GELSINGER TRANSCRIPT

9/27/06

Female Voice: Ladies and gentlemen, please welcome in senior vice-president and general manager, digital enterprise group, Pat Gelsinger.

[Applause]

Pat Gelsinger: Good morning.

Pat Gelsinger: You know, it's always a pleasure of mine to come to IDF to meet with you as an industry to talk about great technologies and really invent the future together. What a six months it has been since we were together last. We've just turned over the entire company, the product line, everything that we're doing in the last six months. In high school, I was a middle distance runner, sort of a half-mile, and you used to pace yourselves as you went through the 800 yards. And those last 100 yards, right, you were always racing for that finish line and the tape. That sprint to the finish line. This last six months, it's just been one flat-out sprint, as we've had an incredible summer of 40 new processors coming to the marketplace. And that has been enabled by our relentless pursuit of Moore's Law.

In the recent proceedings of the IEEE Solid State Circuits Society Newsletter they dedicated the entire issue to articles about the past, present, and future of Moore's Law. And I was pleasured to be able to have one of those articles, and I talked about the '70s, the era of innovation and invention in Moore's

Law; the '80s, the era of scaling; the '90s, the era of manufacturing; and today, the era of efficiency. The entire industry recognizes the importance of Moore's Law and its impact on what we do. Intel is leading the ecosystem, so developers, manufactures, end users, and Intel benefit from our relentless pursuit, our impassioned embrace of Moore's Law.

Today we wanted to look at those results of that Moore's Law and that embrace in our amazing products, how we've done them efficiently, how we've partnered with the industry to deliver new capabilities for the customer. But before we dive into that, I wanted to just take a quick look at what we've done over the last six months, beginning with the Xeon LV, we've launched on March 14, a processor designed for high-density bladed servers, storage devices, telecom, and compute platforms. The Xeon LV.

We followed that on May 23 by the Xeon 5000, our first 65-nanometer, dual-core server processor, 2x the performance of the then mainstream dual-core processors, excellent price performance. We followed that on June 26 by the Xeon 5100, formally known as the Woodcrest, the first product based on our core microarchitecture, a first for Intel, bringing server products as the first to deliver a new microarchitecture to the industry, performance, and energy efficient leadership.

We followed that by the Montecito, the Itanium 2 9000 product introduction, the first dual-core Itanium, 24 megabytes of cache,

a stunning 1.7 billion transistors on this chip, doubling the performance while reducing power.

Then the Core 2 Extreme, bringing the Core 2 processor to the client platforms, outstanding performance for gaming. We introduced it as unlocked and overclocked to greater than 4 gigahertz, over 33 percent scaling.

On July 27th, we introduced Core 2 Duo, based on our core microarchitecture, bringing it to mainstream clients for desktop and for mobile, up to 40 percent boost in performance, 40 percent decrease in power; an incredible engineering feat.

This was followed on August 29th by the introduction of the Xeon 7000, formally known as Tulsa. Dual-core MP, including both high-performance as well as low-voltage SKUs, a stunning 16 megabytes of cache for unquestioned performance leadership in MP servers.

Next was the vPro platform, greater efficiency, productivity, manageability, and security; our first platform to fully embrace the Core 2 Duo, delivering capabilities that IT needs at the heart of their business problems. We see it is three pieces of silicon that form the platform. It's the Core 2 Duo, the Intel gigabit network connection, and the Broadwater or Q965.

Today, we're announcing the Xeon 3000, the Core 2 Duo for the single processor server platform, designed for energy efficiency, workstations, and entry-level platforms.

As Paul described, we're bringing the first quad-core server products to the marketplace, the Xeon 5300, formally known as Clovertown. It may be the most exciting news of IDF. It's the first quad-core processors being delivered in server platforms. The first IA, quarters ahead of competition in delivering quad-core and server platforms.

We'll conclude with the first quad-core in the client platforms, the QX 6700, formerly known as Kensfield, available this quarter for rendering financial service, high performance gaming, multithreaded workloads.

When I look at a wall of wafers like this, I'm just overwhelmed. The amazing technology feats, the capabilities that we've brought forward to the industry, and seeing how you, as our partners, have embraced and worked with us to bring an amazing amount of technology forward to the industry. Ladies and gentlemen, the most impressive wall of wafers you'll ever see in your history.

[Applause]

Pat Gelsinger: The crown jewel of our work this summer has been the Core 2 Duo. July 27th launched an amazing new microarchitecture,

exquisitely executed, taking advantage of the most advanced transistors on the planet. Our simply massive 65-nanometer manufacturing effort. We've brought this architecture top to bottom in our families -- from our high-end Xeons to vPro clients to Centrino mobile platforms, delivering incredible improvements -- 40 percent of performance, 40 percent decrease in power, over 500 client designs taking advantage of this platform, over 200 server designs. The response has been simply overwhelming. One of my favorite quotes is this one from AnandTech that says, "It didn't lose a single benchmark -- not one. This is the unquestioned performance leader in the industry today."

We've applied this to our server product line as well. This is the Xeon 5100 series platform, core microarchitecture, delivering up to 3X performance over single-core platforms, 3.5 times the performance per watt. The Bensley platform, our fastest ramping server platform ever in our history. As Paul said yesterday, over a million of these platforms have already been shipped in the industry; a tremendous embrace and response to a great new product.

Let's hear what customers have to say and how they are taking advantage and seeing value in this platform. If you could play the video now.

[Video plays]

Pat Gelsinger: So, incredibly positive response from customers embracing and taking advantage of this platform. In the spring, I gave a head-to-head demonstration of this versus our direct competition in this space, and we showed a great result. So let's take a look at how we're doing since then. I'd like to invite Chuck up on stage, if Chuck could help me with a demonstration here.

Pat Gelsinger: Good to see you today. What have you got here for us?

Chuck Duvall (Intel Demonstrations Team): Well, you know what, I've got two servers that I have pretty much made as identical as I possibly can. They have the same 2U chassis, same hard drive, same amount of memory, same power supply, and we're going to see how our new Intel Xeon processor compares to the very new Socket F response from AMD.

Pat Gelsinger: So this was their response to our great new product?

Chuck Duvall: We came out with a product. This is their upgrade to compete against.

Pat Gelsinger: So if I remember right, in the spring, we were almost 40 percent faster. I think 38 percent faster and about 5 or so percent lower power.

Chuck Duvall: Across the board correct.

Pat Gelsinger: So how are we doing today?

Chuck Duvall: I've got this application here. Let's take a look at the power real quickly as they sit in an idle state. So you can see, the Intel is at 245 watts, the AMD is at 306 watts, just sitting.

Pat Gelsinger: Sitting there. Okay.

Chuck Duvall: Now, you never want to make a server just sit, so let's start this little benchmark application here. And this is a little application that really does a matrix-to-matrix multiplication table, does it over and over again, increases the size, and is really going to peg the CPU and show us exactly how this application is performing on these particular machines here.

Now, again, we came out with this two or three months ago, and this is AMD's response to our architecture. So you can see both machines are pegged out pretty much at 100 percent. They're cranking away, going through this multiplication table. They'll build one, calculate it, make another matrix that's even bigger, and calculate it out. Now, the Intel machine is done.

