Intel Developers' Forum
August 24, 2005 - 8.30 to 10.15am

Female Voice: Please welcome IDF General Manager Rob Chapman.

[Applause.]

Rob Chapman: Good morning and welcome back to day two of IDF. I hope yesterday was a productive and fun-filled day for you, and I think you're going to find a lot more excitement today. I have two housekeeping items before we get started. First of all, some of you may have received information that indicated that tomorrow morning's keynote was going to start at 9:00 a.m. That is incorrect. Tomorrow morning, like this morning, we start at 8:30 a.m. Second, huge demand for our multi-core architecture and implementation class. We're going to do a repeat this afternoon. That's at 4:30 in room 2002. Again, multi-core, 4:30, 2002.

Before we get started today, I'd like to hit some of the key take [aways] that we heard yesterday. First of all, Paul introduced to us the next generation power optimized microarchitecture. We're taking a usage-oriented approach to that next generation that's going to enable new form factors and new business models. He talked about a 10X performance scale in certain product areas where it's needed, and he also talked about reducing power 10X in product areas where that is needed. This will enable a performance-per-watt advantage that will be based on multi-core platforms that will serve our server, desktop, laptop, and a new area, [handtop] product categories.

All of this will be supported by our new technologies, our star Ts, in areas such as security, mobility, and virtualization. This is a very compelling value proposition for our software developers, and you're

going to hear a lot more about that today as Pat Gelsinger talks about the digital enterprise and as Don MacDonald talks about the digital home.

We also had an update from Sean and Dottie on mobility ubiquity. They talked about Napa as the next platform of choice in the notebook laptop area, and that there was a lot more to come as we roll out the next generation microarchitecture in the notebook space in the second half of '06. They also gave us a great update on the proliferation of Wi-Fi and the mass acceptance of that, and I think they also showed some pretty compelling early demonstrations of WiMAX and the fact that is ready to ramp, that it's starting to be accepted world wide, and that we do need to take a standards-based approach to that to assure mass proliferation.

Finally, in the area of digital health, I think we heard a great story from Mr. Burns that appealed to us both at a professional level but also at a personal level.

I want you to know that we're listening to your feedback. Your feedback is very important to us, and I’m hoping you're seeing some of those changes on site this week. We've increased the access to our technical experts. We have the think tank event tonight; I invite you to join us for that. We've implemented chalk talks, which we're piloting this week, which seem to have met with rave reviews yesterday. That will continue today. We also had the birds of a feather session; I didn't see how it went this morning, but we'll be doing that again for breakfast tomorrow morning, so please join us for that.

We continue to increase industry participation. We have over 70 companies that are participating with us, teaching classes and labs, and we'll continue to increase that mix going forward. We also have a new concept called these community zones. You may have noticed that in the spring. I invite you to visit the [ISA] Village down on the second floor. That's kind of a new concept as we continue to involve our communities and make them more compelling and interesting to our attendees. We're also doing more as far as hands on with these experience zones – the virtualization zone and the power zone, both on the second floor. I'm seeing a lot of people in there. I hope you're taking advantage of that learning interaction opportunity.

Again, content is key. We're continuing to increase the mix of our advanced and intermediate classes, so again, we really look to get your feedback so we continue to improve and evolve this program. I want to be able to exceed your expectations as we move forward with IDF globally. Finally, I want you to enjoy the second day. Make it a productive one. And I also want to introduce really a man who needs no introduction, Pat Gelsinger, the grandfather and godfather of IDF, and the head of our digital enterprise group, please welcome Pat, and thank you.

[Applause]

[Music/Video plays]

[Applause]

Gelsinger: Like you, Intel has a vision to transform the future. It takes a team of talented individuals with passion and vision to succeed and, together, create the technologies that will transform business and change the world.

[Video plays]

Pat Gelsinger: Welcome to day two of IDF. It's great to be back at IDF. I will say, though, it feels a little bit odd, right? You know, this is my baby, and I sort of passed on the baton and moved to a new role. So, I tell you, I felt a little bit funny sitting in the audience on day one just watching the show go on. But, I do want to say thank you to the IDF team – to Rob, to Jeff, to Eric, and everybody's who's putting on just an absolutely – and maybe our best ever IDF. Maybe we can just have a round of applause of thank you to all the people who worked very, very hard to make this happen.

[Applause]

Pat Gelsinger: But, the grandfather of IDF? [Laughter] Golly, I'm not that old, right? I just love IDF; I love the technology, I love the people, the passion of the show, the energy that we have here. Just seeing the industry's innovations just everywhere you go. I just love walking the show floor. I think this is my 16th time on stage here. So many of you I count now as personal friends as well as partners in the industry.

Recall last spring, at IDF, where I introduced the Digital Enterprise Group. We talked about our mission where we were off to develop an end-to-end view of the enterprise, exploring what we're doing in clients, what we're doing in servers, communications, and finally in the storage area. What you saw Paul talk about yesterday was describing this end-to-end strategy where we're leveraging technologies across each of these four business areas to deliver a value proposition that's consistent for developers; that delivers value to customers, and ultimately allows our industry to grow and deliver more to our customers. From big machines like Itanium through communications like ATCA through next generation server platforms and finally, innovations in our business client, we have it all today and we're going to describe those areas to you before we're finished in our time together.

At last spring IDF, I described work that had been underway for several years, exploring what it is that business customers were looking for from us and us together. What are they interested in to change and improve their business productivity and cost of ownership and operation. We presented this model at Spring IDF of the four "pillars" as we called it. In those four pillars, pervasive connectivity, hooking it all up, and you heard much about that yesterday from Sean Maloney and WiMAX and Wi-Fi and new mobile platforms. Today, I'll cover the other three pillars in the course of our time together.

