Patrick Gelsinger; Sr. VP. Digital Enterprise Group

[Beginning of recorded material]

[Video plays]

Female Voice: Hello, Pat Gelsinger. Welcome to the Intel Developer Forum 2008. You have 2,626 guests.

Pat Gelsinger: [Laughs] Welcome to my home, my Internet home.

[Applause]

Pat Gelsinger: Welcome again to IDF 2008, the world's greatest geek-fest. As you saw on the opening video, one of the topics I want to take up today is this next generation or the fourth stage of the Internet evolution. [The] video described [it] this fundamental idea from the ITU paper, the Internet of things. And I like to call it the embedded Internet.

The idea of the embedded Internet is a really big deal, and I'm going to start my keynote by talking about this idea. Then we're going to explore the Nehalem microarchitecture, the world's first fully dynamic, scalable microarchitecture. Then we're going to conclude with a discussion about visual computing. So let's get started.

From this stage over the years, we've had many Intel executives and technologists give future predictions. In 1998 at IDF, Andy talked about a billion connected computers, and at the time, a huge, huge statement, a billion of anything connected to a network. Craig Barrett revised this and looked forward to that billion-connected computers understanding the implications and benefits of the Internet as well as mobility. And Paul Otellini in 2003 revised those numbers yet again to look forward to 2.5 billion connected handsets, the role of mobility as well as 1.5 billion connected platforms.

Today I want to look forward, again, at what is the next big thing beyond those statements. A decade ago, the late Mark Weiser, CTO of Xerox PARC described the future vision of ubiquitous computing transforming our lives by these computing devices all around us. He said, "We are in a path which increasing availability of processing power would be accompanied by its decreasing visibility." As he observed, the most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it.

This, we believe, is a powerful aspect of the Internet of the future, the embedded Internet where every human on the planet is connected to the Internet 7-by-24 in every modality of life -- how they work, how they play, how they learn and even when they rest. This idea of the Internet we believe, of the future of the embedded Internet, will result in a stunning 15 billion devices connected to the Internet by 2015, an explosion in new devices and capabilities for the Internet of tomorrow. Let's take a look at a few of these devices and what they will result in.

Some of these devices will be personal. Devices such as you saw in the opening video, maybe infotainment system in the car or communications based over the Internet, possibly home security and energy management. Some of these devices will be associated with public places as well as revolutions in the infrastructure to support it -- smart highways, smart power grids, health networks that even help us while we sleep. The embedded Internet of tomorrow, pervasive throughout our environment and our world.

We do see there are a set of challenges to accomplish this vision. Things like reliability and long life. These devices, these sensors get embedded and you install them and never want to touch them again. They have to live forever. Software scalability deployed at hundreds of millions or billions of connectivity, low power or low cost, [reissuing] new issues in privacy and security, addressability and standards.

Intel believes this is a great market opportunity for us, and us as an industry together we're developing specific products to target this area of the embedded Internet, a big deal for our customers. And in particular, things like long life, having long-life processors available to our markets and, in fact, delivering new low-cost platforms as well such as Intel-based SoC platforms. And the last SoC that Intel did was the 386EX, the chip that I helped design many, many years ago. And long life for it meant that we shipped our last batch of 386's last year, almost two decades later. Now that's long life, and that's what this type of embedded market requires.

As we become more personal data, and our environment is more associated with the Internet, things like encryption become more critical, privacy and security required for this Internet of tomorrow, addressability. In fact, we've seen that there are predictions that say that by 2011 we will have exhausted the IPv4 address space, and how will we be connecting these devices of the future? We need to move to the next-generation Internet architecture, IPv6, something we've described for many years but are still [unintelligible] in its deployment.

And finally, many standards, standards that exist today and apply to embedded and new standards that must be developed for tomorrow. And at IDF we're known for this venue being a platform to drive and deliver those standards of the future.

Intel has long invested in the embedded Internet marketplace. And we've been able to have great success as our market has expanded, and we've seen increasing amount of revenue and business in this area.

In fact, last year we celebrated the Intel 30 years in embedded. And for that we developed the Embedded Chopper, and we're going to show that now. This is the most advanced and powerful bike that's ever been built. We partnered with Orange County Chopper to demonstrate the feasibility of embedded IA for the motorcycle. And it has things like GPS, fingerprint, Wi-Fi, WiMAX, front-end cameras, rear cameras, built-in PC that's dockable that's part of it. And immediately following the show here, Doug Davis, who runs this division for me, is going to be showing his Evel Knievel skills by jumping it over the Moscone immediately after the show.

