

Hector Ruiz, PhD.
President and Chief Executive Officer of AMD
LINUXWORLD Keynote
"Playing Well With Others"
New York, NY
January 22, 2003

DR. HECTOR RUIZ: Thank you for being here and giving me the privilege of sharing some ideas with the Linux community.

The theme for today's LinuxWorld is "Where Open Minds Meet". This is appropriate because when you have an open mind, you find it easy to share and work with others toward a common goal.

The values of working together and creating an open exchange of ideas are so fundamental that we all learn them at an early age.

As we get older and become adults, we seem to lose these powerful instincts. Of course, I am referring to adults who are not part of the Linux community, which is an excellent example of what can be done when you work well with others. Look at all the creativity and innovation that you are responsible for sharing with all of us.

If we take our cue from this community, perhaps we need to start measuring success with something other than traditional criteria. Perhaps the skill of "Playing Well With Others" should be a standard metric for success in the technology world.

Today, I would like to further explore this notion of "Playing Well With Others."

What does playing well with others mean? At AMD, we believe it means being a good partner, listening to customers and taking users and their needs into consideration when we design our products.

But what does it mean to the Linux community? ... and the world outside this building? We decided to find out.

In order to do that, we sent Nate Mooney, AMD's own roving reporter, into the streets of New York and the corridors of the Javits Center to find out.

(Roving Reporter video plays)

DR. HECTOR RUIZ:

None of us had any idea that playing well with others had caught on so well in New York!

But seriously, the notion of "open minds meeting and playing well with others" is a simple, yet incredibly powerful, statement that makes this community so interesting and influential in the technology world.

All of you here today know that collaboration, shared values, and group effort are powerful drivers of creative thinking. Otherwise, you would not be here.

The best creative thinking can inspire the best in new technologies, new applications and new solutions that benefit all of us.

Today, we would like to showcase some of that creative thinking. You will hear from people you know and respect; developers, enterprise users and industry leaders, who are dedicated to Linux and inspired by the force of open source creativity. We will introduce you to some of the creative solutions, breakthroughs and applications that result when Linux and AMD processors play well together.

First, though, I would like to tell you why Linux – not just the software, but the idea – is very important to us at AMD and the future of technology in general.

If you are like me, you became interested in technology because you saw its potential for changing the world. You saw the promise of silicon and software and the awesome effect that it could have on people and you were hooked.

Then you joined the real world. And, after a little bit of time, you came to know that there are some subtle conditions that must be maintained in order for technology to have the kind of impact you thought was possible.

One of those conditions – arguably the most important – is to preserve a healthy sense of competition. Like Linux, AMD is a poster child for competition.

In my view, competition is that great “invisible hand” that keeps the engine of technology evolution rolling along. It is what leads to an evergreen environment of new ideas, products, categories, markets and opportunities.

Without a healthy sense of competition, we believe that customers and end users would not have access to what we call “true innovation.”

What is “true innovation?”

True innovation is innovation that makes the greatest possible technology available to the widest possible audience.

It is not just a factor of bigger and better, but of availability and relevancy to end-users. Availability includes not just price, but compatibility and accessibility as well.

Relevancy means that more people than just the engineers who designed the technology find it important, useful and significant.

Today, we are seeing a growing number of technologies that hold little relevance to people’s lives. They cost far too much to buy, are too complicated to be successfully adopted.

Although they may help one company succeed, they serve very few others.

Allow me to provide a timely and enlightening example.

In a new era of “true innovation,” what would you call the next generation architectural approach...

... that blocks legacy programs...

... that dictates a complete applications overhaul...

... and forces customers to take a huge new risk and hold on for dear life?

Some might call it arrogance.

Others might refer to it as a step in the wrong direction.

If you are Intel, you call it...Itanium.

Is this an example of technology that is truly innovative? We say "No." At AMD, we believe that the time has come for semiconductor companies to adopt a more enlightened view of the difference between technology... and "true innovation."

