver a 10x increase in performance. I'm afraid to say we are not on track on the left side, on low power. We were wrong. We're two years early. We'll deliver [up to] a 10x reduction in power next year with a product called Menlo. Menlo is a platform comprised of our Silverthorne microprocessor and the Poulsbo low-power chip set that goes with it.

As we worked on these products, though, we found we needed something else. We needed to focus on not just the low power, the power state when the machine was working, but also the power state when the machine wasn't, the idle state. What matters in a battery device is that listen time, standby time before you're talking or before you're surfing. And we found that this was something that needed some work. So we are committing today another factor of 10 change. Here we are going to reduce the idle power by a factor of 10 by 2010 with a product called Moorestown that Anand will talk more about in his speech tomorrow. I think when you look at the continuity of these devices in terms of the roadmap over time, you get a very exciting glimpse of the future of what's possible in ultramobile machines.

Let me shift to notebooks, markets that are large and that are here today. [Intel Centrino Duo and Centrino Pro processor technology] is the current generation of Centrino [product] that's shipping. We've been shipping since May. We've now shipped 10 million [Intel Centrino Duo processor technology]-based platforms into the marketplace. That platform will have a refresh in early '08 with a version of Penryn, the 45-nanometer microprocessor that is tailored for that platform. It also has better integrated graphics at the time.

In mid-'08, we're going to launch a new platform called Montevina. Montevina has a number of new features. It has an optimized Penryn microprocessor for that platform. It has a new chip set called Cantiga, which has better integrated graphics and other features. It has integrated Wi-Fi and WiMAX that I talked about a little bit earlier. It is a chipset that will support in native mode both HD-DVD and Blu-ray, so both of the high definition formats. And we're going to take the form factor, the footprint of this down by 50 percent. And for those of you working on very, very small form factors, I'm also pleased to announce today that we'll have a 25-watt version of this available when we start launching the product.

Montevina is out of the lab. As I said, we're getting ready for shipments next year, next May. And there's a demo of it today in perhaps the ugliest notebook ever built. The guys were clever that it's in Plexiglas. It's just a motherboard, and it's in debug mode, but we're out getting ready to start sample the products to our customers and you can see it's working. This is a bit ugly, and I've challenged Dadi, who's talking to you tomorrow, to see if he can get this brand new silicon into a real notebook form factor for a real demo tomorrow. But stay tuned and watch that space.

Okay. Let me shift to extreme mobility. There's a new class of devices that's starting to emerge that we call mobile Internet devices, or ultra-mobile devices. They're growing in popularity, but there are still a number of limits to these devices. You can't get connected, in many cases. It's great on Wi-Fi. You leave the Wi-Fi network and the service is intermittent and/or slow. And the Internet compatibility, while it can be made perfect, isn't perfect today. Any part of the Internet can be transcoded to run on any device.

But that's a big job. The Internet's large. From our perspective, having Internet compatibility from the start, across the breadth of the Internet, is a really exciting value proposition. The solution to this that we see is integrating Intel architecture along the lines I talked about, with small ultra-mobile devices and WiMAX.

Now, in the video, when we kicked off the speech, you saw a guy doing some jumping off a cliff and then an ultramobile device and ultimately landing hopefully safely. I thought I would show you where that was actually taken. I take you from Moscone Center here over to Utah, to Zion National Park out in the middle of the desert there. And this is where that jump was done. Now it turns out the jump team is still there. And I'd like to take you live, via satellite, to this site. Hello, Steve.

Steve: Hey, Paul, how you doing today? Good to see you!

Otellini: Good to see you. I see you and Jeff are there, huh?

Steve: Yes, we are. Actually, I'm way up and Jeff's way down there. We're actually coming to you live via satellite, as you mentioned, right outside of Zion National Park in Southern Utah, not far, actually, from where we did that opening video. And if you're wondering why I'm dressed like this, well, I'm going to do my own base jump as part of this demonstration.

Otellini: Uh, that sounds awfully risky, Steve.

Steve: Oh, thanks for your concern, but I'll be fine.

Otellini: Actually, I wasn't worried about you. I was worried about the computer.

Steve: Oh. Oh. Well, no problem. I'm sure that my body will cushion the blow.

Otellini: Oh, good.

Steve: Besides, Paul, with a WiMAX enabled MID like this brand-new device I have from Compal right here, I've got the world's knowledge base right in my fingertips. Which, as you can imagine, is pretty useful when you're out in a remote location like the one we're at here. In fact, Jeff is down at the base camp right now. He's been surfing the Internet for some last-minute tips for my jump, so Jeff, what kind of information do you have?