Pat Gelsinger: It's done in 40 seconds.

Chuck Duvall: Forty seconds.

Pat Gelsinger: We're still waiting for the competition here.

Chuck Duvall: It's still cranking away.

Pat Gelsinger: And you ran this beforehand, and what were the numbers?

Chuck Duvall: Well, it's going to run to a little over a minute, like a minute seven or so.

Pat Gelsinger: Okay.

Chuck Duvall: But the important thing, also, is let's take a look at the power while we had that full workload going. You can see that we peaked it at 370 on the Intel machine but the AMD machine went all the way up to 440 watts. So really, almost 70 watts higher.

Pat Gelsinger: Okay, so we're at 40-plus percent faster and 17 percent lower power compared to their best offering.

Chuck Duvall: Exactly. It's an awesome system.

Pat Gelsinger: There was no response.

Chuck Duvall: Not yet, no.

Pat Gelsinger: Okay. Unquestionably the winner. Paul talked yesterday about quad-core. Well, I know the audience is impressed with the 5100 and that's looking great.

Chuck Duvall: It looks very good.

Pat Gelsinger: And we've told customers it's directly socket upgradeable. So show me the quad-core.

Chuck Duvall: Right here?

Pat Gelsinger: Right now.

Chuck Duvall: Right now. In front of everybody?

Pat Gelsinger: Yeah. A little bit of risk-taking here. Upgrade the machine. You said it was socket compatible.

Chuck Duvall: Can you give me a couple of minutes?

Pat Gelsinger: How many minutes do you need?

Chuck Duvall: Give me seven, eight-ish.

Pat Gelsinger: Seven minutes? Okay. We're going to watch and Chuck's going to upgrade and we're going to do the quad-core demonstration on exactly the same platform and see how we do.

Chuck Duvall: Okay, I'll get on it.

Pat Gelsinger: Okay, go for it, Chuck. [Applause] So get your stopwatch out and we'll see how Chuck does here. The response from

customers has been great. The response from our partners and OEMs has been fabulous as well. And it's a pleasure today to be joined by Susan Whitney, our longest-standing partner relationship in the industry. Please join me in welcoming Susan Whitney, the general manager of System X for IBM.

[Applause]

Susan Whitney: Good to see you, Pat, and congratulations. That is a very impressive array. Congratulations.

Pat Gelsinger: I love that stuff.

Susan Whitney: That is great.

Pat Gelsinger: Thank you, Susan.

Susan Whitney: Thanks very much for having me. I really appreciate the opportunity; this is a great topic, a great event obviously for our colleagues. Also want to thank you for the opportunity to be at your Woodcrest announcement back in New York City. That was also very good.

Pat Gelsinger: Thank you for joining us there.

Susan Whitney: Well, that was also fun. And it really helped both of us launch our products and it turned out to be a very exciting summer. As you said, Paul quoted yesterday. In fact, what I'm excited about

is the acceptance of our new Woodcrest-based products. And when I sit down across the table and talk with customers, what they are telling me is they are experiencing world-class leadership performance -- and with 40 percent less power. So that is a great testimony to Intel, so thanks very much.

Pat Gelsinger: A great proof point. But we're also entering a new phase of the relationship of the companies as well.

Susan Whitney: Absolutely. And I think it would be important to put our partnership, which has been going on for years and years, in a new context in terms of what's contemporary and what's the marketplace looking for. And so I thought a way to position it is to talk about IBM's systems agenda. This is our commitment and our agenda, the things that we're focused on, for servers and storage. And it's based on three pillars: Collaboration, openness, and virtualization. And our partnership, Pat, has been going on and continues to go on and we have very strong proof points across all three dimensions. In fact quite honestly, we couldn't do it without Intel.

So let me just talk for a moment about collaboration for innovation. At IBM, we're smart enough to know that we can't do it alone. We do depend upon partners such as Intel. And I think the perfect, probably maybe one of the best proof points is our blade center. We started working on that together, probably now five years ago. Came out with a design that I think set the

standard in the marketplace. And the fact that it continues to enjoy market share leadership, I think is a good proof point.

Now you add to that blade center, you add Woodcrest performance, 300 percent increase and a 40 percent decrease in power and cooling requirements, and you have a winning combination and a winning solution. But we didn't stop there. We're equally, both of our organizations and our companies are committed to openness, because that delivers real value to you. And so we took our blade center design and we published the specifications. We published them and made them available. So what does that mean? That means, today there are hundreds of developers around the world building content for blade center. They're building accelerators, they're building solutions, snap-in capabilities that allow you to have more choice. And that's at the hardware level.

Then we went even further. We said, okay, IBM and Intel became founding members of blade.org. Blade.org says let's not just collaborate and publish standards at the hardware level, let's go further up the value stack. And let's work in an industry consortium to build solutions that leverage all of that ecosystem and leverage all of that content.

And then another example would be PCI and PCI Express. Clear cases of where we work together on industry standards. And I believe, based upon what customers are telling us, they value as a result.

And then virtualization. Clearly, to help you build the best infrastructure, get the best return on investment, we have to virtualize. And you take Intel's enterprise-class reliability with IBM's mainframe experience and you get systems that are architected specifically to be scaleable and to virtualize. In fact our Tulsa-based system now -- we've done hundreds of studies, and what we've been able to demonstrate is that we can manage three times the amount of virtualized workload, running at a peak performance. That's because of the capability and the scalability of Tulsa.

Pat Gelsinger: Yeah, that's fabulous. Our blade work, our dual-processor work, and we're really excited about how IBM is embracing the 7100 or the Tulsa. The performance improvement of that, our delivery of that product, and you guys were ready as we brought it to the marketplace ahead of schedule.

Susan Whitney: Yeah.

Pat Gelsinger: Right, the capabilities and performance, the improvements in performance per watt, the platform compatibility -- you can drop that immediately in so customers can take advantage of it quickly. RAS features, 16-megabyte cache safe technology. And we've seen great performance from 10 to 60 percent performance improvements on that, and I think in the performance area, you have some exciting news for us today, Susan.

Susan Whitney: Yes. I'm very excited to say that we just published a benchmark, a TPC-C benchmark, which in fact delivers 20 percent greater performance than any four-socket X86 in the market.

Pat Gelsinger: World record.

Susan Whitney: World record, I agree. So we are excited about that. [Applause] In fact, that really represents the 101st number-one benchmark, and it's because of the Tulsa, the 7100, because of the architecture, because of our success in the marketplace with enterprise-class applications that, Pat, we're committing -- we are going to invest in the fourth generation of our enterprise X architecture. We're going to be ready for quad-core, and as you know, this is uniquely, exclusively on Intel-based processors.

Pat Gelsinger: That's fabulous. And, you know, our joint commitment to this portion of the marketplace, to the MP marketplace, is substantial and we've seen great success together in that area. But one thing I wanted to ask you about, Susan, you know, power is really an issue to our large customers as well, and what's IBM doing to address that problem of our customers as well?