In order to foster and take part in true innovation, we in the semiconductor industry cannot continue to operate on a business-as-usual basis. We must consider a new business model, one that is more collaborative and connected: more like the Linux and open source community.

When we, at AMD, started planning for our next generation processors, we did not start with the technology. We started with you, our customers and your customers.

We listened.

The result is our upcoming new family of processors: AMD Opteron™ for servers and AMD Athlon™ 64 for desktop and mobile computers. Many of you in the room have given these your personal stamp of approval and we thank you for that.

Many of you have selected the AMD path already. You saw the promise of control and lower total ownership cost that the AMD and Linux combination provided. As a result, we have already seen great acceptance of our AMD Athlon MP processor running in Linux clusters around the world.

In Germany, Daimler Chrysler depends upon AMD Athlon and Linux for both engine and performance for crash testing for the Mercedes line of vans and trucks.

In Russia, the country's most popular Internet portal, Rambler, is using AMD-powered servers to satisfy over two million search requests per day, 50 every second.

In the Netherlands, Dow Benelux has realized gains in research and development efficiency by using an AMD-based Linux cluster to model fluid dynamics.

Here in the United States, Rackspace, the world's second largest web hosting company, depends on thousands of AMD servers to deliver reliable, cost-efficient web hosting to their customers.

These leaders in their field have chosen AMD for the most demanding real-time data intensive application environments that you could imagine. They might have chosen us for our performance, but we believe they will stay with us for our approach.

At AMD, we long ago made a strategic decision to use compatibility and open standards as a design objective because we see the best path to the future grounded in the best of the past. We believe in preserving our customers' data from platform to platform and application to application.

That is why AMD's next generation processor architecture is the only one that is both fully compatible with today's 32-bit systems and completely 64-bit capable when you have the performance needs to fully upgrade. We've simplified your transition to 64-bit Linux environment through native high-performance support for simultaneous 32- and 64-bit applications. The 32-bit capabilities of our forthcoming 64-bit processors are already supported by numerous 32-bit Linux distributions.

In addition, 32-bit applications ported to 64-bit in some cases can even show increased performance on the 64-bit Linux running on our 64-bit processor. Maybe that is why many key ISVs and IHVs are porting to x86-64 Linux.

And, unlike our chief competitor, you do not need a special invitation to participate in the x86-64 Linux port.

It should come as no surprise that 64-bit computing is especially important to enterprise and databases in particular. One of the greatest advantages of x86-64 bit processing power is the ability to address more than four gigabytes of memory. You can see that power in action by visiting either the AMD or Computer Associates booth, where a demo of CA's Ingres Enterprise Relational Database is being shown on AMD's x86-64 platform.

I have more exciting enterprise database news. I am happy to share with you for the first time today the news that AMD and IBM are announcing Beta availability of the 64-bit version of IBM's DB2 database software for the forthcoming AMD Opteron processors. A demonstration of that beta DB2 database release running on a SuSE Linux operating system and an AMD Opteron processor is being featured in our booth.

Providing the beta DB2 release to our customers and partners will help ensure that DB2 is tested and optimized when systems based on [AMD] Opteron processors are available in the second quarter of this year. With AMD's x86-64 technology, DB2 customers will have a powerful alternative for 64-bit database computing and the ability to support both 32- and 64-bit applications in an x86 environment.

Another leader in the field that has chosen not just AMD, but our next generation AMD Opteron processor family, is Cray, Inc., a company whose name is synonymous with the term "supercomputer".

When Cray came to us, they were looking for a partner who could help them develop one of the most powerful and reliable computers for the people in the US government's famous Sandia Laboratories.

Sandia plays a critical role in our nation and in our world's defense. They were asked to develop the means to rapidly accelerate the ability to simulate advanced computer modeling and testing of nuclear weapons. Sandia turned to Cray and I am proud to say that Cray turned to AMD.