We'll talk about embedded IT and how there's half a trillion dollars of manageability support costs that we think we can change for how we deliver technologies to address. How seamless collaboration as teams

become more dispersed, what can we do on our platforms and technologies to improve the work experience as they cross time zones and geographies. And information assistance, the flood of data that is emerging in the corporate environment. What tools, what capabilities can we deliver to those users to be able to manage? At Intel storage and storage platforms, our number one area of IT spending growth, and can we do something to make that data more useful for our customers?

Let's start by exploring embedded IT in a bit more detail. As Paul Otellini mentioned yesterday of the IT budget, 90% is being spent on maintenance; and of that, 50% associated with manageability costs. Security risks are rising; we see that there are enormous challenges facing the IT manager of today. Our fundamental strategy is that we want to embed IT -- just a little bit of the IT manager directly into each piece of silicon that we build and ship into the marketplace, and thus reduce the risk and increase the efficiency of IT. In the spring, we introduced active manageability technology, or AMT, a hardware-based approach to bringing new embedded IT capabilities in the platform, and started to, at the fundamental level, addressing these spiraling costs, and difficulties of the IT manager.

Here's a picture of Stacy Smith, the CIO of Intel. I like Stacy a lot, and ultimately what we want to do is take Stacy and stuff him into silicon. But in doing that, and stuffing a little bit of the IT manager into silicon, fundamentally creating new capabilities that allow us to remotely provision IT without the IT manager ever having to come

and physically touching the PC. Asset tracking, and delivering new capabilities, and this is what AMT is all about.

We are just delighted, since the spring, of the amount of enthusiasm we have gotten from the industry around using AMT technology. What you see here is everybody who does manageability technology has signed up and said, "This is cool. We can take advantage of AMT and improve our products and our offering to our customers." And you see many of these who have now committed and now support for AMT. In fact, if you go down to the show floor you'll see many of these demonstrated and many of these products that they're committing delivery to the industry, many of them shipping, in fact, this year. So just an incredible amount of software support for it.

The second thing I'm very happy to show you today is the incredible hardware support. These are the OEMs and we have our wall of AMT platforms shown here, of all of the OEMs who are now delivering AMT-enabled hardware to the industry. Literally worldwide support for delivering this new T, advanced manageability technology, to the industry. And many of you are represented here together. We're enabling and delivering the first generation of embedded IT technology to the industry.

The most exciting thing that's happened since spring IDF is we went and approached the global IT outsourcers, system integrators and the like. People who traditionally haven't found the bleeding edge of our technology to be in their best business interests. They have jumped all over AMT. In fact the four names that are listed here collectively

represent five million PCs that they manage and support for their customers as IT outsourcers. They have embraced AMT in an incredibly strong way. Highly motivated to differentiate their offerings to their customers, reduce the cost of their offering, and improve the service level to their customers. So all four of these have signed up as using and driving their customers, OEMs, and suppliers to AMT as well -- very, very exciting developments in spring IDF.

Finally, I'm very happy today to announce the most recent member to the AMT family, and that's Cisco systems. If you were here yesterday, Charlie Giancarlo of Cisco gave you a preview of this. But, you know, how Intel as a client provider is working with Cisco as the infrastructure provider to make the two work together and deliver a new level of capability. And to not just talk about that, I'm very happy to demonstrate that today, and joining me in doing that is Jayshree Ullal from Cisco Systems, Senior Vice President. Thank you, Jayshree.

Jayshree Ullal: Thank you, Pat. It's a pleasure to be here.

Pat Gelsinger: Well it's a pleasure to have you here, Jayshree. And you have a big job at Cisco, and I just appreciate that you'd take time out of your job and come and talk to us a little bit about how we're working together.

Jayshree Ullal: Well you know, talking of expanded responsibility – congratulations to you as well.

Pat Gelsinger: Oh thank you.

Jayshree Ullal: It's better than having an expanded waist. Um, it's pretty exciting here. I would say that, um, security is no more just an issue for the industry. I think it's truly an issue in terms of technology. And many of the things we're doing is very reactive, but what's exciting to me is, I think, two key technology vendors from a [networking] perspective and a system perspective are really now building OS-independent tools, that can proactively solve the problem rather than react to it after the fact.

Pat Gelsinger: But it really is something neither one of us could do by ourselves.

Jayshree Ullal: No, I need you and you need me. [That's a big plus, yes].

Pat Gelsinger: Absolutely.

Jayshree Ullal: And you know what? That picture you were putting into the box there – Stacey?

Pat Gelsinger: Yeah.

Jayshree Ullal: Well I'm hoping at some point we can trap security that way.

Pat Gelsinger: Okay, just stuff it inside there as well. So, Jayshree, can you give us a look at how this actually works?

Jayshree Ullal: Yeah, let's take a look at that. So what we are talking about here, I think, is two key technologies. NAC, Network Admission Control –

this is really an industry initiative. And we have about 60 partners, but Intel, I think, is a very strategic and unique partner.

Pat Gelsinger: [Lots of capabilities], and that capability is something that's part of Cisco's products today.

Jayshree Ullal: Yeah, NAC ships in our switches and routers and all this networking stuff, you know? And what I'm going to show you is I'm going to plug this switch port over here. There you go, I'm still capable of that – linking into that RJ45. I want you to see here on the screen, it says, "Quarantined firewall [posture] non-compliant."