[Laughter]

Thirty years in embedded for Intel, and we're quite excited about the progress that we've been able to deliver. And why has the Intel architecture been so interesting to embedded customers? Its scalability. Many machines are very high-end, like servers for Xeon, as well as low-power and low-end, purpose-built solutions as well, long-life requirements and a rich set of software tools required for this marketplace.

Today we've been selling and delivering our traditional IA products into the embedded marketplace -- Xeon, Core 2, Centrino products with great success, embedding and modifying and extending the life and support for the embedded customers.

With Atom, we open a bold new opportunity for embedded applications of IA. Atom, a fantastic, new low-power processor aimed specifically at Internet-access devices enabling a wide range of compelling Internet devices from MIDs, nettops and netbooks as well as broad uses for embedded customers as well. Atom delivers a fan that's less than 5-watt operation, specifically enabling small-form factor applications today.

In the short time, in less than two quarters since we announced Atom, we have over 700 design engagements underway in utilizing the Atom processor for embedded customers. Most of these are new business for us that we've never seen before, new opportunities for us working together as an industry, a diversity of designs in medical, entertainment and defense fields. Huge opportunity for us.

It's also purpose-built products such as the Tolapai processor that we announced last month, our first IA SoC since the 386 in 1984, purpose built specifically for communications and security applications with significant customer benefits available. And it's not just the [point] product, but it's a roadmap. And we see our commitment for all of our products becoming embedded products into the future as well as a roadmap of Atom as well as purpose-built SoCs.

And here in my pocket I have one example in our embedded Menlo products, which we'll be bringing to the marketplace beginning in the first quarter of next year as well.

What I'd like to do now is explore a couple of the specific embedded applications that we are looking at for the focus that we have around the embedded Internet of the future.

First, let's look at the demonstration that we began our keynote with. I came through the door as part of an Internet security system that was installed in the door. It did face recognition to me on the other side. And as soon as I entered and came to the camera, it was able to recognize me and allow me to enter the room. It has a full Adobe Flash implementation that's possible and something only available as part of the Atom family of products.

We also then have a complete home control system here as well that's possible. And in this case, let's go to my energy section. And it's plotting real-time our energy usage in here. You guys are generating a lot of body warmth, so we're turning the energy requirements down in the room here. Or we could go and flip over to our monthly statistics, so monitoring both locally as well as available over the Internet, energy stats and usages, also remote management of this over the Internet.

In addition, we see it revolutionizing the area of communications. An IP media phone is an example of this as well. And what we see is that new devices like Internet-enabled phones and media stations will simply revolutionize the communication category in the future. And if we can just give an example of the IP media phone by running the video now.

[Video plays]

Patrick Gelsinger: Very simply, for communications of the future we want to replace 100 percent of these with new devices like these, the open IP phone of tomorrow. Another area that we're exploring for embedded applications is in vehicle infotainment. Literally, the average American spends hundreds of hours per year in the car. We have 155 million OEM and after-market telematics are predicted by 2013. And in-vehicle infotainment is about the convergence of media -- audio devices as well as telematics for the platform of the future.

As we launched our open infotainment platform effort, we've been delighted by the amount of interest and enthusiasm from the industry in working with us to deliver services, middleware, boards, operating systems, and complimentary silicon specifically for the auto marketplace. As we've approached the key players in the industry, names such as BMW, Nissan, and Daimler, we've been delighted by their response and enthusiasm for the simple idea of the open infotainment platform, enabling them to quickly bring their new innovative apps to customers and to the marketplace.

To have us explore this a little bit more carefully, please join me in welcoming Alex Bush from the BMW Group. Alex?

[Applause]

Alex Bush: Good to see you, Pat.

Patrick Gelsinger: Very good. Thank you so much for joining us today, Alex. Can you tell the IDF audience what BMW Group thinks about this idea of open infotainment platform?

Alex Bush: Sure. The BMW Group is embracing this concept and is already working with Intel and other technology leaders to develop an open infotainment platform. Our new approach is based on a very simple idea of utilizing key hardware and software from the personal computing world. And the Intel Atom processor is a good example of one of the building blocks needed to do this for in-car entertainment features.

Patrick Gelsinger: That's great to hear, Alex. What is most important to BMW Group as you develop your IVI systems platform?

Alex Bush: The most important thing is to build a multimedia system that is reliable, can be easily updated, is intuitive for the customer to use, and of course one that can be designed very quickly. Open infotainment platforms will allow us and our vendors to be much more efficient in our design process, and then develop applications for our customers much quicker.