The result is a joint effort to build the world's most powerful computers based on Massive Parallel Processing Architecture. To give you an idea of the sheer scale of the machine, it will be powered by more than 10,000 64-bit AMD Opteron processors working in concert together.

AMD was chosen to provide the processors for such a challenging feat. It will easily rank in the top engineering achievements of the next few years because of the way we listen and work with our customers to ensure that they got just what they needed.

Let us take a look at what this new super computer is all about.

[Cray video]

DR. HECTOR RUIZ: Joining me now to discuss the Red Storm supercomputer in more detail is Wayne Kugel, Executive of Operations for Red Storm and Professional Services at Cray.

What does Red Storm mean to Cray and to the high-performance HPC market?

WAYNE KUGEL: One of the things that Red Storm does for Cray is that it acknowledges

that our expertise and technology is very important and vital to the HPC community.

Even more importantly, we are seeing a resurgence of scientists, engineers and other groups throughout the world looking to capability types of machines and massively parallel, highly balanced machines to solve tomorrow's grand challenge problems. They are also looking to off-load some of their heaviest computational work in today's work environment.

In the automotive industry, to ensure a quiet comfortable ride, they are working on aero-acoustics, which is simulating the sound within a cabinet of a car given different velocities and tire and engine noise. In addition, they are also simulating the engine compartment with all the moving parts: the liquids, the combustion and all the different pieces to optimize the power and the fuel consumption of the engine. Most importantly, they are also doing safety by high fidelity simulations of crash and airbag dispersion.

If you look into the global environment for weather, they are actually starting to forecast weather at the global scale. Last November, there was an announcement from Japan that they were doing this. They were actually modeling atmosphere, atmospheric warming, global warming, and looking at the characteristics of El Niño and El Niña all the way around the globe down to your local park, for which you can get a forecast for this weekend to make sure that you can have a good picnic.

In healthcare, they are simulating millions of compounds in the human chemistry and determining at the molecular level how to reduce the toxicity of cancer drugs and increase their effectiveness. Hopefully, one day we can actually simulate a cure for cancer on these supercomputers.

The problems run on these types of systems are dictated by the applications suite that is run on the systems and the characteristics that they bring. We look at processor speed, but more importantly the processor architecture. We also look at bandwidth in several ways: from the processor to memory, from the processor to processor, from the processor to the network and from the processor to the storage devices.

We also look at it from Cray's perspective as what we call "by section". You take a slice through any of the computer systems and actually see what the weakest point was. For instance, if you were to take a slice out of Red Storm at its weakest point, it is actually 1.5 terabytes per second.

As we look to the types of computing, we also look at the numerical treatment of those problems. In some cases, you need massive vector super computers, as most people would know Cray, or you would need a highly balanced parallel system like for Red Storm. In Red Storm, we are looking for a 10,000-processor machine. The first six rows by as far wide as it is here is the first instantiation of Red Storm. It is actually designed to go as far back as the cameras back here. The first instantiation is 40 trillion calculations per second.

To make sure that we could do appropriate simulations on these kinds of machines, we blended two basic things together: the AMD Opteron processor and an innovative, interconnected high-speed network.

DR. HECTOR RUIZ: That's amazing. We're flattered that you chose the [AMD] Opteron processor for all the reasons that you mentioned. Can you tell us what role that plays in some of the work you do?

WAYNE KUGEL: Linux is prevalent throughout the HPC industry and many of the users at Sandia National Labs are part of the Linux community. We wanted to make sure that when they go to their next system, they have a very familiar environment. So their user

environment will be Linux. Many of the components and features of Linux, such as the high-performance file system, are Linux-based.

The community is constantly bringing to bear new innovation and functionality, things that are required for Red Storm and will bring us into the future. From a Cray perspective, Linux is opening brand new avenues that we have not entertained yet. We are excited that Linux is bringing us to the next step.