Susan Whitney: Very good question, and clearly on everyone's mind. It starts with the processor and the 40 percent reduction, but then there's an opportunity, architecturally, to build on top of that. And so recently we announced our cool blue portfolio of offerings, which really helped to manage power. We call it our

energy-management framework. And it's built on the Intel demand-based switching capability and technology.

This is three dimensions, if you will. The first is the ability to budget. And we use a power configurator, an online tool that allows you to budget how much power you're going to need. Not based upon a typical configuration, but based upon your specific configuration, so that allows you to budget. Then the next step is, okay, now you're actually executing and running, how can you ensure that you can maximize the performance within any given power envelope? Not just throttle down, but within whatever power envelope you have, how can you ensure that you can maximize the performance? And we announce Power Executive to be able to do that.

And then the third element is, in addition to the processors and the servers, what other components can we add to physically reduce the power requirements? And a good example of that is our rear-door heat exchange system, which takes chilled water from the air conditioner and, on a rack basis of a few thousands dollars per rack, to be able to cut down the power and cooling by an additional 55 percent.

So those are the three elements, and I'd like to just spend a moment on the power exec, because this is an industry-unique capability. A CFO has the tools and the metrics to be able to manage cash flow. This is for a CTO or a CIO to be able to manage energy. And what it does is, it tracks how much power

are you using -- not against what the labels add up to, but based upon your configuration and your applications, how much power are you actually using? How does that power trend over time? And you can measure that in 8-minute intervals or months, so you get that management information. You can do it at the server level, or you can do it at the rack level.

And the last piece is it can also measure, and it does measure and provide you information on temperature, ambient and exhaust. And again trends that. So it is, like, a CFO has lots of tools. It's time that we have the tools for the CTO and the CIO to be able to manage this critical resource -- not cash but energy, which translates into cash. So we're excited. We've got the processors, we've got the architecture, we've got the management tools, and we're excited because it's resonating in the marketplace. And that's the most important thing.

Pat Gelsinger: Oh, this is fabulous, Susan. We're happy to see you take advantage of our dual processor, our blade work together, our MP products, the virtualization, and the power work as well is uniquely tied to our demand-based switching architecture at the platform level as well.

Susan Whitney: Absolutely.

Pat Gelsinger: Overall, a tremendous partnership.

Susan Whitney: Pat, thanks very much, I really appreciate it. Thanks. Thank you.

[Applause]

Pat Gelsinger: As we've suggested, we're on track to deliver the industry's first IA quad-core processor, 100 percent socket compatible, exceptional engineering, we've accelerated this offering into the marketplace. This is the right way to do quad-core on 65 nanometers. This offers a whopping 8 megabytes of smart cache in the platform. We utilize smart-die selection technology through our manufacturing; allows us to select the high-frequency or the low-power parts, producing the best quad-core product for the different segments of the market at a lower cost, compared to a monolithic solution, 20 percent higher die yields, greater than 10 percent lower manufacturing cost. It allows us to leverage our enormous manufacturing capacity of the Core 2 Duo platform, given its socket compatible, but also offers investment protection through direct upgrade and ultimately delivering millions of units before the competition even responds. But, in result, expected greater than 4.5x increase in performance compared to our single-core platforms and provides over 50 percent performance increase compared to our dual-core processors. And it drops into the same platform. So let's see how Chuck did.

Chuck Duvall: Hey, Pat. Well, we're done.

Pat Gelsinger: We're done? How'd we do?

Chuck Duvall: Well we did actually really well. We're up and running again. You can see we've got all eight cores showing. If you want, I can run the benchmark again and see how we compared against the AMD system. And I think I finished in under my 7 minutes.

Pat Gelsinger: Okay. How long did it take?

Chuck Duvall: I didn't start myself, but, hey, I think it was under 7 minutes.

Pat Gelsinger: Okay, so we did the full upgrade into the same running system here, a little bit of risk taking, we're now back running the exact same benchmark, and we'll see how we do. Now, these were the low-voltage parts as well.

Chuck Duvall: Right, so, yeah, what I've put in here were the low-voltage parts, so these were actually almost 900 megahertz slower than Woodcrest or the Xeon system that I had in here earlier.

Pat Gelsinger: The same thermal envelope as we had the dual-core parts, we've now put the quad-core parts in place.

Chuck Duvall: Exactly. And so you can see we finished quite a bit faster, almost 20 percent faster on a system with the same thermals. If you want, we can go to the power here and see that we're at about 380 watts, so we gained like 3, 4 watts, but we're right in where we were at before, and our performance got 20 percent better.

Pat Gelsinger: This is stunning. So now we're 17 percent lower power and greater than 50 percent higher performance than the competition.

Chuck Duvall: This is what you needed, right?

Pat Gelsinger: Absolutely. You rock, Chuck.

Chuck Duvall: Good to hear. Thanks, Pat.

[Applause]

Pat Gelsinger: And as Paul promised you yesterday, well be delivering this in volume. We'll ship hundreds of thousands of these units before the end of the year. In fact we're going to ship over one million units before the competition ships one unit. This is an incredible way to deliver value to our server customers.

In July we launched the Itanium 2 9000, the Montecito, a new level of mission-critical performance, increasing up to 2x in performance at 20 percent lower power. We also, 366 days ago, launched the Itanium Solutions Alliance, and since then we've seen a great response from the industry. We've now crossed 10,000 applications in the Itanium platform, 2x what we had a year ago. The next generation of those applications is being announced and brought forward to the platform. In particular we've had Oracle announce this week that they are bringing

their next-generation database as well as their widely accepted fusion middleware onto Windows and Linux with the Itanium platform. We've also this week just had our 100th ISA member join, and that's JDA, databases for retail customers.

And we've seen the industry is continuing to innovate on this platform. And let's take a look at one of those innovations here. And I have a Hitachi Itanium server here. And what we have here is, well, we're just using a little bit of the rack, right? We have four dual blades here, so we have an eight-processor or an eight-socket system here, right, and as you can see, we can fit eight of these into the rack, right, so we've got eight times what I'm demonstrating to you this morning in this rack here. And with it we have two physical partitions, isolated platforms, inside of this rack. One of these, as you see in the picture, is running a hypervisor, which then has on top of it seven SQL servers, so four processors running seven SQL servers in this hypervisor environment. And the second, with the same virtualization manager, right, takes four processors and is running a set of Linux workloads. So we see seven Windows machines, and if we go to the first slide here, right, you see the full picture of the manager. Right, so seven of these are running the SQL server database. Five of these are running Linux stacks on top of the platform with three of those running the Transitive Quick Transit software. On those we are running native SPARC code, so unmodified SPARC binaries are running on three of those partitions. So, on this thing you see the system partition of all 12 of the virtual machines. If you think

about the amount of work that's going on in that density – an incredible consolidation of a CIO's workload is going on.

In the second picture we see the Windows physical partition. If you look at the performance in the lower graph we're seeing greater than 90 percent performance utilization on these homogeneous platforms. You see the load that each of those seven SQL machines is seeing in the top graph. So, high efficiency, looking absolutely great.

Now let's go look at the second partition. On this we're seeing the five Linux partitions. And because of the heterogeneous nature of those we're getting greater than 60 percent CPU utilization. That's pretty good. The other graph is showing the response time, which remains very good. This is the web query response time or [for] hitting the website. In the lower graph you're seeing the overall performance of the platform.