Patrick Gelsinger: Now, I know that we've been working together with BMW and Wind River on some of these concepts. And we have a long reputation at IDF that we demo things, right? We don't just talk about them; we show them, so can we do that here?

Alex Bush: Of course we can.

Patrick Gelsinger: You know, Alex, this is a pretty good-looking car here. I think it's about time for my BMW to get an upgrade. What do you think, man?

[Laughter]

Alex Bush: We have to talk about that later. Okay. What you can see here is the proof-of-concept IVI system developed by Intel and Wind River with support from the BMW Group. It is based on an Intel Atom processor and the Wind River Linux Platform for Infotainment, which is a Moblin-based automotive Linux [deck]. Let's take a closer look.

Patrick Gelsinger: I see Chuck here from our demo team is in the platform. Can you tell me a little bit about the in-vehicle experience, Chuck?

Chuck: Yeah, Pat. I've got a new navigation system here, the 3D navigation system from Planet 9 Studios. And we had to take your car for a little trip the other day, so we recorded our trip over to the Golden Gate Bridge, and recorded it for your audience. Our 3D navigation system goes way beyond just standard 3D navigation. It actually looks at points of interest along the way and sends all those points of interest to the back seat, and even goes to the Encyclopedia Britannica and puts that information on the seat also for our users.

[Demo plays]

Chuck: What you're hearing is that the navigation system is recognizing points of interest in San Francisco and sharing information with us to keep us more informed as we drive.

Patrick Gelsinger: That sounds pretty great. But when you have the family along driving around the city, you also have who's in the back seat. What's that experience like, Chuck?

Chuck: Well, Pat, if you're going to be in the back seat of a car, this is it. Great computer, great computer interface, Internet access, the ability to play and download movies. In fact, we're going to play one right now.

[Video plays]

Alex Bush: Since this system is based on an open infotainment platform, it is much easier to add applications such as email, calendaring, and social media apps on top of the current hardware. Since this car's IVI system shares the same technology base as the PC, it also gives us access to applications and services developed by the PC ecosystem.

Patrick Gelsinger: So, a real win-win for the automotive industry and the consumer. We're able to incorporate multimedia into the car faster; consumers stay more connected, informed, and entertained. Very impressive possibilities about the future of computing. And the Internet apps and video streaming that were shown here were done over a WiMAX network, as well. In fact, one of our other embedded applications was actually the WiMAX base station that we use to stream to the car here based on the [total pi] or the EP8579 integrated processor we introduced earlier today.

Alex Bush: The BMW Group recognizes the need to provide a broadband [pipe] to the car, and is already working with Intel to evaluate WiMAX. We will also continue our collaboration with Intel in the automotive infotainment space. This demonstration car is a great example of how we can leverage from commercial technology, and tap into a bigger pool of applications that allow us to build multimedia systems that keep our customers safely entertained and informed in their cars.

Patrick Gelsinger: Well, Alex, from the IDF audience I want to say thank you for the collaboration with Intel and the great progress we're making working together, as well as for the auto industry overall. And in particular, the last thing, my last assignment for you as you leave is deliver this to Nancy Lane in Beaverton, Oregon. Okay?

Alex Bush: We will consider it. Thank you very much.

Patrick Gelsinger: Thank you, Alex.

Alex Bush: Take care. Bye. [Applause]

Patrick Gelsinger: Very simply, in summary, we see the embedded Internet as being transformational in everyday life, every mode of operation, every human on the planet -- work, learn, play, and rest -- resulting in a stunning 15 billion connected devices by 2015, the fourth wave of the Internet, the embedded Internet.

I'd now like to look at the dynamically scalable architecture for computing of tomorrow, and this is the Nehalem microarchitecture. We talked much about the tick-tock development model for Intel, how tick-tock allows us to have predictable, effective innovation for the industry, how tick-tock enables us to deliver new technology exciting and yet predictable for the industry to build upon.

Before we dive into a look at Nehalem, the next member of that family, the next tock, let's do a quick update on a few other items quickly. First of that is the Intel Itanium processor family, or tock-tock family from mission critical computing. And our IPF roadmap is alive and well with Tukwila, Poulson, and Kitson, delivering the highest reliability systems and RAS, moving to our next-generation system architect with quick path interconnect, as well.

We can zoom in up here. We have a Tukwila wafer onstage. We're delivering this technology to our OEMs later this year, and expect their system delivery to the industry in the first part of 2009. Itanium processor development on track, doing well, and delivering for our customers and for the industry.