Jeff: Hi, Paul. I'm here at Zion base camp, and I'm checking the latest wind and weather conditions on my MID from BenQ. I'm accessing a Java-enabled webpage which is something you can't do on most smart phones.

Steve: Paul, something else that's cool. Check this out. I've got my USB webcam right here attached to my helmet, and I'm going to connect that to my mobile Internet device, when I get Internet, as soon as I'm done, for the whole world to see. Maybe the people in the audience would like to see that. In fact, to give you an example, before I go, Jeff has one of the last ones that I sent up. Jeff, why don't you show them that one?

Jeff: Sure. WiMAX isn't just about MID, so why don't I show you this next example on my WiMAX-enabled notebook. Here, we've upload the video to YouTube and I could've easily done the same on my MySpace page or Steve's video blog. Now, another cool thing we could do is we can watch our home television programming using Slingbox*. Here, we're watching a 900-kilobit stream over the WiMAX link out in the middle of the desert. Now, I think that's pretty cool.

Steve: Jeff, I got to tell you, that is pretty cool. That's a lot of bandwidth there. I can't wait to get to the bottom and watch some TV with you. That looks cool. Do you have anything else for us?

Jeff: Yeah, one more quick example to note. Internet banking is something you want to be able to do anytime, anywhere, and because of the encryption that's required, it requires a little more performance than your standard hand-held device. So what I'm going to do here is I'm going to log into my home banking site and I'm going to make a payment to make sure that Steve's medical insurance is up to date.

Steve: Oh. Thanks a lot for the vote of confidence there, my friend. While you were talking, I just ordered a pizza to be delivered to my GPS coordinates at the landing zone to kind of celebrate, you know what I mean? Oh, hey, you know what, Paul? The wind conditions are absolutely perfect right now for my jump. I've got to go. I wish I had a lot more time to tell you about all the cool things you can do with a device like this, but basically when you have an [Intel architecture]-based platform and you connect that with and combine that with WiMAX, well, you've got the real Internet right in your pocket, and that is pretty cool.

Paul Otellini: That is cool, Steve. Thanks, and good luck with your jump.

Steve: Oh, hey. Thanks a lot, Paul. I appreciate it. Okay. Here we go. One, two, three, Geronimo!

Paul Otellini: Wow.

Steve: Are we clear? That was great, guys. Awesome. You smell that? You smell that? That's major promotion. That's the smell of --

Paul Otellini: Steve! Steve! You're still on!

Steve: [I thought we were off!]

Paul Otellini: Man. I've got to get some new demo guys. [Applause] In all seriousness -- we had some fun with that, but we did that to show you -- I mean, that was an extreme idea, I mean, WiMAX in the middle of the desert, ultra mobile devices, maybe or maybe not a real jumper doing a real jump. But it shows you what's possible, and that's really what this forum is all about -- what's possible, and what our job is, is to make that from possible to probable over the next few years and make those devices real.

Let me shift and talk about entertainment. We've seen a fundamental shift in the entertainment industry in the last two decades as it's shifted from analog to digital. That's changed the way movies and records are made, it's changed the way these kinds of entertainment vehicles are distributed and consumed by us as individuals. That revolution is now starting to happen in our homes. It's happening with the evolution of consumer Internet, which is becoming increasingly focused on highly visual applications and highly-interactive applications. This moves beyond movies, moves beyond videos. It's focused on social networks, it's focused on user-generated content, it's focused on 3D graphics, on virtual worlds, on virtual economies, and on games. And I thought we'd explore a little bit of this in the area of games this morning. Games used to be a rather niche market. They're becoming mainstream.

Internet gaming is driving much of the gaming growth today. If you look at the United States today, over 80 million online gamers participate in that entertainment on the PC, versus 7 million who play online in consoles. In Asia, it's even more extreme -- 114 million gamers online, on the PC, versus 2.7 million online on consoles. PC and Internet gaming is where gaming is going. Now, when you look at the gaming environment, rich gaming experiences require great processors and great graphics. And if you look deeper at the extreme gamers in this area, they've long had a demand for the fastest systems. They want insatiable performance. These are the guys who over-clock the processor, who over-clock the memory systems to get that competitive edge in the games with their friends.