DR. HECTOR RUIZ: We are delighted that you came here. Thank you.

WAYNE KUGEL: Thank you.

DR. HECTOR RUIZ: Let me now shift to another kind of supercomputing: computing that results in the creation of super heroes, super aliens and beloved cartoon characters.

Rhythm & Hues is one of the largest and well-respected visual effects houses in the entertainment industry. They have won numerous awards for their work in film, television and advertising, including Coca Cola's popular polar bear commercials.

What may interest you most of all is that the talented people at Rhythm & Hues are creating dazzling visual effects in a predominantly Linux environment powered by AMD Athlon MP processors. Here to tell us more is Mark Brown, Vice President of Technology at Rhythm & Hues.

Mark, tell us a little bit more about Rhythm & Hues and who your customers are.

MARK BROWN: Rhythm & Hues creates computer-generated visual effects primarily for feature film, although we do commercials and theme parks as well. We are located in Marina del Rey, California and have between 300-450 people depending on the staffing we need for our work load. Our customers are all of the major Hollywood studios – Warner Brothers, Twentieth Century Fox, and Paramount – and the major advertising agencies and big theme park venues.

We have been nominated for a number of Academy Awards and won the Oscar for Best Visual Effects for the film "*Babe*," which, at the time, was groundbreaking material.

DR. HECTOR RUIZ: Why did you decide to run your business on an AMD/Linux environment?

MARK BROWN: For many years, most of the top facilities that did computer-generated imaging used proprietary SGI hardware and software. That had advantages and disadvantages. One of the biggest disadvantages was that it did not allow us the flexibility of using commodity hardware. With the decreasing prices of that hardware, it was getting hard to justify spending five times as much for one third of the computing power.

Having come from the SGI environment, we are very Unix-centric at Rhythm & Hues, so moving to a non-Unix environment was not an option for us. We had too much invested in the infrastructure to even consider that. The move to Linux allowed us to leverage off all the commodity hardware and take advantage of the huge support and development base of the Linux community. In addition, it did not significantly change the look and feel of our environment and, because Linux is open source, it allowed us to make customizations and enhancements that we could not with the SGI proprietary platforms.

Rhythm & Hues is always looking for technology that delivers a superior price/performance ratio, so the choice of AMD was a natural one. When we started looking into Linux at the desktop, we went through a rigorous benchmarking process and found that in our environment, where we use both proprietary and third-party software

products, the AMD processors were performing as well as the competition, and in most cases better, for less than half the total system cost.

DR. HECTOR RUIZ: What systems are you currently using? What do you think of the move to a 64-bit world?

MARK BROWN: Our current preferred desktop machine is a dual processor AMD Athlon MP 1800+ or an AMD Athlon MP processor 2000+. We have a gigabyte of RAM and we use an ATI FireGL graphics card and have a 40-gigabyte IDE hard drive.

What is unique about our environment is that these are not deskside configurations or workstations. We worked with Angstrom Microsystems of Boston and came up with a rack mountable 3RU unit. It is pretty non-descript, but it is very functional for us. We can house that unit in a machine room environment and then run the monitors remotely to the artist workstation, using a Cybex KVM box. That allows us to keep the noise down at the artist workstation and monitor the processors and power condition in a monitored environment.

As to the move to 64 bits, our scene complexities are going through the roof. We cannot even tell you the large memory requirements we have. We are excited by the concept of 64-bit AMD Opteron-like processors and that we get huge per-process memory footprints, we basically do not have to do any re-engineering of our code, and can drastically enhance our scene complexity.

DR. HECTOR RUIZ: I understand you have brought some of the exciting work Rhythm & Hues is doing. Can you show us something?

MARK BROWN: Absolutely. I brought some visual effects from "*The Sum of All Fears*" and character animation from "*Men in Black II*," "*Scooby Doo*" and "*Cats and Dogs*."