Also, I wanted to take a quick look at Dunnington. Dunnington, our six-core wonder, is the final member of our Penryn family, our 45-nanometer tick family. Dunnington has been shipped to customers since July of this year, and we'll be launching it in the marketplace next month. This completes our delivery of the 45-nanometer family or products from Intel. And if we could just zoom in on the Dunnington wafer up there, our wafer I showed you last time.

But today we're even more excited, because we're presenting to the industry the first performance results from the Dunnington processor, as well. As you can see, across a wide variety of workloads -- database, ERP, Java, and middleware, virtualization, we are delivering best-in-class results, superior results from Dell, from Sun, and Fujitsu. In particular, I want to draw your attention to the TPCC number from HP, a new record four-socket platform.

But today I'm most excited to talk about a new eight-socket IBM TPCC result, the first ever one million TPCC result from an Intel architecture platform ever delivered, breaking the one million TPCC platform -- an industry first. And if you look at the scalability from a four-socket HP system to an eight-socket IBM system, you see extremely good scalability, showing the benefits of the caching and the memory architecture of the Dunnington processor, unquestioned expandable server leadership for tomorrow.

Now I want to turn your attention to Nehalem, the world's first dynamically scalable microarchitecture. Just last week we announced a new brand for this, the Core i7 for our desktop products, building on the Core microarchitecture and its strong momentum in the industry, and extending the brand family with the Core i7.

What it means for developers is a great new product; for customers, a product they know now getting better. The i7 is the premier member of that core family of products. We'll be delivering the first products of the Nehalem family in Q4 of this year and filling out the rest of the family as we go through 2009.

In fact, one of the things I was very excited to show you today is the first silicon on Nehalem EX, and this is the expandable segment of the Nehalem family. So, this is our high-end expandable server, an eight-core wonder. So, the first showing of the eight-core Nehalem EX silicon. We'll be excited to talk to you more at IDF about this as we go through next year.

As we talked about last year, Nehalem, we disclosed many of the key features of Nehalem -- new instructions, new capabilities as part of that. Today, I want to give the next set of disclosures around that: grounds-up new microarchitecture, great energy, efficient performance; introducing hyperthreaded native, a four, an eight-core, monolithic designs, new innovations as well as part of that.

One that we haven't talked about yet and I want to spend some time on today is in the area of power management, and in particular a way cool feature that we call "turbo mode." To help me discuss the details of this, I'd like to invite Rajesh Kumar. Rajesh Kumar is an Intel Fellow. That means he's a certified smart guy. So, he's really smart. And more important than that for today's discussion, he's the key architect of the power management system on Nehalem. Rajesh, thanks for joining us.

Rajesh Kumar: Thanks, Pat. [Applause]

Patrick Gelsinger: Rajesh, what's this breakthrough in power management that's part of Nehalem?

Rajesh Kumar: Well, the key idea in power management is actually quite simple, is to shut things off when they are not in use. And we've been doing that for quite some time with something we call "power gate." The issue is that power gate only removes switching power. They don't take care of leakage power, which has become a dominant source of power in modern process technologies. So, that's what we have done. We have now invented something which can take care of all power. When things are not being used, power goes to zero.

Patrick Gelsinger: Okay. That's a pretty simple idea, shut things off. Why haven't you done it before, Rajesh?

Rajesh Kumar: Thanks, Pat. First of all, he's right. The concept really is truly obvious and has been there for at least a decade. But actually doing it was really hard. We had to develop an entirely new process technology around this idea of how to build a perfect power switch, something which has extremely low resistance when it's on, which is this M9 stuff, and something which has extremely high resistance when it's off, which is ultra-low leakage transistor. That's what we have done.

Patrick Gelsinger: This sounds like a key technology breakthrough that could only come as a result of our collaborative work between process and product design and development.

Rajesh Kumar: Absolutely. If you want to be on the cutting-edge in this area, you've got to be working like this with your process technology architecture and design.

Patrick Gelsinger: That's very cool. What about the overall system that it fits into on the chip?

Rajesh Kumar: Right. So, the power gates [are the drive frame]. We also have several innovative sensors to actually measure real-time power. Perhaps the most interesting piece is the brains behind the operation, which is this new power control unit, which is an integrated micro-controller which only works on power management. The algorithms are sophisticated enough that we took more than a million transistors just to build the power control.

Patrick Gelsinger: So, more than a million transistors just [in] the power control unit.

Rajesh Kumar: Yeah.

Patrick Gelsinger: Right. That's more than my whole 486. My design achievement isn't even the power control unit of Nehalem.

Rajesh Kumar: Well, there you go.