Pat Gelsinger: Okay, now, so we've plugged in our IBM, our [ThinkCentre] from [Lenovo], into your router, and all of a sudden a bunch of IT gobble-de-goop comes up, so what's going on?

Jayshree Ullal: Well, it's really meaningful gobble-de-goop, let me tell, let me talk more about that. So what you see here is three major components -- the hardware platform, the IBM platform, and then the critical AMT technology, and then the linkage into the Network Admission Control. And key to making that happen is what we call a "trust agent," a Cisco trust advisor agent, which has a variety of posture plug-ins.

What's happened before is you have a network and you can have anybody wild and wooly enter in at any time. You have PCs, laptops, wireless access points, video surveillance gear. By virtue of this technology, what you now have is plug-ins that can go in -- an AMT plug-in and a firewall plug-in -- that gives you a method to how you

access the network. So let's see how that really happens in action here for a minute.

Pat Gelsinger: Okay, so, if we can dive in and take a deeper look at that, that would be great.

Jayshree Ullal: So what you have here is the three major components: the platform; the critical hardware, the AMT technology; and then the software that runs on the network that provides this intelligent access, right? So the platform interacts with this hardware, with these plug-ins, and it has two major components: the firewall plug-in that establishes the policy, and then the manageability plug-in, which is the AMT plug-in. What's unique is –

Pat Gelsinger: It's very important, right, that the AMT plug-in, that this is acting out of band, right?

Jayshree Ullal: There you go.

Pat Gelsinger: [That the OS] could be up or not, right, the platform could be plugged in or not, right? As soon as it's plugged in, out of band and it's interacting with your capabilities --.

Jayshree Ullal: Right, what's really unique there is you get the NetBIOS, and you get the configuration settings, and those are things you don't get until you go into an OS-independent system. So you can see here the trust agent is collecting the information from the posture plug-ins, if you will, and

now it's just basically establishing the process so you can have all the right security parameters.

And then from there, let's see, you now go into really understanding – you have basically a NAC-enabled firewall policy, that allows you to say, "Okay, how does this interact with the corporate network?" And you're now talking into the corporate network, you go into, the network, you actually look for the firewall posture, establish the policy, and then you make the connection, if you will, between the end system and the network. And look what happened here – it didn't allow you in.

Pat Gelsinger: Okay. So before, I could do anything nasty to the network. Voila!

Jayshree Ullal: Yeah, so now what you have is a method. You can be permitted, you can be denied, or in this case, you've been quarantined, because now you're getting a chance to remediate yourself. You can fix it. Before, you were just allowed without any method, right? So now you have a mechanism to fix yourself, and that's what we're going to go through now.

The firewall policy is notified in two places. Your IT administrator is notified and then you can also be notified of what to do with the device system in terms of server updates and patch management. That's what you're seeing here. You're sending that firewall update now back into the network, and also into the [AMT] technology, so you can fix it.

Pat Gelsinger: Okay, without IT ever needing to be involved?

Jayshree Ullal: Ever intervening. This becomes a much more automated process between the network and the end system, and a lot less manual and a lot less reactive. You can see here now we're going to transform from this denied access state, which we're in, to now the admission control server and the network granting you access. All of this happened real time. All of this happens with the magic of the [NAC and CTA and AMT] technology.

Pat Gelsinger: Okay, so Intel and Cisco coming together, a fundamental new capability for the industry.

Jayshree Ullal: I think it's very exciting. As you know, the hardest part of all this, Pat, is the planning and the innovation and the work we did together eighteen months ago, right? And it's very exciting to see this in fruition here today.

Pat Gelsinger: That's fantastic. One more question before I let you get away here, Jayshree. Now this is cool, and we're all geeks here so we enjoy this and how it works, but when are we actually going to deliver this to customers?

Jayshree Ullal: I was going to ask you that question.

Pat Gelsinger: Well, I think we're working together, and it's this year, right?

Jayshree Ullal: I think so. It'd December 32nd. [Laughter]

Pat Gelsinger: December –

Jayshree Ullal: No, no, no, I'm sorry!

[Laughter.]

Jayshree Ullal: It's going to be end of this year. Are you committed?

Pat Gelsinger: Absolutely.

Jayshree Ullal: We are too. Thank you. It's pleasure to be here.

Pat Gelsinger: Thank you, Jayshree.

Jayshree Ullal: Thank you. Bye-bye.

[Applause.]

Pat Gelsinger: We're very excited with the partnership with Cisco we're announcing today. Yesterday you heard Paul and Steve Ward from IBM Lenovo describe how they were, in their platforms, delivering and utilizing the AMT technology as well. They were going a step further in their exploration of this, actually combining VT and AMT. Since we're a little bit geeky now, we can probe a little bit deeper on what you saw yesterday as well as look a bit forward into some of the next steps of the embedded IT usage model.

What we did, and what you saw yesterday, was AMT, the [out-of-band management engine], but it was also acting in conjunction with VT technology. Fundamentally, what we did is we actually virtualized the NIC connection into our EIT virtual appliance, and then the NIC is represented as a virtual NIC into the actual running operating system and platform, so a little driver, you know, [shim] driver that actually makes it appear like you have a NIC, but the real traffic is going through this out-of-band manageability engine that's now virtualized using the VT capability, so this allows client isolation, allows [healing, and] as you saw yesterday, with Lenovo, their antidote engine to run on top of this.

The next step in this is EIT is going to evolve even further. You see in this picture that we've taken two more steps in our evolution of our EIT – embedded IT capabilities – we've also expanded the virtual machine capability so we virtualized a bit more of the platform. You can imagine that we virtualized storage or other elements of the platform as well, so we have virtualization capabilities which might enable [additional] usage models as well. We've also added LT and TPM, so we've hardened or secured that out-of-band EIT virtual appliance, making it even further secure.