The thing that I wanted to point out to you -- and this is the greatest part about this demonstration -- this is running, through the Quick Transit code from Transitive, native SPARC binaries. What you are seeing on the WebSphere MQ and DB2 is the highest performance SPARC machine that has ever been demonstrated. Ladies and gentlemen, the highest performance SPARC machine in the industry, the Itanium platform that you're seeing here today.

[Applause]

Pat Gelsinger: Somehow I think Scott McNealy may not enjoy that as much as I did. Amazing products done efficiently.

What we saw in the early part of this decade, as we started to see the mobile market emerge and the importance of battery life, was we had no good way for the industry to look at how to measure thermal and battery life. In fact, there was not just confusion, but probably a bit of obfuscation going on by different vendors to look at this problem. Much like the auto industry with miles per gallon, we needed some way to look at this in a consistent and metric manner. We worked with the industry through BAPCo to form MobileMark. This is now an industry accepted way to look at performance and battery life. We have no equivalent today for clients.

We're very happy that this week BAPCo has announced that they are developing -- and we are working and part of the committee is working on this inside of BAPCo -- to form EcoMark, which will again look at performance and energy consumption in a consistent and structured manner. We're also excited that ECMA, the European agency, has joined with BAPCo to agree on this as a standard for both Europe as well as for the BAPCo solution. It's great to see the industry coming together, as you see by the quote from their announcement here.

We're also working to establish those same metrics for the server platform. How do you measure energy efficiency? Is it TDP? Is it idle? Again, confusion. We're working with the SPEC group, and we formed, inside of SPEC, a working group to specifically agree on how we measure energy efficiency and performance in server platforms as well. We expect the results of this to be available in the first half of '07.

Let's look just a little bit more closely at EcoMark. What this is, is trying to take a representative way of how a real PC would be used in the workday of a user and also measure and model the energy consumption in that workday as well. So it's real workloads through a real workday measured at the platform level trying to report energy costs in a meaningful way. As you see by the quote here, the industry is excited. We need some consistent way to look at this. Let's back up a slide.

What we are doing now is taking a workload -- and we've suggested SYSmark, the BAPCo model, and broken it up over the day. People work, take breaks, take lunch breaks, and go home at night. We look at how that workload would occur and the actual measured energy costs on that platform through this workday. So in this case, if you look at performance and energy cost, we see that if we take both the SYSmark score for these platforms, as well as the annualized energy costs for the platforms, we result in these kind of results. We see that there are annual energy costs of $23 a year -- and this is measured using the North American average of 10 cents per kilowatt hour

-- and the representative performance levels of those platforms, and providing a consistent way that customers could look at both energy cost as well as the performance of those platforms as they make buying decisions.

Intel knows that it's going to take a while for the EcoMark to be finished. What we're committing to do today is to begin measuring ourself and publishing these results -- and we're releasing a white paper today on this modeling methodology, as a way that the industry can start measuring us and our platforms and partners against the EcoMark directions for the industry.

Similarly, we're working with the industry on server energy efficiency, and we're very happy today to have with us Mark Bramfitf from PG&E. If you happen to be a Northern California customer, it's interesting that you can qualify for this program today. If you are upgrading a data center today, you can receive, just like when you get a rebate on buying an energy-efficient refrigerator, you now get a rebate on buying an energy-efficient data center in Northern California. You might not have known that. But we're working with PG&E as well as with our customers, and these Xeon 5100 products qualify for that up to a $350,000 rebate coming from PG&E for an energy-efficient data center upgrade. So we're working on programs as well as metrics for energy efficiency at the platform level for mobile desktop as well as data center and server products.

Working with the industry around our amazing products done efficiently. Collaboration is what IDF is all about. That's why we hold the forum, that's why you're here. How can we work together to deliver the next-generation technologies, collaborating with the industry around instruction sets?

As we were thinking about this speech today and pulling together our discussion, one of the things that we hearkened back to was the debates around RISC versus CISC. How many of you remember these debates? And this is a picture of John Hennessey, who was my professor at Stanford. When I was working on the 486, he and I held a public debate on RISC versus CISC. This is a picture of him at that time and a picture of me, two charming young men there, as you see. John's done pretty well for himself as well; the president of Stanford now. And the debate was minimize the instruction set, the thesis of RISC, versus evolve and enhance the instruction set, the thesis of CISC. Which one won?

Well, over the decade since then, we've evolved the architecture substantially. You might have heard the Intel architecture has done pretty well since then. We've extended it with new instructions, we've added new capabilities like floating point to it, gone to virtual memory paging, 32-bits, flat address spaces, 64-bits, SSE, we did chipset integration -- not in the new idea, all the way back to the 386SL -- and more recently, with SSE and 64-bit extensions.

We're working with customers to understand the next evolutions of the instruction set. What are their emerging workloads and how can we enhance the instruction set of the platform to embrace those new usage models and deliver solutions that handle the data sets? High-performance computing, media processing, search/compare, data integrity, pattern recognition -- all of these are new, emerging workloads that the platform and the instruction set needs to comprehend.

As we've described our two-year tick-tock model, we see that we need to be more embracive, aggressive, and change the way that we've work with the industry in describing and setting the direction for instruction set extensions. Today, we're announcing the new generation of instruction set extensions for our 45-nanometer products. So these will be the products that we introduce over the next two years. Today, we're publishing a white paper describing the instruction set extensions that will be part of the 45-nanometer products. This is SSE-4 and more. Things that address vectorizing compilers, media, string and text processing, application-targeted instructions. In all today, we're announcing 50 new instructions that we'll deliver as part of the instruction set over the next two years. Today the direction and the mnemonics, and we'll be disclosing the specific op codes and the semantics for those instructions as soon as they're proven in silicon, a change on how Intel works with and communicates our directions with the industry.

But of course we don't just do this alone. In two of the many ISVs and OSVs that we've worked with are disclosed here, Adobe as well as Microsoft, working with the industry in instruction-set extensions. We've also seen critical enhancements at platform capabilities. We've described the work in virtualization as a key usage model where VT is critical hardware enhancements to support that usage model as well. Novell announced on Monday the first production enterprise Linux taking advantage of the Zen capabilities allowing the hosting of unmodified guest OS's in the platform fully utilizing VT technology.

On this stage, we described our virtualized ASAP program in the spring where we had Diane Green from VMware and myself announcing the virtualized ASAP program for ISVs, offering resources to develop best practices for sizing guides, performance and tuning and application configuration recommendation and testing. The response from that has been stellar, and these are the companies that have joined with us as part of the virtualized ASAP program. The industry has internalized how important this is and the strong support we have from them.

We're also working with the industry in terms of new usage models, new platform models, and one of those partners, SAP, and it's a great pleasure to have Ranjan Das, the senior vice president for Emerging Solutions at SAP join me here today.

[Applause]

Pat Gelsinger: Ranjan, it's a pleasure to have you here with us, the partnership with SAP. We also have Jeff Word here who's going to help us with the demonstration. So what have you got?