Patrick Gelsinger: Well, okay. [Laughter] What can you do with this?

Rajesh Kumar: Right. The most obvious stuff is energy efficiency. We can go for battery life improvement. We can go for [idle] server or datacenter reduction. All of that stuff is obvious. But the one cool thing I want to talk about today is how we convert this power [headroom] back into performance. That's what we're calling turbo mode.

The idea is if you have four cores and you detect that only one or two cores are actually active, we turn the power grid off, the power of those cores goes to zero, then the power control unit takes all of this power, gives it back into higher voltages, higher frequency, and we get a boost in performance as a result.

Patrick Gelsinger: So literally we're able to turn, dynamically, that available power budget into more performance?

Rajesh Kumar: Absolutely.

Patrick Gelsinger: Turbo mode. That's really cool. Now can we demo it? Can we see this in operation?

Rajesh Kumar: We should be. There is a demo. There you go. All right. So what this is going to demonstrate to you is, again, the concept that I described before. We have something [unintelligible] all the cores running, and we detect that all the cores are not really needed, but the application needs the performance, so the power control unit gives it back into performance. And as you can start seeing there, the frequency starts jumping up and that's what turbo does.

Patrick Gelsinger: So in this case, two extra bins of performance that come dynamically anytime the workload allows.

Rajesh Kumar: The more constrained you are, the bigger the performance boost will be.

Patrick Gelsinger: Wow, that's exciting. So, you know, the whole power system in Nehalem is very powerful, capable. Can you just sort of summarize it to the audience?

Rajesh Kumar: Yeah. I think this is a multigenerational thing. We're not building something for one generation. I think this is going to be a breakthrough infrastructure which will survive for quite a few years to come, which is innovative sensors, innovative [drive trains], and brains behind the whole operation. And, Pat, just a plug for your audience. If you guys care about what's exciting, cool, and innovative in Nehalem, beyond the area of power management, and want to know in more detail, there's a whole bunch of technology sessions. I'll be leading off with an overview of Nehalem innovation at 3 o'clock today. So come join us.

Patrick Gelsinger: Thank you, Rajesh.

Rajesh Kumar: Thank you.

Patrick Gelsinger: You get these guys out of the lab, they turn into salesmen. It's scary. So that's a look at a key new capability of Nehalem, its power management system, and a cool new feature, turbo mode. What I wanted to do now was look first at servers and then at desktops and some of the unique applications that are being enabled by this compelling new Nehalem microarchitecture.

In servers, Nehalem will deliver the world's most adaptable server platform, and I wanted to focus on three particular areas as we looked at that. One is in high-performance computing, next in content creation, and finally in virtualization.

High-performance computing demands incredible performance capabilities, but it needs so with high density and energy efficiency. It typically consumes incredible amounts of bandwidth and requires low-latency interfaces to that memory. And Nehalem delivers a three-channel DDR3, which provides a tripling or more of memory bandwidth into the socket.

It also gives a large footprint, up to 288 gigabytes of memory capacity in a 2P server, a stunning amount of workload possible. The world's most adaptable server platform. Looking at that specifically in high-performance computing, we go to customers like NASA. And NASA's thirst for computing is literally insatiable.

And I'm pleased to announce that NASA has just completed the first phase of a multiyear project called Pleiades. NASA is anticipating the delivery of thousands of Nehalems for the second phase. And NASA will be using this for petaflops of performance for key problem areas, like space exploration as well as climate modeling, which is the simulation that you're seeing running on the screen right now. And today I'm very happy to have Rupak Biswas, director of NASA advanced super computing division, here in the audience. And Rupak, just say hello, right? Working with NASA on some of the world's most computationally-demanding problems.

We're also seeing great enthusiasm from the content industry working with the key content creators, each of them seeing the opportunity that Nehalem offers and embracing and working with us. One in particular has been exciting -- Kevin Mack. And Kevin Mack is an Oscar-winning digital artist at Sony. Just to pick on a piece of his quote here, Nehalem enables the creativity of vast worlds of new complexity. So let's just see how Sony Artist Kevin

Mack is creating mind-blowing, three-dimensional animation experiences. If you could show his animation now.

[Animation plays]

Nehalem will enable artists using new usage models that literally are computationally out of bounds for today's hardware and software. This particular animation that we're showing, you know, its complexity of abstract 3D animation, right, is showing a black box containing forces [that are about an explosion] of evolving capabilities inside of this visual world. Now, seeing the output, you can only sort of envision the computational complexity. What I'd now like to do is show that rendering being conducted in real time. We're going to compare a Harpertown system to a Nehalem system.