You can imagine as part of this that there are a number of business opportunities and applications that emerge from the introduction of this EIT virtual appliance. You can see that there are other third-party applications. You can see maybe automated self-recovery becomes a service offering. Maybe backup, fail over, and restore capabilities. In the future, we envision that there would actually be uses of this not just

in big corporate IT, but also in SMB down the wire services, or even service and management offered delivered to home users as well. We see this as a rich area. In fact fundamentally reinvigorating and adding a whole new venue to the business model for PCs today. We're quite excited for the directions of embedded IT.

In addition to management, one of the things we've heard very strongly from IT managers is stability. We want to be able to qualify something, deploy it, and not have it change. In the interest of that, we developed a program, launched it three years ago, we launched the stable image platform program. You think about it as 1-4-1. We have a quarter of qualification, four quarters of stable platform available in the industry, and one quarter of end of life as we transition to the next generation. So the yearly cadence to the platforms that we're delivering and that are stable for corporations to deliver and deploy. This program has been extraordinarily successful for us. Over 100 million units have been shipped under this program to date, and over 50% of all of our business platforms today utilize this, together delivering platform-level stability to IT managers around the globe.

As we looked at this program, we saw that there were some corporate IT users who simply bought transactional. “Give me this exact PC, this day of the week. I don’t care about stability.” And, there was the other, larger portion of the customers who wanted the stable capabilities, and that was our fundamental platform offering to them.

We also found, as we began engaging more with IT users, there’s also a class of businesses that wanted the best platform as well as stability,

not wanting to give up those best, newest technology capabilities. With that in mind, we announced earlier this year the Professional Platform where it includes stability but also the best technology that we have to offer. This includes AMT, VT, Mainstream and Stable. This was the desktop; we launched the Lydon platform this year. We'll have the equivalent Professional Platform launch for the mobile platform with the Napa launch in 2006 – the notebook cadence following the desktop cadence by a few months.

We found, again, very enthusiastic support from the industry for this Professional Platform launch. I list here all of the OEMs who joined us in the launch of that Professional Platform; an incredible amount of industry support for that, noting specifically those who are including AMT as part of those platform offerings. We continue to look at that as an opportunity as we move to a yearly platform cadence of capability.

This is a glimpse of what you can expect in the 2006 Professional Platform offerings that we're bringing to business, as we bring dual-core into the mainstream of the business platform. We'll also have a second generation of the AMT technology where we'll add additional features like circuit breaker, full integration into the chipset, and new presence and agent technologies as part of it. Of course, it will be stable. We call this the Averill Platform for 2006, for client, for desktops, and we'll have the equivalent of early 2007 for the notebook platforms. You'll see this generational, yearly cadence that we'll bring to the business Professional Platform.

So, this was a look at what the embedded IT capabilities are for the client. Let's take a look at the server usage models for embedded IT, as well. In particular, I'd like to explore virtualization technologies as they are changing the way that servers are delivered and deployed. Virtualization technology allows you to run multiple OS instances on a piece of hardware, and literally, through the virtual machine manager where you can have an OS and a different task running in a particular virtual machine environment on a server.

The interesting thing about this is by doing this, you actually create a new level of flexibility as well as robustness on the part of the IT manager. As shown on the slide here, you could be running a database server. You could have an active second server running. If there were any reason that needed to be migrated or upgraded, you simply could, in fact, quickly migrate that service to a fail-over service that's running for that same web service. It could be running on the same hardware, or in fact, it could be running on a different set of hardware, as well.

At Spring IDF, I was very happy to be joined onstage by Jim Allchin. We talked about a number of technologies together, and in particular, Jim emphasized their support and use for VT technology and their plan to make that part of the Longhorn product offerings. But, better yet, I'm very happy today to announce, for the first time ever, that we will have VT support as part of the Microsoft Virtual Server 2005 Enterprise Edition software. Better yet, we'll give you the first demonstration of the capability today.

What we have is two VT-enabled development platforms. I'll tell you what those are a little bit later in the speech. On them, we're running Microsoft Virtual Server. These are both VT-enabled – the first time that we've ever shown Virtual Server, now with VT capabilities. We have a set of virtual machines running on this platform and you can see we have NT legacy, SQL, and the web server. What I'm going to do is migrate one of those virtual machines to a second piece of hardware. In this case, it could have been another place on this virtual machine, or it could have been another piece of hardware somewhere in the cluster.

So, I picked this particular virtual machine. I right-clicked on it, and told it to move that group. Had their been more than two machines in the cluster, I would have had to identify which machine I wanted that sent to, but it's now packaging it up, so it's encapsulated that entire OS application environment, put it in cellophane, shipped it across the network, and as you see it's now popped up and running on our second virtual server. First ever demonstration of Microsoft Virtual Server with VT capabilities running on our VT-enabled server platforms, which I'll tell you a little bit more about those later, delivering robust, improved, and simplified virtualization to the enterprise.

This is one example of how Virtualization Technology can be used. Let me give you a second. We see that in datacenters there's a large amount of legacy applications, maybe running on a particular OS configuration, with a particular application environment. The guy who wrote that particular application is probably not with us anymore. So it's very, very hard to imagine migrating that environment forward.

The result of that is that you've got a bunch of hardware there that are underutilized, and drive up the manageability costs of the environment. So what's happening is you can do application consolidation using virtualization technology, where you'll have legacy OSs running these applications, but able to bring them into a modern platform environment.