Ranjan Das: Thank you, Pat. Jeff and I are here to show this box. That's right, this box. It is a cool new solution to a very complex business problem. The problem is how to get knowledge and business insight into the hands of decision-makers in real time. Historically, companies have had mountains and mountains of data but no business insight. And the bottleneck has been performance. With this box, SAP and Intel collaborating, we have solved that problem. We have many satisfied customers. I was sitting in the back room, reading my email and just received a message that a very large customer has bought this solution so we're very excited about that.

Pat Gelsinger: Excellent.

Ranjan Das: This is called the Business Intelligence Accelerator, BIA, and it is part of a new generation of information appliances.

Pat Gelsinger: So what is an info appliance anyway? So give us a bit more?

Ranjan Das: This obviously is an example of an information appliance. Basically an information appliance is a server product combined with software and it's for a specific purpose. Things to

remember about this: It's easy to install, easy to maintain, has very good performance, amazing scalability and is cost-effective. This BIA, this one here, as I said, it's a first of a new generation of information appliances built in memory technology. The beauty here is that it also plugs in nicely, snaps on to existing data structures, those mountains and mountains of data. Plugs in very nicely and it's optimized and designed for Intel's technology. But let's show you a real-life example. Jeff?

Jeff Word: All right. So what you're seeing here – actually one of our mid-market customers, British Petroleum, it's one of our smaller ones --

Pat Gelsinger: One of our smaller customers. They've probably heard of them, yes.

Jeff Word: Just as Ranjan described, they have terabytes and terabytes of data, but what they want to do is get some business value out of it. So what you can see here is on the top I've got all of my retail store point-of-sale data for all of BP's gas stations all over Europe. They're actually quite a big retailer in that sense.

Pat Gelsinger: And how big is the data set here?

Jeff Word: We're talking roughly around a billion records here.

Pat Gelsinger: Wow, okay.

Jeff Word: So a billion records on the system.

Ranjan Das: And that's not uncommon. In most of the large-enterprise customers, the data set is like this. In fact, what we've seen is that the data set is growing; it's actually doubling in size every 12 to 18 months.

Pat Gelsinger: That sounds like Moore's Law kind of to data sets and working sets.

Ranjan Das: I'm sure you know a thing or two about Moore's Law.

Pat Gelsinger: Just a little bit. So Jeff, what is it that you're showing us here?

Jeff Word: All right, so like I said, on the top of half of this, you can see here's my kind of global perspective on my sales data, I've got five quarter's worth of retail sales data by country, and I've also, on the bottom here, got a little blade monitor for the techies in the audience to see the actual queries hitting the blades. So if I go ahead and click on the United Kingdom right here, you're going to see everything changes. That's the actual query hitting the blades, and it brought it back a query of -- selected 34 million rows in just about a little less than a quarter of a second.

Pat Gelsinger: But that's fast, and you're connected back to Germany for this, but what happens if hundreds of users started to hit this?

Jeff Word: Well, one of the things we found with the Clovertown chips that are on this is you can actually simulate and you can get huge performance gains, so if I simulate really quickly, just clicking four or five of these queries at the same time, simulating five or six users, you can see the performance time is phenomenal on the new Clovertown chips.

Pat Gelsinger: Well, that's great, we're seeing great scalability with you on that, but what's the business value of this?

Jeff Word: Well that's the thing. You know, everybody can look at the big-picture data, but what the real people in the business want to know is they want to drill down into and see trends. They want to see actual sales data by hour, by convenience store. So if you look at this screen, I've got one store called up in London. I can see here's my sweets, because I'm the product manager for the sweets, which is always fun, and you can see here's the individual products. I can see right here the different times of day and the different point-of-sale data for the different times of time per that store. So if I want to quickly click here on the deli chocolate, which is their profitable chocolate.

Pat Gelsinger: Chocolate? BP sells chocolate?

Jeff Word: BP sells a lot of stuff in their retail stores, their convenience stores. You can see here, as a business user, I can get right down into the detailed information that I want, I can look really quickly and see I've got some trend data here, there are

different times of day that I'm selling a lot of this chocolate, so as a business person, I can maximize my profit potential by making sure I have the right product for the right place at the right time.

Pat Gelsinger: Very good.

Ranjan Das: Excellent. Knowledge and business insight.

Jeff Word: Yeah.

Ranjan Das: So, Pat, what's interesting here is that we have seen linear scalability across different generations of Intel's microprocessor. Performance has doubled from, you know, the last three generations, so what this means is that end users can process larger amounts of data in real time to make business decisions.

Pat Gelsinger: So that's beautiful. And the end-user impact of that?

Ranjan Das: I'm glad you asked. Customers are very happy. They are benefiting from 10x improvement in performance, the cost of acquisition is 10x lower, and managing this, like I said, is very easy -- it has gone down by a factor of 20x. Even better, even better, the time to install this has gone down to two hours, so customer satisfaction is very high, going through the roof.

Pat Gelsinger: Excellent, excellent. You know, I've worked with Shai a lot on this, your boss --

Ranjan Das: My boss. His boss too.

Pat Gelsinger: Right, and he was pretty excited about the solution and maybe we can just hear a few words from Shai right now, and his enthusiasm for what we're doing here together.

Shai Agassi[video]: Thank you very much, Pat. I'm delighted to have the opportunity to speak to all of you. My name is Shai Agassi. I'm the president of products and technology at SAP.

You've all used Google at least once in your life, and you know that you can send a request and get every answer that you want to every question in the world. You get an experience that is second to none. We've all seen the same experience today in what we call the Business Intelligence Accelerator at SAP. When you want to go with data to the masses, you have to have an engine that can take you to every single person in your organization. We've done that together with Intel to the degree that we can now put billions of records, terabytes of data in the hands of every single individual in the company. We've also now worked on the ability to extend that into search, the ability to search documents and find the pinpoint, the needle in a haystack, as well to extend that to master data and share and collaborate across one organization or multiple organizations, as well as one application or many applications with shared master data.

We started working together with Intel; we put together an effort around Intel's 64-bit platform, and we created this set of hubs, the data hub for BIA, now expanded into other areas. We ended up with 100x price performance difference, which was 10 times better than what we expected than when we went into this experience. And that's even before we put in Woodcrest and next generations of Intel's chips, and I expect by the end of this year, we'll be at a 500x price performance difference from where we started. So this is tremendous work. You can't do that on your own. This is the result of combining the efforts of two world-class engineering organizations, SAP and Intel. And the end result is obviously tremendous for our customers. We'll see this innovation going into hundreds and thousands of customers over the next year, year and a half. And we think that by doing so we're not just changing the landscape for our customers. We may be changing the landscape for this whole industry.

[Music playing]

Ranjan Das: It's clear that Intel and SAP have done a tremendous work together today. We're very happy about the partnership and the collaboration. I'm here to reiterate SAP's excitement to work with you going forward, working on your multi-core, many-core technologies, as well as virtualization, manageability, and security aspects of your platform. Combining your expertise in microprocessor technology and our experience in enterprise software, we are taking information appliances to the next level. The future looks really bright and I'm very excited about it.