So I thought it would be fun to have a view of extreme processor performance by bringing out Charles Worth. Charles is the founder of a company called Extreme Systems. Good morning, Charles.

Charles Worth: Good morning.

Paul Otellini: How are you? [Applause]

Charles Worth: Good, good.

Paul Otellini: So, Charles, I understand you're an expert in pushing our processors to the limits?

Charles Worth: Correct. I'm very well known for pushing Intel processors.

Paul Otellini: To the limits.

Charles Worth: To the limits, very much so. Over-clocking to the extreme.

Paul Otellini: Good. Maybe you can tell us what you've been doing this morning.

Charles Worth: Oh, let me just turn on my system real quick for you. I have an Intel desktop chipset with DDR3 and an ultra video card. My video card is [unmodded] and running at stock speeds. I do have a phase change cooler on here, and if you give me a few minutes, I'd like to get started. This morning I broke a few world records for you, four, and if you give me two minutes, I'd like to break three more [live] for you.

Paul Otellini: You can break three records in two minutes?

Charles Worth: Easily.

Paul Otellini: You got it. I just happen to have a handy dandy basketball timer.

Charles Worth: Okay. Tell me when to go.

Paul Otellini: On your mark, get set, go.

Charles Worth: Did you start already? Okay. As you can see, this is a SuperPI*, it's one-meg calculation, and currently Team Japan owns the world record at 5.9 gigahertz. And as you can see I'm an eight-second Quad-Core processor, a new world record. Moving on.
[Applause] This is AquaMark 3*, it measures both CPU and GPU performance. It's quite a popular benchmark for us. It's easy to push, and let me get started here. All right. The current world record's at 267,000. And like I said, a lot of things you can do on this benchmark is you can reduce image quality, you can crank up the video card. The video card's running at stock speed, I'm running at high image quality, as you can see on the screen. In just a few seconds later, this will be finished up and we should have a new world record. 3-2-1. Okay. So, we have 273,000. New world record for you. All right.

[Applause]

Pat Gelsinger: You've got a minute to go.

Charles Worth: Up next is Cinibench 10*. It's 32-bit, and I'm running well over 5 gigahertz for this benchmark. It's quite [unintelligible]. I'm also running a 32-bit. What makes 32-bit a little bit unique is it's a little bit harder; 64-bit does quite a bit better, and here we go.

All right. Extreme Systems -- we push systems to the absolute maximum performance. And with the previous generation of processors, this is actually a leap ahead. When we saw the [kimer] processor was huge compared to the previous generation. Now, this is the fastest processor we've ever seen.

Like Paul said earlier today with a gate leakage -- we can directly see that by pushing voltage into the core. The more voltage we push into this core, the faster we can push this processor. I have not reached the top speed of this processor yet, and I did reach 5.56 gigahertz already on an unmodded board. This board has no voltage modifications of any kind.

All right. So, we're finishing up here. 2-1. Oh! [I shouldn't have] talked so much here. That's okay.

Pat Gelsinger: You're in overtime. Wrap it up.

Charles Worth: All right. I'm just shy of 20,000, a new world record for Quad-Core processor.

[Applause]

Pat Gelsinger: That's fantastic. Congratulations. Charles, you mentioned you had a rather unique cooling solution that you're using for this.

Charles Worth: Correct. This is a refrigeration system. I'm running at 160 below zero right now. It's a three-stage cascade cooler. This is pretty typical of what you find in extreme cooling, or what we compete with in competition. What me to pop the case here real quick for you and show you what it looks like?

Pat Gelsinger: Sure.

Charles Worth: This is made for long duration runs. It's running on refrigeration technology, and it's basically three refrigeration compressors chained together to reach 150 below zero. Actually, I'm reaching 160 for the demo here.

Pat Gelsinger: That's amazing. Don't try this at home, right?

Charles Worth: Thank you, Paul.

Pat Gelsinger: Thanks, Charles.

Charles Worth: All right.

Pat Gelsinger: Call me up when you hit six gigahertz.

[Applause]

Pat Gelsinger: They'll make that quiet in a second. Okay. To address this market -- and we see this as a growing part of the market -- Intel has developed a new chipset of platform technology called "X38." X38 is a combination of the highest end of our extreme edition Quad-Core microprocessors, which today is 65-nanometer Clovertown product. But in November, there'll be a 45-nanometer version of this called Yorkfield, which 12 megachache onboard that is able to be overclocked into systems like this. And in fact, that was a Yorkfield that Charles was running here.