And on my right, we have the Nehalem system. And on my left, we have the Harpertown system. And now we're actually going to go render that image in real time as we're going on here. And as that starts to develop on the two screens over here, as we benchmark this, we see that on this type of application, computationally very demanding, bandwidth very demanding. We're seeing the Nehalem systems result approximately 2X faster. And when digital artists look at this, like Kevin Mack, they say this enables new things that I've never been able to do before.

Nehalem is also designed to deliver next-generation virtualization capabilities. We demonstrated a flexible workload management, and Nehalem extends that capability with VT and flex migration. We're working to co-optimize the Nehalem software and hardware platform to enable new capabilities like dynamic power management, improvements in the software ecosystem, and performance management to the virtual machine environments. We're also working to improve I/O virtualization.

And one of the key usage models for virtualization is consolidating multiple VMs onto a single hardware instance. And when you do that, you very quickly run into an I/O bottleneck, as the software, the VM, has to handle the I/O of multiple virtual environments. And to help me discuss this problem, please join me in welcoming [Rich Fromer] from VMware. Hi, Rich. Hey, can you just tell the audience a little bit about how's the collaboration going on with Intel and VMware?

Rich Fromer: Oh, it's been great. Intel and VMware have been working together on a number of new technologies, including the demo we gave you at last IDF on flex migration, working with enhanced Vmotion.

Patrick Gelsinger: So, you know, I talked about this I/O virtualization problem. Can you talk about how we're solving it?

Rich Fromer: Sure. One area to address is CPU overhead. Now, with our existing virtual network adapter emulation inside of VM, we're able to easily handle multiple screens of network traffic at full-line rate. But that CPU overhead starts to become noticeable for high-speed I/O devices, such as 10-gigabit adapters. Now we will be showing today a new technology from VMware called [VM Direct Path], which reduces that CPU overhead. Now, we expect that technology to be available in 2009.

And what it does is it allows the virtual machine to bypass the hypervisor emulation for the ritual network adapter. It allows that VM to directly interact with the underlying physical device for frequent and performance critical operation. What this means is that there are now CPU cycles available for that VM and other VMs. So this could be of great benefit to network-intensive applications. But, you know, we couldn't do that without Intel VTD, which provides the necessary address translation and protection.

Patrick Gelsinger: So just to get it right, this is next-generation software working with our next-generation hardware and just showing the depth of collaboration between the two companies. So can we show this?

Rich Fromer: Absolutely. We have a demo up here. We have two identical Nehalem EP servers. Inside each of these servers is a VM running on an experimental version of VSX. Inside each of these VMs is a compute intensive workload, spec JBB.

Now, on one machine, over here, we're running a standard virtual machine using our standard virtual network adapter emulation. But on this machine we have directly assigned the physical network adapter to the virtual machine using VM Direct Path. And we're targeting both of these systems with heavy network traffic.

Patrick Gelsinger: So how much does it result in?

Rich Fromer: Well, as you can see, the virtual machine that is using VM Direct Path, its transactions per second inside spec JBB are at least 1.7 times that of the other VM. That's because there are additional CPU cycles available to the virtual machine that's using VM Direct Path. And this direct assignment is all made possible by Intel VTD.

Patrick Gelsinger: So very simply almost a doubling in performance specifically as a result of our collaboration and network hardware features. This is great, Rich. Thank you very much.

Rich Fromer: Thank you.

Patrick Gelsinger: And the systems that we're demonstrating this on, provided by Super Micro working closely with them, and also one of the other attributes of virtual machines is lots of memory footprint. And in the machines here that we're demonstrating, we were running, actually, the first demonstrations of DDR3 MetaRAM devices, an 8-gigabyte DIMM, as well as a 16-gigabyte DIMM were used as part of the demonstration here.

The large memory footprint of Nehalem systems being demonstrated in addition to cool new virtualization features. In addition to working with VMware, we're also working with the other hypervisor vendors and in particular working closely with Microsoft and Hyper-V. And they're aggressively taking advantage of some of the advances in P states and T state technologies that are part of the Nehalem platform and making that part of their offerings as Hyper-V.

So we've looked at the server applications for high-performance computing content creation. Let's now look at the applications of Nehalem technology specifically for the client offerings. What we see here is the Havendale systems. And we have a Havendale wafer on stage here that we're showing [a board] our new Ibex Peak platform. And what this shows is the migration from our three-chip partitioning to the two-chip partitioning. As we described at the last IDF, with the move to the Nehalem system architecture, we bring the integrated memory controller and graphics capability into the CPU socket and introduce a next-generation chipset, Ibex Peak, as part of that platform. This results in higher integration, leading performance, lower power, smaller footprint, and lower-cost system design.