Today I'm also very happy to announce and demonstrate to you the first time that VMware technology. VMware is a leader in virtualization technology in the industry. We announced their plan to be time to market with us on VT for the workstation product. Today, I'm very happy to announce their commitment for the ESX product line, the server product line with VT technology. Let's give you a demonstration of their product running on a VT-enabled platform again.

Here we have a VMware prototype, VT enabled. You see we have three operating system environments, now we're running at 64-bit Windows Server, 2003 Enterprise, so a full 64-bit OS virtualization. Let's click over here. Go to an Oracle environment. So this is running on Red Hat Linux, so running on the exact same piece of VT-enabled hardware with VMware technology. We'll go over here to a legacy Windows NT environment as well. So three very different operating system environments, all seamlessly running on the same platform.

Let's go do something nasty here. Let's crash one of these machines. Not that that would ever happen in real life, but we have a special application that does that. We'll go blue screen, but while that's blue

screen we can happily bounce back and forth to the other operating environments running on this platform seamlessly. With VT technology our industry partners delivering robust virtualization to server platforms in the future.

VT makes virtualization more reliable, more stable, simpler, capable, and responsive. We're launching VT technology in our products in the second half of this year. We'll begin shipping those with our desktop platforms that will be followed shortly by our Itanium platform support with Montecito platforms. We'll also have it in our next-generation mobile platforms in the beginning of 2006. We'll also have our next-generation DP and MP platforms as well. Literally we're delivering virtualization technology top-to-bottom across the Intel product line. VT becomes a core capability of the Intel platforms going forward.

Virtualization today: software only. With VT we make it simpler, and more robust virtual machines. We're accelerating the performance capabilities of that with next-generation developments, and then we'll expand the virtualization domains of the platform to IO, graphics, and other capabilities in the platform as well. What you'll see from Intel is a rich and robust virtualization environment for the future.

So we've looked at the embedded IT leg. Let me now go on to leg two: seamless collaboration. More and more, the challenge of enterprise is working across time and distance. My organization at Intel has almost 50 sites. How do I make my teams effective literally as they work

around the globe, 7x24? It's about how we bridge time and distance for our users.

VoIP is the first area that we're particularly focused on with regard to seamless collaboration. It's a key industry focus that we believe will hit mainstream in 2006. We see this as sort of a beachhead to many other service over IP offerings as we look to the future. We see our technologies and the PC as uniquely well positioned to benefit and deliver these capabilities. We see it delivering reduced cost to the enterprise, as you have international bridging, multi-party calls. It enhances usage models, making our workforces more productive as well. It works over wired and wireless, and also once you've gotten voice into packets, it's just more data. We can start doing wonderful things with it, like speech-to-text applications and many other things in the future.

But the thing about voice is it's old. In fact, the public switch telephone network operates up to 3400 Hz. It's been that way since the 1930s when it was first architected. It's one of the oldest and most antiquated pieces of technology that even I use on a daily basis. Can't we do better than that?

So VoIP isn't just about doing voice over IP, it's doing better voice, and we call that business-class audio, where we can significantly expand the frequency range of the technology and just deliver a better user experience. We think this idea of going to wideband audio is in fact a game changer for how voice is used in the future. What I wanted to do is give you a quick example of that. The first clip will be a

normal phone call. And the second clip will be a wideband phone call, or a business class audio phone call. Let's just listen to the two, and I'll let you decide which is a better experience.

[Phone clips play.]

Pat Gelsinger: Isn't that stunning? The difference between those? And if you think about that, imagine yourself in a multi-party call, all the noise and other things that distract in those. Well imagine an international call, or maybe as you go to people with English as a second language or you're crossing language barriers? You go to phonetic languages. What an extraordinary difference that can make in comprehension and effectiveness of disbursed teams.

So we're working to drive and make business class audio a reality. VoIP is growing extremely rapidly. There are over 1,100 VoIP providers now in the U.S. alone. Skype exemplifies this as they've been just an incredible engine of driving this into the marketplace. As you can see here, 150 million downloads so far. Skype alone represents 46% of all VoIP minutes across the globe today, and I'm very happy today to announce a partnership between Intel and Skype to make their clients better on our platforms, utilizing our software technologies, our Codec technologies, dual-core platforms, and improving the number of participants in the calls and the quality of calls as well. Intel and Skype -- driving and creating new experience for seamless collaboration. If you've not gotten the chance to look at this yet, take a look in the enterprise zone downstairs.

So we've looked at embedded IT, we've looked at seamless collaboration. The final area I want to explore is information assistance. Corporations are dealing with an incredible amount of data, and the rate that data is being generated is becoming even more rapid, particularly as we go to autonomic data sources, sources that you'll operate and create data without the intervention of users -- like RFID or sensor networks. We also see that business intelligence is becoming more critical as we go to analyze and make use of that data. Things like personal Google Desktop or X1. I just live on X1 as a way I keep track of my information and management. As you deal with this, you need to have highly scalable solutions that just demand performance in incredible ways. I love these types of applications, because they're just MIP-sucking applications. That's sort of what we live for as silicon guys, things that just demand incredible amounts of performance.

All of that is what Information Assistance is about. What can we do, and how can we deliver that capability in our platforms. The first leg of that is just performance. You heard Paul Otellini talk about this quite a bit in our shift to multi-core. It's just the better to deliver performance. In fact, as I described in the spring, the fastest way of performance improvement in the history of Moore's Law is going to occur over the next five years. As we transition to a multi-core, transitioning all of our products to delivering dual-core and beyond. We have, as we described in the spring, 15+ dual-core products underway, we have 10+ quad-or-higher-core products under development as well today. So dual is just the beginning. Next year when we're on stage, we'll start showing you the ramp race for our

quad-core and beyond product lines as well. This is a major shift in how we develop and deliver performance inside of the thermal envelopes of our platforms.