The chipset is the X38. It will launch October 10th. We are shipping production parts to our customers today in anticipation of that launch. It has unlocked bus ratios, so they can crank it up. It has easy memory overclocking capability. We have doubled the bandwidth of the PCI Express Bus to graphics on the chipset. And we're incorporating into the chipset a new software tuning utility that allows you to tune directly to the BIOS to get the optimal system performance, so that you can win game benchmarks like this.

Beyond processors, we think graphics is increasingly important in the visual computing environment. And today, Intel is the world's number one supplier of PC graphics through our integrated graphics products. Integrative graphics from Intel today can play all the popular games and deliver Hollywood-quality video into the marketplace. But we've been delivering our integrative graphics on a silicon technology that is one to two generations behind the leading edge at Intel. And you can see this here. I plotted the video performance -- relative video performance of our products in 2006 on 130 nanometers, and 2007 on 90, while Intel is well into the 65-nanometer era with our microprocessors.

Well, we're changing that pace. We're changing that pace starting next year. We'll bring 65 nanometers out early next year into the graphics, and take the graphics up by a factor of a little more than two here. In the second part of that generation, in 2009, you'll see us move integrated graphics to 45 nanometers. And at this point in time it becomes incorporated into the microprocessor, so it becomes part of the CPU, matching the technologies for graphics and microprocessors in the same silicon generation. That gets us more than a 6x boost from where we were last year.

In 2010 is where we actually hit the design point from scratch, where the microprocessor and the graphics are designed from scratch to hit that same early launch point on a new technology, in this case 32 nanometers. At that point in time, our integrated graphics performance will be up by a factor of 10 from where we are now. This is good, but it's not good enough.

What we are focusing on is another new product called Larrabee. In the last developer's forum, Pat told you about Larrabee. It's a highly parallel, many-core product comprised of an array of Intel architecture cores. We are well into development of this product today, and it's our intention to be able to demonstrate that product in 2008.

What are we going to do with it? We think this brings the benefits of Intel architecture in a many-core array to the high-performance, visual-computing segments of the marketplace. We'll deliver teraflops of performance with this chip. And one of the things that we think is a unique advantage is that it will scale easily for software developers. It not only has the code compatibility of everything you're familiar with in iA space, but will have a shared cache to be able to make that an easy programming model for you in the developing community. We think it has applications in supercomputing, in financial services, and physics and health applications. But it's also got one more thing that it's going to be very good at, and that's graphics. And one of the beauties here is that it is not dependent on a new software paradigm. Again, the same existing programming models that you all know today will be applicable to this device as it moves us into discrete graphics.

And to tell you a little bit more about this, I'd like to welcome a brand-new member of the Intel family. His name is Jeff Yates. He's VP of product management at a company called Havok, which has joined the Intel family last Friday.

[Applause]

Paul Otellini: Good morning, Jeff.

Jeff Yates: Good morning, thank you.

Paul Otellini: How are you?

Jeff Yates: Great, thanks.

Paul Otellini: So I thought you would tell us a little bit about Havok and a little bit about the relationship with Intel.

Jeff Yates: Sure, absolutely. Well first off, it's great to be here. Thanks very much. Havok is a seven-year-old company, and our passion has been serving game developers and film effects producers over that period of time creating software technology that puts physics into games and into the effects in film. So we've been fortunate enough to be a part of productions like "Poseidon"*, "Troy"*, and the "X-Men"* on the film side, and on the game side, from titles, blockbuster titles like "Half-Life 2"* to "Halo"* all the way up to the current big release that's coming out like "Stranglehold"*, and, as we'll see from one of our partners, Pandemic Studios, "Mercenaries 2"*.

Paul Otellini: And how do you benefit from multicore in your product?

Jeff Yates: It's very critical, actually, because most of these creative types are going after reality and immersive interactive experiences. And the kind of compute power needed for that is almost insatiable. Quite often, if you want to have thousands of objects interacting, debris, dust, and sometimes bodies flying around, you need as much compute power as you can get.

Paul Otellini: Yeah. Well, you know, I understand you've got a lot of customers that are already anxious to use this technology?

Jeff Yates: Yeah, definitely. We have over 100 titles released today. We have many more in development right now. And we've actually invested quite a bit in making our physics scalable across multiple cores. So we've done a lot of that work. We've put it out there for the developers in our technology, and we're starting to see the fruits of that coming out today. And with something like Larrabee, we look forward to a great potential for adding the gritty details, dust and debris, that can be added on top of the interactive game playing side of the computer as well.