One of the areas that we'll be delivering as next-generation platform is our continuing thrust to deliver new capabilities such as the Pro platform value proposition. vPro has been extremely successful in the business market and we're expanding this to consumer and small business as well. vPro has hit critical mass with over 60 percent of Fortune 100 companies now deploying vPro systems.

We've had great success as we're working with key customers to enable these new capabilities. A customer like the State of Indiana -- 25,000 vPro clients, 80 percent reduction in desk-side visits, improvements in power expense, over a four-year lifecycle period, 850,000 pounds of carbon dioxide reduction in their footprint as a result of deploying vPro. Huge savings as a result of vPro technology being deployed. Let's hear from one other customer, Verizon.

[Video plays]

Patrick Gelsinger: Very simply, seeing great response from customers deploying, and we'll be enhancing that with Nehalem-based solutions. Nehalem is the unquestioned performance leader; a performance user's dream is the Nehalem platform. It's sort of like Michael Phelps. The only records left to beat are his own and Nehalem is upping the ante on Intel performance records in the industry. And I want to show you just a couple of examples of our high-end desktop applications and what we're able to do with the Nehalem system.

Here onstage I show a demonstration of the next release of Lost Planet Colonies, a popular game. And with it we have two equally configured systems. Both of these are 3.2 gigahertz compared to 3.2 gigahertz. This is a quad-core Yorksfield compared to a quad-core eight-thread Nehalem. Same memory, same hard drive, etcetera, but as you can see by the frames per second calculation that are going on here, generally Nehalem is greater than 50 percent better over this application environment. It's showing the value and strength of the Nehalem microarchitecture.

Next to it here -- let me start the demo here -- we're going to run a well-known benchmark, the [CINEBENCH] rendering benchmark. You'll see all eight threads of operation of Nehalem, and this is running an over-clocked demonstration. And we're comparing the over-clocked operation of Nehalem versus the best over-clock operation ever achieved on a Yorksfield system. And we're seeing well over 30 percent performance benefit as a result of the Nehalem microarchitecture. Simply put, if you want performance, Nehalem is the system in high-end desktop and extreme edition products that you will want.

In summary, Nehalem is dynamic. It is the first dynamically scalable microarchitecture in the industry. And our discussion today on power management and turbo mode gives you just another proof point for that claim of dynamic. It's the world's most adaptable server architecture delivering unquestioned leadership performance. And the only records left that we had to beat were our own, and Nehalem handily beats those records as we begin delivering it to the marketplace beginning in 2008 and continuing through 2009.

The last topic that I wanted to take up in our discussion today, embedded, dynamic, and visual. At the last IDF, we talked about visual computing, and that visual computing requires a complete platform solution. Requires high-performance CPU, high-performance chipsets, advances in high-def audio and media capabilities, and it also ushers in the opportunity for a breakthrough new architecture. And that architecture is Larrabee. And please join me in welcoming Larry Seiler, uh, a key architect of our Larrabee effort, to the stage. Larry?

Larry Seiler: Hi, Pat.

Patrick Gelsinger: Good to see you. Great having you join us, Larry. So, Larry, can you say a few words about Larrabee and why software developers are just so excited about what you guys are working on?

Larry Seiler: Well, Pat, software developers need more programmability than they've been able to get from the graphics pipeline. Now, Larrabee combines the parallelism of a GPU with the general-purpose programmability of an IA-based processor. And the result is that Larrabee gives developers the freedom to do what they want.

Patrick Gelsinger: So what about existing applications? You know, if I had my application on DX or open GL, you know, will they be able to take advantage of Larrabee?

Larry Seiler: Well, of course. Existing games will run very well on Larrabee, but there's more than that. Larrabee lets us extend the Direct-X and OpenGL pipelines to add new features developers have wanted for years.

Patrick Gelsinger: So can you give some examples of those near features, Larry?

Larry Seiler: Okay, well, there's translucency. That adds a lot of visual interest to a scene. But this is what happens if you render in the wrong order. You see the small dragon's wing appears to pop through the large dragon's wing.

Patrick Gelsinger: So it's wrong, so literally it's been rendered incorrectly.

Larry Seiler: That's right and the solution is to sort the surfaces before you render them, but that can be very difficult to put into the application pipeline. So the result is we just do that sorting on Larrabee.