Pat Gelsinger: Thank you. You know, 500X -- that's transformational to business practice. Absolutely great. Thank you, Ranjan.

Ranjan Das: Thank you, Pat.

[Applause]

Pat Gelsinger: Working with the industry and software, we also see the critical need to work on next-generation hardware interfaces as well. Today the industry platform of choice is around Intel's front side bus. You heard Susan Whitney earlier in the talk describe their next generation investments into their chipsets in that platform. Today, I'm excited to announce that Intel is extending our front side bus relationship in the industry to the FPGA vendors, to Altera and to Xilinx, and to enable them to take their FPGA technologies, embrace the front side bus, take advantage of the tens of millions of platforms that have shifted to the industry, and enabling them for application accelerators of tomorrow.

In fact, we've been working in the labs on this for the last several years. What I have here is actually a running prototype of what we call our L3 NIC, where we've actually built one of these FPGAs. This is running at greater than four gigabit per second line rate, so we've actually prototyped this. Now we're taking this and making this available through Altera and Xilinx for the industry to innovate further on the standard front side bus platforms in the industry.

We're continuing to explore I/O evolution for the platform as well. From the early days of ISA and the IBM PC, we enhanced that with ESA in '88 to PCI with plug and play in '92, PCIX as we extended that to 64 bits. In 2001 at IDF we laid out a vision for IO innovations and moved to serial interfaces called PCI Express. Serial interfaces -- one to 32 serial lanes that have come into the platform beginning in 2003 and the PCI Gen 2, which will be coming forward to the marketplace next year. The result of that has been an amazing ramp of PCI and PCI Express and the adoption of those platforms in the industry -- 100 percent of graphics suppliers, storage, networking, and accelerators.

But we have a vision to take this even further. To help me describe that, I'd like to invite Tom Bradicich, the CTO for IBM System X. If he could join me now, we'll talk about the future as well. Tom?

[Applause]

Tom Bradicich: Hello, Pat.

Pat Gelsinger: Hi, Tom. You know, as we were talking, we've been working together forever. You're getting old. [Laughter]

Tom Bradicich: That's true.

Pat Gelsinger: I haven't, but since the 286 we've worked together. I think maybe as two executives we worked together maybe longer than anyone in the industry.

Tom Bradicich: Yeah, the 286 and the 386 days we did simulations. We actually did offloading and accelerated simulation of the brand new 386 processor that you were leading.

Pat Gelsinger: Yeah, that's great. Maybe you could talk to the audience a little bit about some of the work that we're doing together.

Tom Bradicich: Sure. About five years ago, IBM and Intel got together with several industry leaders and developed what is the PCI Express interconnect technology. That has proven to be very successful. It's a low-cost, reliable, high-speed interconnect that has achieved ubiquity across server and client platforms. The momentum continues. We've worked together on PCI Express 2.0, which will be around shortly.

Pat Gelsinger: Very good. This is an area IBM has lots of experience in. What trends and interconnect are you seeing.

Tom Bradicich: There are two dynamics at play that we're seeing here. One is out in the industry we're seeing new computing models putting higher demands on the platform and on the system.

At the same time in IBM, it is the season for us to take the next step in our X architecture roadmap, and that is building proven

mainframe and high-end enterprise capabilities into open and industry standards.

Over on the new model side, we're seeing an increased application of mathematically intensive computations in the financial industry as well as the biophysical and the geophysical industries as well. Also we're seeing intense visualization, not only in just GUIs or user interfaces but also in applications as well as, of course, the gaming industry. And then content acceleration, as well, and processing such as XML. I think I was in kindergarten when IBM pioneered the I/O channel. And back then, IBM applied attached processors, applied accelerators as well as off-load engines to that I/O channel. And we are very, very excited to step up to leading and bringing to bear that experience and that technology to the new open standard that we are calling -- between Intel and IBM internally, we are calling it Geneseo.

Pat Gelsinger: Okay, so Geneseo a set of hardware and software enhancements. And maybe you could describe a little bit to the audience exactly what we're bringing forward or proposing with the industry.

Tom Bradicich: Sure. It's intellectually intriguing as well as will be of great value to our customers. And I can portion the features into roughly four categories. First we are going to address the efficiencies in synchronization and signaling by providing atomic operations that will indeed increase performance and communication

among the elements of the application, the memory, as well as this new I/O interconnect.

Pat Gelsinger: So that's going to significantly improve memory throughput for these application accelerators.

Tom Bradicich: Very important; if you can't get to the data you can't process it.

Pat Gelsinger: Okay.

Tom Bradicich: So it's extremely important. The second dimension we will add is transaction ordering. The data and its associated transactions can be ordered and prioritized, and that will increase dramatically the efficiency of the bandwidth allocated to those particular tasks and threats.

Pat Gelsinger: So huge improvements in terms of the ability to optimize data flow through the interface as well.

Tom Bradicich: Yes, absolutely. Now IBM being a systems company, we're very interested in the systems aspect. So as you mentioned several times, and as Susan mentioned earlier, we're going to work to build thermal and power-management capabilities and control into this interconnect, such that the accelerators in the attached elements may be controlled thermally as well as electrically and power can be portioned and allocated on a dynamic basis and action taken to keep things cool.

Pat Gelsinger: So this is huge, and power-efficiency improvements are incredibly important for both of us.

Tom Bradicich: Yes indeed. And of course a systems approach wouldn't be complete without looking at the software stack. So we must look at that, look at the issues where there are bottlenecks, overhead, and reduce the latency significantly within the software stack. So we're extremely happy to dip into our rich portfolio of technologies and work on the standard with Intel as well as others in the industry.

Pat Gelsinger: Yeah, and if we don't address the software stack, in many cases that is the bottleneck.

Tom Bradicich: Yes, absolutely.

Pat Gelsinger: So we think this enhancement is pivotal to embracing the next generation of application accelerations. It's exciting work that's going on. As you see by the companies listed here, we've had an extremely positive response, and we really look forward to working this with you, Tom.

Tom Bradicich: We do too. We can't wait to keep going and getting it done.

Pat Gelsinger: Well thank you, Tom. And we're going to be taking this forward to the PCI-SIG as well as the DMTF, working to see it standardized across the industry.

Tom Bradicich: Absolutely.

Pat Gelsinger: Thank you, Tom. Great to have you here with us.

Tom Bradicich: Okay, thank you. Bye-bye.

[Applause]

Pat Gelsinger: Amazing project done efficiently with the industry and ultimately for the customer. We've been researching and looking at the issues facing enterprise customers for many years. And we've heard what they've said and talked about these for several IDFs. And we've gotten incredible response from our vPro platform launch. Beginning on September 7th of this year, IT customers have said, "You listened to us and this is the kind of platform that we need to solve those problems."

Some of the usage models that we've been working on in this area, the usage drivers for the platform, have shown incredible innovation. And we've asked our customers what it is, and we've built it and their response has been great -- security, manageability and energy-efficient performance. And we've shown some of the manageability examples of that. And what I'd like to do now is show you another usage model using Symantec. And if I can have Jeff Marek from our demonstration group come and join me here, and we're going to show off a virtual appliance today, right Jeff?