A critical, critical challenge in making this transition to dual, quad, and beyond, is how you develop the software to utilize it. It's this move to threaded applications. We're investing heavily at Intel, and others in the industry with us, with the software tools, the libraries, the training classes – all of the things to make this a reality. This is what Intel Software Tools and Support is all about. You know, maniacally focused to making this transition to threaded applications. Today, we deliver thousands of seed units of our next generation, you know, dual and beyond platforms, to our customers. We have over 200,000 developers that today are using Intel Tools. You might have heard of a few of our software customers using Intel Tools today. You ever hear of IBM? They're one of our Tool customers. Oracle? Symantec, Novell, Adobe, Macromedia – these are all customers of the Intel Software Tools and Support, helping to drive that transition to threads and beyond.

What I'd like to do is take a quick look at our platform offerings, and how we're delivering those as part of our server offerings, taking the first look at Xeon. Xeon today, we have our DP servers based on a Lindenhurst platform – very, very successful platform. Earlier this year we launched the Truland platform, the MP platform of choice for the industry. Over six million 64-bit platforms shipped today. Next step? Dual-core. We are very happy we announced last week the acceleration of our dual-core products for DP, the Paxville DP product

line. That'll be shipping beginning in early Q4. We also announced the MP product. Again, we pulled that from '06 into '05, and we'll be shipping that in volume in Q4 of this year as well. The MP platform needed more bus bandwidth, and so we've taken the front-side bus frequency up to 800Mhz front-side bus, and added the VT capabilities to that platform as well. So delivering dual-core capabilities into today's industry-standard platforms.

Let's take a quick look at what kind of performance you can see, based on the transition to dual-core platforms. What we have here is Sungard, a financial analysis and risk assessment application. So it's a MIP-sucking application --I just love those babies. What it's doing is a Monte Carlo simulation of a financial portfolio. These kinds of simulations would run literally for hours over somebody's portfolio, to do a risk assessment.

We have three platforms here. We have an Irwindale, today's DP platform. Next to it we have a Paxville, so a dual-core DP platform. Next to that we have the Paxville MP. So a four MP system running dual-core. And if we can start the simulation, please. We'll see the four of these running the same workload across the three platforms. On Irwindale, it's an HT-enabled dual-core platform, so you see four threads of operation. Using Paxville DP, we have two sockets with dual-core with HT-enabled, so we see eight threads of operation. And with Paxville MP we have four sockets with dual-core with HT enabled, so a full 16 threads of operation. You can see very quickly the performance difference that is delivered by the scaleable platforms that we have DP, dual-core, and MP platforms as well. The next generation

of dual-core performance delivered in Q4 on the industry standard platforms.

But wait, there's more. Also on that same Truland platform, upgrading it even further with our next-generation MP processor development. We call that Tulsa. Tulsa adds to our capabilities a large-share L3 16-megabyte cache, a whopping 1.3 billion transistors, and I'd like to give you a look at our first silicon on Tulsa. So, 16 megabytes of cache, 1.3 billion transistors per die. So you're looking about a quarter trillion transistors on this one wafer. Tulsa will be a product in the second half of next year, and we'll add it into our Truland family of products, which will further upgrade our MP performance capabilities. Truland, Tulsa, the next generation of MP computing.

[Applause]

Pat Gelsinger: A quarter trillion transistors? Man, I remember those little wafers, a few thousand used to be cool. So that's the MP platforms, let's move to the DP space.

In the DP space, today's Lindenhurst platform has two front side bus loads, and a shared bus platform. Lindenhurst a very, very successful platform. Next generation is Bensley, which we'll begin delivering in early '06, and Bensley has separated front-side buses. This gives us increased bandwidth, because of the dual buses. Essentially you have a switch inside of the chip set. We also add four channels of FBD, so we double the number of memory channels with it, plus with FBD we're able to deepen the memory channel. So the effect of this is front-side

bus bandwidth triples to 17 gigabytes per second. We triple the amount of memory bus bandwidth that's delivered to the platform, up to 17 gigabytes per second. We quadruple the memory capacity. And many, many of the workloads that you want to run on such a platform are significant workloads as well. So this is the next generation platform, Bensley.

So the results of Bensley are incredible performance improvement, but even more importantly, based on our next-generation dual-core, and the next generation microarchitecture, substantial performance per watt improvements. It includes all of the *Ts -- VT, AMT, as well as I/OAT capabilities. It also includes improved I/O capabilities with both RAS and RAID. We also have a version of this specifically for workstations called the Glidewell platform as well.

So let's take a little bit deeper look inside of the Bensley platform.

So we have two generations of our 65 nanometer process technologies. Our Dempsey CPU and our Woodcrest CPU. In it we include FBDIMM. So right here you're looking at eight channels of four gigabyte FBDIMM, so 32 gigabytes of memory in this board. In particular, I want to point out that in the FBDIMM a very strong community demonstrated. We have 11 different FBDIMM vendors showing interoperability here. In this case we're very thankful that Micron announced their four gigabyte DIMM products, and that's what we're showing here on stage, so thank you to them. And this platform it's very, very critical because server customers buy a platform and love to have longevity in that platform. And this exact same Bensley

platform will take Dempsey, begin shipping early next year, and when we introduce Woodcrest later next year, drops into exactly the same platform.