Paul Otellini: Maybe you could help the audience take a look at this?

Jeff Yates: Sure. I'd like to introduce one of our top developers, who's one of our customers, Pandemic Studios, Josh Resnick.

[Applause]

Paul Otellini: How are you, Josh?

Jeff Yates: Hi, Josh.

Josh Resnick: Hi, Paul. Good to see you. Hi, Jeff. Good morning, how are you?

Jeff Yates: Good, how are you?

Josh Resnick: Good.

Jeff Yates: Welcome.

Josh Resnick: Thank you. Actually, first off, I just want to say that Intel's investment in Havok is truly an exciting announcement for game developers like Pandemic. You know, during the last nine years as we've been creating our kick-ass games for the studios, we've relied heavily on Havok for bringing that level of realism and excitement to our products. And, you know, frankly, nothing else comes close to allowing us to do that. And as many of you might know, Pandemic Studios makes the kind of games that go boom. And Havok has been great in enabling that, and some of our hit titles like "Full Spectrum Warrior"* and "Destroy All Humans! 2"*. And that partnership is continuing on Pandemic's upcoming blockbuster hit, "Mercenaries 2"* on the PC.

And you know, when Intel first approached us about bringing their plans for Quad-Core technology to PC games, it instantly made sense to us. You know, frankly, we just could not perfect our vast open world games like "Mercenaries"* and "Saboteur"* without some of that Quad-Core firepower on the PC. So by enabling that technology, Intel was really ushering in a new standard for PC games development.

But enough about that. We actually have a demo that I want to show you. It's two minutes long and it's very early in development, so please bear with me. It's being developed by one of our partners, LTI Gray Matter. It's not coming out until early next year, but I hope that it'll give you a taste of what Quad-Core and Havok will be able to bring to a game like "Mercenaries 2" on the PC. So we're going to get started now and show you something very beautiful.

So here, the Pandemic creators have spent some time creating a very detailed dynamic environment, and what's the fun of creating an environment like that unless you start blowing it up, though? So in trademark "Mercenaries" style, we're going to get a black screen. Look at that! I didn't prepare for that. Well, here we go. Another black -- you can't have two. There we go. And so we're blowing a lot of stuff up, of course, and he's not making any friends right now, so he's going to spawn a helicopter and get out of here, and actually, the mercenaries like calling this "practicing field acquisition", and as you know, "Mercenaries 2: World in Flames" is harnessing the power of Quad-Core and Havok to create a lot of destruction and damage here. And Pandemic is really well known for creating the vast open worlds, where frankly, you can go anywhere, do anything, and most importantly for "Mercenaries 2", you get to blow everything up.

Another foundation for open world games is creating these very dense, detailed dynamic environments, and so you're seeing a lot of that here. And some of our early optimization, as you can see, the vast draw distances, and we have a really high frame rate as well. So there is where Matthias is going to get out and again, trademark "Mercenaries" style, he's going to thank the helicopter for the ride. Now if you look off to the distance a little bit, one of the nice things about "Mercenaries" is you have so many choices in terms of how you clear through obstacles, and for this mercenary, he likes going a little over the top and bringing in a few [air bombs]. Why not, right? And another nice thing about "Mercenaries" is you have so many choices. You have over at least two dozen different types of air strikes, over 130 different types of vehicles and weapons. At least this mercenary likes doing things in a certain way, so he's going to redecorate this particular neighborhood right now. Again, he's going to bring those two [air bombs]. All right. So there goes the neighborhood. I just want to thank you again.

Paul Otellini: Thanks, Josh. Fantastic.

Josh Resnick: -- Mercenaries 2. I look forward to working again in the future.

Paul Otellini: Good luck.

Josh Resnick: Thank you.

Paul Otellini: Thanks, Josh.

Definitely made things go bang. Okay. Shift from gaming to consumer electronics. This is an industry which has gone through a lot of change in the past 20 years. It's shifted from analog, where it was born, to digital in the '80s and '90s as new kinds of devices came out, but they were all fixed-function, they were all rather linear. There is as much change as I think is going to happen in the consumer electronics industry as there has been in the computer industry in the last 10 years, as it evolves over the next five years to incorporate digital technologies, and increasingly, to build Internet capability into every CE device on the planet.