(video begins)

What you see here is a completely computer-generated missile silo. Everything in this is done in the computer. As the bubbles pass, you will see the underlying computer model.

This is a totally computer-generated missile launcher. The background is actually a simple still photograph.

A blast plate ... all of this is computer-generated except for that one car.

This is a computer-generated helicopter and a computer-generated nuclear blast wave. Again, particulate matter.

The helicopter is actually real in this one; we had to move around it.

This is the final scene. There is no city in the background. All the explosions are computer-generated and it was actually filmed in a meadow.

This is a little bit of our character animation and creation. What you are seeing is Lou the Puppy from "*Cats and Dogs*." We are showing you how we drop the skeleton in. We put muscles on it, skin on it and then we put the fur on it.

This is "*Scooby Doo*." This is a completely computer-generated character that we have to put into a live action scene. He has to interact with the actors. He is a major character, so he actually has to be an actor himself. There was actually a man in that dress that we replaced. This is an animator animating Scooby. Here is Scooby

interacting with a live-action environment.

This is Mr. Tinkles from "*Cats and Dogs*", a fully computer-generated Persian cat we had to seamlessly intercut with a live Persian cat.

Scooby really interacting with his environment. High-brow humor.

This is a ninja cat from "*Cats and Dogs*." He is a stylized Siamese cat, all computer-generated.

This next scene we have never seen before of course.

Scooby Doo in a live action environment.

This is the final scene from "*Men in Black II*." This is the live action plate we had to work with. Now, these are the aliens that Rhythm & Hues created to populate the scene at the end. This is just three of them; there are actually six in all. This is the final scene as it showed up in the film.

We will leave you with Mr. Tinkles making a parting comment here.

MR. TINKLES: At what point did you forget that we are trying to take over the world?

(video ends)

DR. HECTOR RUIZ: Incredible stuff. There is a thin line between reality and non-real stuff. Seeing what is possible at Rhythm & Hues, it is going to be unimaginable when 64-bit processing power is put to the test.

I hope you now have a good idea just how well AMD plays with others. One of the first things we learn about playing well with others is the importance of sharing. While AMD is, and has been, committed to playing well with others for a long time, I want to make use of this prestige occasion to introduce and reinforce this commitment even more.

We have been listening and have heard the Linux community say two words: "performance libraries."

I am proud to announce today that AMD is now working with the Numerical Algorithms Group (NAG) to develop and freely distribute the AMD Core Math Library in support for AMD's upcoming AMD Opteron and AMD Athlon 64 processors.

With this library, we will provide developers and customers with another tool to develop and deploy computationally intensive applications for systems based on our processors.

Developers will be able to reduce the overall time and cost of development for both 32- and 64-bit applications, as well as systems running on x86-64 technology.

In addition, although we are giving ACML away for free, we will fully support and maintain it. AMD is listening, we hear you, and we are assembling the resources our customers need to optimize software and hardware for high-performance server and clustered server environments.

There is more, however.

To proactively meet the needs of the High Performance Computing community, AMD will release detailed documentation about AMD Opteron processor that will allow the open source community to develop Linux BIOS for large Linux clusters based on AMD Opteron

platforms. We plan to release this information by the time we launch the AMD Opteron processor.

Why are we doing this? Because we listened to you and it is the right thing for you, for us, and for the industry. We understand that other people have to make money in this business today.

If you succeed, we succeed.

We know how interconnected Linux programming is with the hardware. It is time we gave you what you need to innovate, to create, and to make the best possible products.

We are here to support you, the open source community. That is why we have memberships in Linux International, the Free Software Foundation, the Free Standards Group, and the United Linux Technical Advisory Board. Our goal is to make [AMD] Opteron the 64-bit platform of choice for the open source community.

One of the Linux companies AMD is actively working with to give you what you need to innovate and create is SuSE Linux. SuSE Linux is in the forefront of porting Linux to x86-64, providing all kernel, development tools and basic device drivers to the Linux community.