Another challenge of datacenters is compute density. They want to be able to put more in the compute centers than they have in place today. Many of those datacenters are aging. They have fixed amounts of cooling, fixed amounts of power into it, and they're trying to grow their business at the same time and deal with things like rising utility costs. As a result of that, there's a great deal of focus on compute density.

What you see up on the slide here is the wall of ATCA that I showed you at Spring IDF. Because the same challenge that datacenter guys are dealing with is exactly what network operators deal with. We see enterprise datacenters and Telco central offices, but there's a new category as well, what we call the network datacenter. It sort of looks like an enterprise datacenter, but also sort of looks like a Telco data center. To address this, we're continuing to drive forward the ATCA standard, and we're seeing that grow very rapidly in the industry, as well. Again, density – MIPS per watt are critical factors in delivering dense, high performance computing.

With that in mind, I wanted to show you another processor that we're announcing, the Sossaman Processor. If we could start the demonstration; it takes a little bit. They actually started it a little bit ago backstage. This is a highly computationally intensive rendering

application doing full ray tracing of this entire image. So, we're going to have that running here, and I'll explain what that is in a second.

With Yonah, the Sossaman is based on the Yonah core that Paul described yesterday. So, we've taken this high performance, dual-core, power-optimized microprocessor and applied it to the Telco and data center application, as well. We'll be delivering Sossaman at the thermal envelopes of both 15 as well as 30 watts. So, we'll have an ultra-low as well as a low power envelope for that part. What you see on stage here is that you have today's standard platform, a low-voltage Nocona platform and an ATCA form factor. We're also showing our next generation Sossaman in an ATCA form factor, as well. They have approximately the same power envelope, again meeting the ATCA specs. But, you can quickly see that Sossaman is well outrunning today's Nocona. So, you see the performance benefits of that technology.

We'll also take Sossaman and deliver in rack mount servers such as this 1U server that I show here. We'll also be delivering that as part of high-density data center blade offerings, as well. Sossaman offers substantial improvements in MIPS per watt capabilities for data centers and network datacenters. As you can see here, Sossaman's finishing up substantially ahead of today's standard product with the Nocona offerings.

This is just the first step. We're moving to a next-generation microarchitecture.

Pat Gelsinger: Just as we're making the transition from last generation microarchitecture leaders to the next-generation of microarchitecture from Intel, picking the best of NetBurst as well as our mobile architecture to deliver the next-generation core platforms. For the business and enterprise customers, we're taking full advantage of that.

We'll have two versions of that core that deliver products – the Conroe family of products; highly power-efficient cores, moving the thermal envelope of the desktop client to 65 watts, two cache sizes – the two-megabyte as well as the four-megabyte version of that as the next standard for business clients. So, it's very exciting and offers substantial performance improvements.

We're also taking that technology and applying it to the server space. It's very important for MIPS per watt and compute density that are capable. We'll have the two-core version with Woodcrest. We also have the four-core version of that under development for delivery in '07 called Whitefield. Woodcrest with four megabytes of cache and Whitefield with 16 megabytes of cache --significant performance improvements at reduction in power envelopes; performance without suffering any loss or challenge in the thermal envelopes in the platform. As you can see, Whitefield would be one of our first quad-core offerings to the marketplace as well.

What I wanted to do was show you, we have Woodcrest as well as Conroe running here. Here is a Conroe platform running. You can see the demonstration of Conroe running a high-performance rendering application. We also have a Woodcrest running in Bensley, so this is

Woodcrest running inside of the Bensley platform I just described, and you see that running here as well. Conroe, Woodcrest – the next standard of business computing for servers and clients.

So you see by the chart here substantial improvements in performance as a result. As we move from Irwindale to Dempsey, 88% performance improvement. As we move from Dempsey to Woodcrest, another 50% performance improvement delivered next year. Huge jumps in performance, but even more important and substantial is the performance per watt improvements, or the efficiency of those architectures. And these are measuring at a rack level, so the full system benefits that result from this, over 100%performance per watt improvement as we go from Irwindale to Sossaman and as we go from Sossaman to Woodcrest, another 60% performance per watt improvement -- the unquestioned performance per watt leaders in the industry in 2006. Stunning new products, and together we're excited to be delivering scalable microarchitecture with absolute leadership in the marketplace next year.

So summarizing all of that is this simple little picture. So this is everything we've tried to describe with our server roadmaps, Sossaman, our next generation Bensley platforms, the migration from Truland platforms today and bringing multi-core capabilities. Describing quickly some of the first quad-core products we'll be bringing into the product line, and quad-core and beyond capabilities such as with Whitefield and our Dunnington development, and a number of other quad-core and beyond developments that are now underway.

Finally, I wanted to take a look at the Itanium family of products as well. Building on the success and momentum of this family, wanted to explore what it is that's happened since spring IDF. Poulson, Tukwila, also quad-core or beyond products that we have underway in development today.

What is the value proposition of Itanium? How is it that we're trying to engineer and deliver this product to the industry? Well, performance, reliability, scalability, and finally, choice.

Montecito will deliver a huge step up in performance, more than doubling the performance capabilities of today's Madison family of products. But more importantly, we see the industry boiling down to a two-horse race. That's Power and Itanium. So how do we stack up versus power? These are our estimates and benchmarking versus the power architecture, and as you can see, Montecito looks pretty hot.

In addition to that, we see there's just an incredible amount of interest in utilizing the technology. I wanted to show you an end user example as we started to seed Montecito systems into the customer base, and to help me look at that, I have Chuck Duvall joining me on stage. We'll look at how SGI and Halliburton are taking advantage of Montecito technology.