As they do this, I think there's a number of things that are going to come to bear that will change the industry dynamics as consumer electronics evolve. And it all revolves around the implications of the Internet on these devices. First and foremost, there's a support-for-services new product models coming out. You can easily envision that in everything from a handheld device to say, a new version of a set-top. The interactivity requirement for entertainment is going to be a predominant factor. Content has to be able to be moved around. It has to be portable.

But as these devices look more and more like computers with better form factors, what they'll have is a difference in the development cycles of these devices. Computer development cycles are six to 12 months now. We've refreshed the product line once or twice a year depending on the products. In consumer electronics, that's not necessarily true. The development cycle times were, for televisions, multiple years, three or four years, and they ran them for three or four years. They're coming down to 24 months. We see them shrinking down to 12 months.

But perhaps the most important change is the role of software in consumer electronics. As this happens, software becomes, again, one of the long poles in the tent for development. Getting the software ready for these new devices, ensuring that it has the headroom in terms of performance for not just today's task but the task that are going to be coming down the Internet over the next few years, and also being able to take advantage of the Internet to do updates to the software, to fix bugs, deliver new features, deliver new services, create new revenue models. All of this is going to change the competitive dynamics in the consumer electronics industry, and we see that as we engage with customers -- traditional customers in this industry. What they really want is a product from us that can meet these needs, that can meet the classic needs of consumer electronics, but the emerging needs of performance-based Internet computing.

To meet these needs, we have a new class of products that are coming out next year. We call this iA consumer electronic system on chip. We start with a high-performance, low-power iA core, and then we add around that dedicated blocks that are focused on the things that consumer electronics need. Most importantly was a dedicated audio/visual pipeline to be able to get that high-quality experience. We're investing heavily in this area. I won't show you any demos of this today. We'll save that for CES in January. But we will deliver our first iA SoC silicon, codename Canmore, to the industry in 2008.

Let me shift to enterprise and problem solving. We've gone through the home, gone through gaming, gone through mobile. Now it's time to talk about business, and particularly technology in business. All of us in business see that our company's ability to differentiate itself is increasingly based upon the use of information technology. How we can employ information, data, analytics to solve new problems, how we can do it faster than the competition. It's becoming the new competitive advantage. And rather than talk about this, I thought we'd show you a short video that shows how three companies are using technology to create a competitive advantage for themselves.

[Video playing.]

Paul Otellini: The video shows how our customers are solving what they called extreme problems by applying lots and lots of technology. As we see other customers do this, we've seen a large demand growth for Quad-Core devices, and we've now shipped millions of those in all segments of our server marketplace. Two weeks ago, we introduced a [Intel®] Core™2 [processor] version of Xeon® [processor] into the MP segment of the marketplace, such that we now have Intel Core 2 [processor] and Intel Quad-Core [processor] top to bottom shipping in servers today. All the platforms we're shipping today are compatible with the 45-nanometer versions of these processors that are going to be shipping in the next couple of months.

But as we take a step back and look at what's going on, one of the major trends that we see is the increasing focus on energy-efficient computing. This continues to be important to both enterprises, in terms of power savings, but also to the environment, in terms of energy consumption. And to talk a little bit more about the collaboration between Intel and other partners in this area, I'd like to bring out an expert in energy efficiency. His name is Andrew Fanara, and he's on the EPA's Energy Star* team. Andrew, good morning.

Andrew Fanara: Good morning. It's great to be here at IDF.

Paul Otellini: It's good to be here -- thank you for being here.

Andrew Fanara: You know, the EPA, where I work, is really excited about the IT sector and about the tech sector because we really think that the tech sector can help to drive energy efficiency in many, many different types of products. And of course in IT products.

Paul Otellini: Well, can you tell us a little bit more about Energy Star and what it does?

Andrew Fanara: Sure, well, the U.S. EPA where I work, I help to manage the Energy Star program. And the Energy Star label, which is right up there, that's on products that are in the top echelon of energy performance, and those products are available in the marketplace. Those products do a couple of really good things. First off, they save dollars. They help reduce your utility bills. Second of all, they help reduce CO2 emissions, which are critical to climate change. And so we're very excited, in fact, that our newest Energy Star computer specification was just launched this July as Energy Star 4.0. Energy Star 4.0 means that the computer's more efficient in all its modes, it's got more robust power management, and it's available in the marketplace right now. So we're very excited about that. But there's also the Climate Savers* initiative, which is helping to even raise the bar even further.