Now, here's an example of a video that's going to show us the kinds of scenes that game developers can create when they're able to use large numbers of translucent surfaces without having to worry about sorting. And as you can see, this allows gamers to visualize more of the environment that they're moving around in.

Patrick Gelsinger: So if I was a game developer at this point, I'd be pretty excited about the kind of capabilities this will enable.

Larry Seiler: Pat, the game developers I've talked to certainly are excited. Hey, here's another example. Now, shadows will help us to understand how objects relate to each other in space. And so the standard shadow algorithm can produce jagged-edge artifacts as you see right here. Now, developers have to detect the cases this might happen and then find ways to work around them. But there's a recent algorithm called irregular Z-buffer, which always draws smooth shadow edges. The problem is this requires complex data structures and adds extra stages to the graphics pipeline.

Patrick Gelsinger: So can you show us a real example of this?

Larry Seiler: Well, sure. Here's a video that demonstrates moving around in a scene with lots of complicated shadows. And Larrabee's general-purpose programmability makes it possible to use an algorithm like irregular Z-buffer to guarantee that all those shadow edges are absolutely smooth. And as a result, software developers are free to focus on their content and not worry about trying to fix the bugs in the image.

Patrick Gelsinger: So what else does Larrabee support? There's lots said about new techniques, like ray tracing? What about those?

Larry Seiler: Oh, well, absolutely. There's no end to the possibilities that Larrabee enables. Now, I was at the SIGGRAPH computer graphics conference last week to present a paper on Larrabee, and the developers and researchers there were absolutely excited about the opportunities that Larrabee provides.

Patrick Gelsinger: Yeah, there was a lot of press and interest coming out of SIGGRAPH last week. Is it possible for you to summarize for the audience what you talked about there?

Larry Seiler: Pat, at SIGGRAPH, I did it in 50 seconds.

Patrick Gelsinger: Fifty seconds.

Larry Seiler: Fifty seconds.

Patrick Gelsinger: Okay, everybody, right, you ready? Can you do it for us, Larry?

Larry Seiler: Okay, are you ready?

Patrick Gelsinger: Fifty seconds.

Larry Seiler: This is fast. Listen carefully.

Patrick Gelsinger: Go.

Larry Seiler: Now, people say that GPUs are more computationally efficient than CPUs. Larrabee combines Intel processor cores with wide-vector processors for high-computational density. People say that bin rendering doesn't work as well as a standard rendering pipeline. Larrabee uses binning to remove serialization bottlenecks. And as a bonus, this reduces the memory bin needed for rendering.

People say that graphics requires more programmability, and they are right. Just like a CPU, Larrabee can compile and run shaders and, through our native interface, can run standard C and C++ programs and it supports virtual memory and fully coherent caches. Are we crazy?

Patrick Gelsinger: Well, it looks like you did it in about 40 seconds, Larry, and, no, you're not crazy. You're not crazy at all. In fact, historically, what we've done with the Intel architecture is we've been evolving it over time. And if you think back over the history, new workloads appear, special-purpose hardware often appears for them, and eventually we upgrade the platform, extend the instruction set, to comprehend those new capabilities like floating point, like media processing on the platform. And that's now exactly what we're going to do with Larrabee and its new instructions as we bring it into the platform and into IA over time.

Larry Seiler: That's right, Pat. And that's why Larrabee will fundamentally change computer graphics. So show up at our class on Thursday and you'll hear the latest news, that Larrabee's the way for emerging GPUs and CPUs. Thank you, Pat.

Patrick Gelsinger: Thank you very much, Larry. Thank you, very good, bye. Okay. Like I said, you let them out of the labs, they turn into showmen. So in addition to the additional classes that you have here, I'll point you to Renee James's keynote tomorrow, where she'll be specifically talking about software tools, the developer environment that we have for visualization, not just for cores and our multi-core and many core roadmap, but also how we're extending that toolset for Larrabee as well.

In summary today, we've talked about three very important and profound topics. First, the fourth generation of the Internet, the embedded Internet, 7-by-24, touching every modality of human life, to every human on the planet, and enabling an enormous explosion, 15 billion devices connected to the Internet by 2015.

We then looked at Nehalem, the first dynamically-scalable microarchitecture and the huge benefits that it will bring to both clients as well as server infrastructure. On track for first production in Q4 of this year.

And finally we explored visual computing, a revolutionary new architecture for visual computing, Larrabee, and how it is generating huge excitement and interest from the ISV community. And three areas today, we discussed waves of new innovation coming and the opportunity that it presents for all of us working together as an industry. Thank you very much.

[Video plays]

[End of recorded material.]