Chuck: Hey, Pat, how's it going?

Pat Gelsinger: Good, Chuck. Thank you. So can you tell me how SGI and Halliburton are utilizing Itanium?

Chuck: Absolutely. Really there are three companies that are changing the way we look at, discover, and recover those hydrocarbons or oil and gas from the ground. Now we've got Sean in the corner over there, and he's going to run an application from Halliburton called GeoProbe. GeoProbe takes terabytes of data and packages it all up in a way we can actually look at -- really a 3-D model of the ground beneath our feet. Now while he's going through all that data, we'll look at another application real quick called VIP.

Now VIP takes all those hydrocarbons that are in the oilfield and really tries to calculate how they're going to flow to the oilfield. It's just a mess. We have varying temperatures, varying pressures, varying viscosities, varying thickness of the [oil shale].

Pat Gelsinger: It sounds hard.

Chuck: Yeah, it's a mess, like I said. So with the help of SGI and Prism, today we're running that application on our Montecito dual-core Itanium platform.

Now up behind us you can see some numbers, and that thing is going to crank away for literally hours, because it's very complicated to do a whole oilfield. It can be miles, miles wide. So we've actually cheated a little bit. We've already run that application, and we've incorporated that VIP into our Prism SGI system. Let's take a look at not only that

seismic data, but all that VIP information, and look at that oilfield on the screen. So pretty soon we're going to see a graphic here, we're going to see all the oilfield, actually how it really looks. Now this is compliments of SGI, and you can see that 3-D model that we talked about, that's actually underneath the ground.

Now down below, in the very back, that block he's pulling forward, that's actually the VIP data that was calculated on a Montecito system. And that looks at the viscosity of how those hydrocarbons are going to flow. You can see the pipelines that have been drilled into this oil system, and that's actually – we have to calculate where those are going to reside that, because these oil wells are extremely expensive to drill, $90 to $100 million.

Pat Gelsinger: $100 million a hole.

Chuck: $100 million a hole, exactly. We have to make sure we put it in the right spot. Because not only do we have to extract that, but we might actually have to pump in actually air or water or steam to help extract and push those hydrocarbons to the surface. So we want to make sure we get it in the right spot. So SGI and Prism allows us to look and plan much better than we did in the past.

Pat Gelsinger: So what are the results? How did we do?

Chuck: Well, before, we take, we took these tools and systems, we had a hit rate of about 20%.

Pat Gelsinger: 20%. Hmm, 80% holes.

Chuck: Not so good. Now because of systems like this, and SGI and Halliburton solutions, we're actually pushing above 70%, so the return on investment is very, very high.

Pat Gelsinger: So 20% to 70% on $100 million holes?

Chuck: Exactly.

Pat Gelsinger: Wow! That buys a lot of Montecitos.

Chuck: It does. You better raise your forecasts. Thanks Pat.

[Applause]

Pat Gelsinger: We're also delighted with the Itanium's success in extending and enabling reliability. Since last IDF, HP announced their "Integrity," the non-stop family of product lines. They promised seven nines platform, not five nines, sort of the industry norm, but seven nines. And this just looks at the difference in downtime for that. In other words, these machines are built to never go down in their lifetime. Itanium is now delivering unprecedented levels of reliability.

It's also scalability, and this is shared memory scalability. Some applications don't scale over clusters; they need a shared memory environment. Comparing against the power, 64 versus 128 – larger SMP systems can be built with Itanium. Furthermore, as you look at

HPC applications, 512-node SMP application – oil, gas, life sciences, things like we were just demonstrating here. Together with our IPF partners, delivering unprecedented scalability in the industry as well.

The really exciting thing that's happened since Spring IDF is the continuing excitement of end-users as they begin applying this technology to their problems. Today, 43 of the world's 100 largest customers have Itanium deployed. Let's take a look at some of the customers who have deployed it today.

So we have end-users, we have incredible momentum in delivering scalability, reliability, and we have hardware. In fact, this is seven tons of Itanium. We have the SGI system that we were running just a minute ago, I described the HP nonstop systems, at the last IDF we showed Hitachi and Unisys Itanium systems. Since the last IDF, Fujitsu has announced their "Prime Quest" product line, and NEC and Bull, which had already announced their Itanium family of products. Literally, the industry has lined up, all but one, with Itanium as the platform of choice for risk replacement and for mainframe platforms for the future. An incredible amount of industry strength and momentum in delivering IT, IPF to the marketplace.

[Applause]

Pat Gelsinger: And it's also choice. This stuff is just cool. Big iron makes you feel good, right? It's also the choice you get in operating system environments, and through those many hardware providers. We see

unprecedented level of operating system support, as shown in this slide as well. Some of those the proprietary offerings of those vendors, others industry standard offerings like Linux and Windows. And Microsoft's continuing commitment and support for Itanium as part of their Longhorn products as well.

But it's also the choice of applications. In the last year we've seen a doubling in the application availability, Itanium and we estimate over 5,500 applications now available in the architecture. In the last year we've added applications like Veritas Cluster Server, McKesson Horizontal Patient Porter, and Symantec ESM. Major applications extending the platform offerings. Together with developers, we're delivering applications that business need.

Today we've seen the four pillars. How we, together as an industry, are delivering the embedded IT vision, the first steps in seamless collaboration, and fundamentally delivering the power and the tools for the information assistance. Together we're doing that. We also highlighted a number of opportunities. New business opportunities. New technology innovations that are occurring. New usage models. Finally, together, let's transform the enterprise and change the world. Thank you very much.

[Applause]