Paul Otellini: I'm glad you mentioned Climate Savers. It's an initiative that we did with Google and the EPA this last summer. It's focused on driving energy efficiency not just into the data center with servers, but also into PCs around the world. And it's gone from being an announcement -- rather small announcement -- to now having 100 companies involved worldwide, and we're very, very happy with the progress on this.

Andrew Fanara: EPA really worked very closely with all of the members that initiated the Climate Savers program, and we're very excited because to inspire even further innovation in computing energy efficiency for the desktop, Energy Star 4.0 is the baseline for the Climate Savers initiative starting immediately.

Paul Otellini: Now, I understand you've done a recent study on energy efficiency in the data center.

Andrew Fanara: Right. The EPA was asked by Congress late last year if we would investigate the energy impacts on and from data centers on national energy consumption. Data centers are really critical, vital to national infrastructure, they're really important to business and consumers increasingly because of the things that they provide. And so we did this report. We found out that about 1.5 percent of national electricity use is used by data centers. That's roughly doubled since about '01. We expect it to double again in approximately 2011. You would expect that because the demand for what data centers do has continued to grow in this Internet-driven culture and this digital culture that we find ourselves in.

Paul Otellini: Yeah. Now I understand we're doing some collaboration on making this even better.

Andrew Fanara: Yes. As a matter of fact we thought what was the best way to demonstrate that you could have an impact on data center energy efficiency and help to drive that? So let's take a look at the Eco Rack.

So the Eco Rack was actually the brainchild of a colleague of mine at Lawrence Berkeley National Labs right across the Bay here. Mr. John [Coomey], who I think is somewhere around here -- maybe, John, you can raise your hand? There you are. Great. Thank you very much, John, for all of your hard work. [Applause]

John and your Intel engineers tried to figure out a way how they could develop a demonstration project that would be highly visible and really show data center operators but other folks within the industry how you can use standard technology -- today's technology -- to make a rack of servers a lot more efficient. Because of course there are lots of challenges in power space and cooling that data center operators are experiencing.

So they came up with a basic set of components for the system. There are a couple of key things that they tweaked. The first thing that they did is they're running DC throughout the rack itself. That gets rid of a lot of the messy AC to DC losses and improves the efficiency. So that's great. They did other things like tweaking the memory, which can be a big energy drain. Which is great.

And the most important thing that you have to remember is that in addition to this being standard, off-the-shelf equipment, there is no sacrifice in performance in terms of what the servers can do. It doesn't cost any more, which is also excellent, and I think that the last and most important reason why I'm here is to say that you can get 18 percent improvement in the actual energy efficiency of this rack. So that's great.

Paul Otellini: That's pretty amazing, because even in a rack like this that's over $40,000 of electricity savings per year per rack. You think about a data center that's got hundreds of these things. . .

Andrew Fanara: Right. I like to say that the cost of energy for a data center operator is no longer a rounding error on the financial --

[Crosstalk]

Paul Otellini: That's absolutely true, yeah.

Andrew Fanara: It's significant. So I think that this can really help drive TCO down and we're looking forward to showing this and talking about it with folks in the industry and to data center operators.

Paul Otellini: Great, and thank you for the collaboration and thank you for coming.

Andrew Fanara: Thank you very much. Appreciate it.

Paul Otellini: Bye-bye, Andrew.

[Applause]

Paul Otellini: The last topic I wanted to cover is extreme inclusion. And one of the things that happened this year, 2007, was that the industry finally met Andy Groves' projection of a billion connected computers around the world; a billion connected computers on the Internet. And as we did this, it becomes more and more important to think about what's next. How do we connect the next billion people, and then the third billion people after that? We've got some good line-of-sight on that next billion people. Their aspirations, their affordability, the connectivity, the incomes, the economies they live in are rising at rapid rates to where this is a realistic goal to be able to be included in the digital economy. But crossing that digital divide to that third billion is going to take a lot more work. And Intel is committed to this. We have a number of programs in this area. I'd like to talk about three of them today.

The first is a series of things that we call government-assisted PC programs, where we work with governments, country by country -- typically in less mature markets around the world -- and we work with their service providers and local PC manufacturers to create bundles to get PCs out to the citizens in a very affordable fashion, very often paid for over the monthly bill for their telephone. We've created 170 bundles now. There are 80 service providers that have committed to this. It's active in 60 countries around the world, and it's got good track record. This year, 10 million people will be brought into the world of computing through this program using four million new PCs sold through the program.