Jeff Marek: Yep. All right, Pat. Are you ready to have some fun?

Pat Gelsinger: Let's absolutely do that.

Jeff Marek: Okay, here's the scenario: We've both just received an email telling us to go online and update our bank account information. And the email even included a helpful link to the bank's website. Sounds a little fishy, but we're going to go with it. Okay?

Pat Gelsinger: Okay.

Jeff Marek: Now to keep things interesting, we're each going to use a different PC. The one I have is brand new HP dc7700 PC with Intel vPro technology. It's running the embedded Symantec Virtual Security Solution. You, on the other hand, have this generic PC that has none of that.

Pat Gelsinger: So I get the beige, boring PC.

Jeff Marek: Yep.

Pat Gelsinger: Okay. This doesn't sound very good.

Jeff Marek: I like it. Okay. So, now, both these machines are running a slightly older version of the Internet browser that has a known exploit, and we're actually going to take advantage of that here in a second. So I've already clicked on the link, and I've gone to the bank website. And I'm going to go ahead and click on the

login button. Okay, immediately what you see here is the Symantec solution has blocked me from going any further. Now, what's happening behind the scenes is this is a network intrusion prevention system. It's doing deep-packet inspection of all the network traffic, looking for activity that's trying to take advantage of the exploit. When it sees it, it blocks it.

Pat Gelsinger: So this is a virtual appliance, so we have the NIC through the AMT engine, through a virtual partition that's doing this deep-packet inspection.

Jeff Marek: Yes. This is a tamper-resistant appliance. It's resistant to attacks that are directed at the security software itself, as well as things that a well-meaning user might do, like try and turn it off.

Pat Gelsinger: Okay, so is it my turn now?

Jeff Marek: It's your turn.

Pat Gelsinger: Okay, so I get the big beige board.

Jeff Marek: Yep, go for it.

Pat Gelsinger: All right, so I click on here.

Jeff Marek: Okay, now what's happened on your machine is, a brand new piece of malware is taking advantage of the exploit on your

machine to install some software and run code. And it's chosen to run a key-logger application. Go ahead and click there.

Pat Gelsinger: Don't send. I never do that.

Jeff Marek: Don't want to send, okay. Now, it's redirected you to the real website, and up here it's now logging all of your keystrokes.

Pat Gelsinger: Okay, so I'll restart.

Jeff Marek: So why don't you go ahead and log in? Go ahead and log in. You weren't supposed to use your real password.

Pat Gelsinger: Oh.

Jeff Marek: All right.

Pat Gelsinger: Okay, it looks pretty good.

Jeff Marek: You know, I'm not entirely sure you're going to have that money for very long.

Pat Gelsinger: $49 billion, Bill Gates look out. That's great.

Jeff Marek: So that's --

Pat Gelsinger: You don't think I'm going to have that money much longer?

Jeff Marek: No. I think that I have a stop after the show.

Pat Gelsinger: Okay. Thank you very much.

Jeff Marek: Alright, thanks a lot, Pat.

[Applause]

Pat Gelsinger: Since the launch of vPro, we've had incredible embrace by the industry -- system integrators, ISV, OEMs -- all of those seeing the value proposition that we're creating. One of those partners, in particular, I'm most excited to have with us here today. Please join me in welcoming Todd Bradley, the executive vice-president from HP's personal system group. Todd?

[Applause]

Todd Bradley: Hey, Pat, how are you?

Pat Gelsinger: It's great to have you with us today, Todd.

Todd Bradley: Thanks, Pat, it's a great opportunity to come join you on stage.

Pat Gelsinger: You know, we see vPro as having a great impact on end users, software developers, service providers, and, you know, you're sort of all of those. What's been your initial response?

Todd Bradley: Well we are. I certainly agree with you. And as you can see from some of the quotes that are coming up behind us, not only are we great endorsers of vPro, but really using vPro to meet our customers' business challenges, to help them drive down IT cost, and make sure that those IT challenges are closely aligned. Like everything we do for our customers, we're focused on how do we help them drive productivity, how do we increase security and manageability, and at the same time, do that while driving down the total overall cost.

Pat Gelsinger: Now you've just introduced a new business PC with Intel vPro technology, right?

Todd Bradley: We have. We just introduced the DC 7700 with vPro technology. And what this really does is it enables our customers to discover and fix their systems remotely. It allows them to spend more time on running their businesses and not necessarily running their technology. But really using that technology, you know, as the productivity tool that it's meant to be. You know, we were very early to market with this, with you guys, and the customer feedback has just been extremely positive, and we're very excited to push this even further into the market.

Pat Gelsinger: Yeah, and we've seen great embrace from the ISVs, and, you know, Symantec that we just had demonstrated here, is one example of that with their security appliance, but HP is also an ISV.

Todd Bradley: Well, we are. As we've talked about, you know, this is providing not only new opportunities for all your developers, but I think at the same time, as we look at embedding our OpenView software suite, we're able to provide more cost-effective deployments. We're able to provide, you know, products that are easier to manage with HP and vPro together. So by keeping this connected to the PC, even when it's turned off, what we're allowing IP managers to do is use OpenView to deliver a new level of configuration management, one that provides faster remote diagnostics, easier repair capability, and our customers get even more functionality and continue to drive that cost of ownership lower.

Pat Gelsinger: vPro has been a powerful enabler for those IT service providers. And how is it affecting your services business?

Todd Bradley: Well, first, it fits very clearly and very well into our multi-layered services solutions. They're focused on both physical and data security through services. And our systems take proactive measures to really build in defensive capabilities and allow agent monitoring. So vPro, really, through features of remote remediation that we provide, up-to-date inventories, being able to schedule those after-hours wake-up types of capabilities within the PC, all provide a much better experience to our customers, and you're a big enabler of that.

Pat Gelsinger: So HP is a great solution for hardware, software, and services.

Todd Bradley: Well, you know, as we look at how we deliver that total best in class solution, how do we help tackle some of these very, very difficult IT problems that our customers face, we know that working with products like vPro, embedding products like vPro in our solutions, enable us to provide smart client devices, take advantage of the service infrastructure that we've created, and really bring that services expertise to bear. So I can't tell you how excited we are to bring vPro to market. You and I talked about it this morning. About 18 months ago, we were just starting to talk about all the great products we could build on this platform, and it's great to be here today driving that, because it really does push our products into the sweet spot of this market, to create compelling, connected products.

Pat Gelsinger: Todd, it's great to have you here. Thank you for your partnership on this product.

Todd Bradley: Always good. Thank you, Pat.

[Applause]

Pat Gelsinger: This is vPro today. And we see a future, as well, as we rapidly evolve this platform to respond to user inputs and additional challenges. In the first half, Intel AMT will become part of the Centrino platform as well and we'll deliver the vPro value proposition in our mobile platform offerings as well. IT buyers, both desktop and mobile platforms. We're also enhancing the platform for the enhancements to the Intel AMT engine as well

as the Intel Trusted Execution technology, securing the boot and BIOS of the platform as well. It is also being enhanced to support next-generation standards such as web services, providing the ability from a single console to manage the entirety of the enterprise resources.