But looking forward we think that education and children is really where a lot of the activity has to be to get that third billion. And we have had a focus for a number of years -- a flagship program that's part of our foundation called Intel Teach. Intel Teach has now trained four million teachers around the world who have touched probably half a billion kids in their classes over time, to be able to bring technology integrated seamlessly into the classroom and improve the learning process. We are committed to continuing this and will have 10 million teachers trained by 2011.

But as we look forward, the real target -- the threshold for that third billion -- resides in the schools. These are the kids that are going to need this first, and be able to take advantage of the computing most over their lives. There's 1.2 billion children around the world in K-12 schools. Only 50 million PCs are available in schools to serve 1.2 billion. A rather long waiting line, and very often no computers at all.

So we think that one of the things that needs to happen is a new category of laptops has to be developed to bring computers into schools in a very seamless, cost-effective fashion. What we've done for that is we've built the Classmate PC, and we've shown this before. It's now in production. It's going into production around the world with OEM partners essentially purchased by school districts and governments around the world. It's tailor-made for kids, meeting their cost points. It's got teacher control software, etcetera, etcetera.

A second program that we've been pleased to join in the last few months is shown up here in the green computer, the OLPC, the One Laptop per Child Foundation. Intel has joined the board of this foundation. We are contributing technology and funds to this. And the good news is we share a common goal. The common goal is that we both believe very strongly in getting all these kids the opportunity to participate in the digital economy and the digital age.

And then the third thing that is newer that I wanted to talk about is how the industry is responding to this. This happens to be a new notebook from ASUS. It's called the ePC. It is a sub-$200, fully-functional notebook aimed at principally around first-time users and principally in emerging markets. So new price points for dedicated machines for the OLPC program, new technology, and now new capabilities coming into laptops around the world, all focused at addressing bridging this digital divide for the second and third billions.

So, I've talked today about a lot of things. I've talked about essentially a new area of growth for our computer industry. I've focused on mobility, on entertainment, on problem solving, and on inclusion. The common theme, in case you missed it, throughout all this was Intel architecture and increasingly a communication capability built around WiMAX. From my perspective, having been in this industry now for three decades, we've matured. Very often we're called a mature industry. But there's no shortage of opportunities that lie in front of us. In order for us to capitalize on those opportunities, though, all of us as an industry need innovations. Growth only comes through new products.

And I wanted to take you back six years to IDF 2001. This was after the Internet bubble had burst. The industry had its first down cycle ever in 2001. Remember, people were saying that all the servers that would ever be needed for the industry had been shipped, and it just was a matter of repurchasing them on eBay. The world changed. And I talked about a gentleman named Charles Kettering. And Kettering was the founder of Delco Electronics, which was then acquired by General Motors. And Kettering ran research and development GM for decades.

In the midst of the U.S. Depression in 1933, someone asked Kettering about the prospects for growth. And he said, "I believe business will come back when we get some products that people will want to buy." Six years ago, I said this is what we needed to do. We needed to break out of our mold. We needed to think differently. We needed to think about new products, new market opportunities, and figure out where people need computing to be applied in their lives and in their businesses.

We've accomplished a lot. The industry has grown by a factor of two in those intervening six years. But inside that are some more fundamental growth. Servers have also grown by a factor of two in those six years. But the capability of a given server six years ago versus today is night and day in terms of going from single core to quad core, from the energy inefficiency to the energy efficiency that we see in all brands of servers today. Perhaps the most stunning growth, though, inside this was mobility. Mobility grew by a factor of four in those six years, essentially because we as an industry invested in creating this market and in taking advantage of this market opportunity.

The second big change that happened in that time frame is focus on the environment, not just on the energy efficiency of the products that we all build, but also on the cleanness of them, moving to lead-free and now halogen-free. So as I look forward at the opportunities that we as a company have and you as an industry along with us have, I get pretty excited. Much of what we do that's really good -- making things small, making things affordable, making things in mass production, connecting them to the Internet -- is where new markets are going. It's where the world of handhelds is evolving, it's where the world of consumer electronics is going, it's increasingly what business needs. So I get more optimistic about our growth opportunity as an industry today than I was even six years ago.

But growing all this new technology and moving it into the mainstream doesn't come for free. It requires persistence, it requires innovation, it requires collaboration, and it requires vision. From Intel's perspective, you have my commitment that we'll continue to do our part in delivering this collaboration, and what this forum's all about this week, is asking you to join us in that. Thank you very much. Enjoy the forum.

[Music plays.]

[End of recorded material]