One of the challenges that we've seen for an IT customer today is that he has increasing challenges in areas such as data security and how to protect and manage that, challenges being foisted on him by Sarbanes-Oxley and other things. Today, he really has two computing model to look at. One is a thin client model, which promises higher security and better manageability. The other is the thick client compute model, which provides a richer user experience, mobile and remote connectivity. Both have their place in different business needs and we sell silicon to both, so we're happy for both. But ultimately, that's far too rigid of a picture.

We see another emerging model in the industry. This is called "software as a service." What we see as the future for vPro is enabling it to be the best SAS client possible, balancing the computing overhead between the data center and the client, delivering a rich experience of voice, video, collaboration, rich user interface, and maintaining, though, all the values of the thin client model by securing the data in the data center. It drives a breakthrough in data center security as well as TCO.

A SAS architecture, in the simplistic view, is something like this: You take a naked client PC attaching and authenticating itself to a service. After it's authenticated itself, then based on the credentials it presents to that service, the policy manager then evaluates and licenses the application and streams down the data and application. Not a complete desktop install; the data and application required for that particular user and service. And then the code is executing on the client, taking advantage of the client performance and richness, and at the conclusion of the application, on exit, the data is delivered back to the server and data centers, where it is secured and protected based on the data center policies.

And we see that this SAS model is emerging and getting a huge amount of interest from the industry and a large number of players in the business already are embracing this and we're seeing it embraced by industry players like those listed here, as well as by customers who are looking at this.

What we'd like to do now is take a quick look at one of the early providers of the SAS solution. Please join me in welcoming David Wadhwani, the vice president of product development for Flex products at Adobe Systems. David?

David Wadhwani: Hi, Pat. It's a pleasure to meet you.

[Applause]

David Wadhwani: As you know, Adobe's been enabling designers and developers to build service-oriented clients for a long time with Flash and the related technologies. We're not just talking about SAS; we've been doing it since the late 1990's. A couple of years ago, however, we significantly increased our investment in this space with the introduction of the Flex product line. The product line is made up of three related technologies. The first is the Flex SDK which provides a command-line compiler, debugger and a rich user-interface technology. The Flex SDK is available completely for free for both commercial and non-commercial use. We also have Flex Builder for those who prefer more of an integrated development environment. And Flex Builder is based on Eclipse and provides a visual design debugging and coding environment. Lastly we have Flex Data Services for those with more of a complex data management need, and it greatly simplifies the effort required to manage data across the client and server tiers.

Now while Flex has been shipping for about two and a half years to-date, we have nearly 2,000 enterprises and ISVs worldwide that have already adopted it and great partners like SAP and NEC. In fact, you saw SAP earlier today demonstrating their new technology with a Flex front-end.

Pat Gelsinger: So can you show it to us, Dave?

David Wadhwani: Absolutely. So here we have a brand-new PC, never been used before. The only thing on this PC is a browser and the Flash

player. I'm going to go ahead and launch a sales force application on the PC.

Pat Gelsinger: Okay.

David Wadhwani: I'm going to go ahead and log in as Joe.

Pat Gelsinger: Okay, and that's it?

David Wadhwani: That's it.

Pat Gelsinger: So if you took a naked PC and now we see the application running?

David Wadhwani: Exactly. The application was streamed in real time which obviously lowers the cost of provisioning licensed management and updates. The other thing to notice about this is pay attention to the rich user interface. As I click around here, you start to see a great visual dashboard, you see fluid vector graphics, and you see integrated audio/video conferencing. In fact, Fred, one of my regional sales managers is logged in right now as well.

Pat Gelsinger: Wow. That's impressive. Now it looks like you've got a problem here in this month. Something's screwed there.

Fred: It's fine. I'll fix that right now.

David Wadhwani: Thanks, Fred. So you saw that Fred fixed the data on his machine. And Flex Data Services managed all the data synchronization across the client and the multiple servers so that immediately, in real time, it was updated on my machine. So now that we have the real data, let's go ahead and recalculate our sales forecast using a Monte Carlo simulation.

Pat Gelsinger: So Monte Carlo simulation, this is a pretty heavy workload, right?

David Wadhwani: Right.

Pat Gelsinger: On a large data set, and these calculations are being done on the client side.

David Wadhwani: Absolutely, exactly. So pay attention to -- just talk a little bit about these screens. On the end you have the client machines running the Flex application; in the middle you have a server. Pay attention to the server CPU utilization. It's nearly sitting idle. On the other hand, if you look at the client boxes, each of those cores is running about one million calculations of Monte Carlo. We're obviously taking full advantage of the Core 2 Duo processor.

Pat Gelsinger: Wow, that's very impressive. Now when you exit, when you finish with this, where is the application data rend?

David Wadhwani: The application data in this implementation is all residing on the server for, obviously, to enable more corporate data security. We're leveraging Flex Data Services to provide the data security and, obviously, the Intel vPro to manage and create a more secure client.

One more thing I wanted to add quickly was that Adobe's continuing to innovate in this space. And I wanted to give everyone here a sneak peek of Apollo. Apollo's a code name for our cross-browser OS. And you can imagine a single Flex application running both online and offline and running in the browser or directly on the desktop, as pictured here.

So as we start to continue to push the envelope, we need and will continue to need more and more [client-side] processing. So we're thrilled to be working with Intel to change the future.

Pat Gelsinger: This is fabulous. Thank you so much, David.

David Wadhwani: Great, thanks a lot. Take care.

[Applause]

Pat Gelsinger: We see SAS as an emerging and important trend, and we see that through collaboration with ISVs, enhancing the vPro platform, extended trusted execution technology into the platform, and new technologies like Robson, which allow cache and [impending] of applications on the client side, key

innovations to make SAS a mainstream volume deployment for client computing in the future, solving the dilemma between thin and thick computing. In fact, this week we see that there's emerging innovation and interest in this topic. SAS Con, hosted just across the street -- and Intel was one of the sponsors for this and a speaker at the conference -- the first SAS conference in the industry. And we see that this is a great opportunity for the industry to innovate and key usage models in hardware and software for tomorrow.

It's been an amazing six months since we saw you last. Right? As we started our talk today, a wall of wafers that simply is overwhelming and stunning in its breadth and depth.

Pat Gelsinger: Thank you. We've seen that we've done these efficiently -- efficiency in everything that we build, not just great performance. The customers -- BMW, BP, Bindview, CCG, NHN -- that we've talked about today, all embracing these new platforms. But IDF is all about the industry, it's all about how we partner together at delivering the tremendous potential that Moore's Law offers us today, as well as in the future. But the pace of Moore's Law is unrelenting. Will just handed me the first silicon out of the fab on some of our next-generation products. I'll be talking to you about this one in the spring. Time to get back to work. Thank you very much.

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

Intel’s results could be affected by adverse effects associated with product defects and errata (deviations from published specifications), and by litigation or regulatory matters involving intellectual property, stockholder, consumer, antitrust and other issues, such as the litigation and regulatory matters described in Intel’s SEC reports. Please refer to Intel's most recent Earnings Release and most recent Form 10-K or 10-Q filing for more information on the risk factors that could cause actual results to differ.