Joining me now is Richard Seibt, the new CEO of SuSE Linux. Richard spent a great deal of his career working with operating systems at IBM, so he certainly knows this business and I know he's excited to take the reins at SuSE.

Richard, what are your perceptions of the AMD 64-bit architecture and the upcoming [AMD] Opteron processors?

RICHARD SEIBT: Many of us have been eagerly awaiting the performance enhancements as well as the lower total cost of ownership we expect to see from the AMD Opteron processor. An affordable 64-bit computing environment opens up all kinds of possibilities for Linux users. I think the applications that will soon become available will be staggering.

Let me ask you a question, Hector. Why have you taken the time to be here at LinuxWorld and why are you so interested in 64-bit support on Linux?

DR. HECTOR RUIZ: I was not expecting a test here.

Linux clearly has inherent characteristics that make the transition to 64-bit very exciting, for instance its portability across platforms, which will enable much more rapid deployment of 64-bit technology by taking advantage of current skill sets to build on more advanced platforms.

In addition, there are the advanced clustering capabilities that we have seen demonstrated in Linux; this is certainly where we see many of the really high-performance applications, like Red Storm and Boeing's Delta IV cluster, coming into play. With the combination of 64-bit and Linux in clustering, the possibilities are exciting and limitless.

Beyond this, there are obvious advantages that the open source community brings to the Linux table. Every day thousands of people are working to advance Linux, not because it is their job, but because like most of us, they love what they do.

RICHARD SEIBT: Do not forget the inherent security and scalability of Linux. The community of folks out here in this audience makes Linux what it is today: a robust, mission-critical operation environment. The people in this audience make Linux the

important, disruptive technology that it is and of which we are all proud.

DR. HECTOR RUIZ: Very true, and SuSE Linux should know. I understand that you just celebrated your 10th anniversary as a company, one that has crafted a reputation for high-quality engineering, on-time delivery and a commitment to open source ideals.

You have joined the company at a very exciting time.

RICHARD SEIBT: I agree. With the advance in technology like the 64-bit AMD Opteron processor, the embrace of open standards by companies and governments all over the world and the need for cost effective technology and solutions, I believe we at SuSE are on the cusp of great things.

SuSE engineers have been working all along with AMD colleagues on x86-64 technology since August 2000. We booted a 64-bit kernel on the hardware simulator only four months later. Within a day of receiving our first sample AMD Opteron processor in early 2002, we were booting a full-fledged kernel. Having run on actual AMD Opteron processors for about a year now, SuSE has been very impressed with their flexibility, stability and performance— and you know how hard it is to impress engineers. You probably hear that I have a German accent and know that German engineers especially are sometimes difficult to convince.

We are not the only German company benchmarking the AMD Opteron: another big software company that delivers business applications has gotten great results benchmarking the [AMD] Opteron processor compared to others. They have not disclosed it yet, but we work closely with them because they use the SuSE enterprise server as well. They really appreciate what they got from you.

I think we are particularly impressed with how easy it is to shift from 32-bit to 64-bit applications and maintain high-performance with the AMD Opteron processor when running SuSE's 64-bit Linux. In our experience, most 32-bit applications run just as fast, or even faster, under 64-bit Linux as they do under 32-bit Linux. This is what we all want. All of which means great flexibility for our customers.

It has been a pleasure to work with a company that is as enthusiastic about Linux as you are. It is important to us that when AMD develops new technology in the future – and you will – you immediately and openly introduce it to the Linux community. That allows us to explore together with this community the best way to enable every aspect of the device in GPL source code.

This significantly speeds our time to market and enhances what we can bring to market.

DR. HECTOR RUIZ: Thank you for your kind words and support. We look forward to continuing to work with you. I also want to tell you that from my own experience having a factory in Germany, German engineers are tough to impress.

Perhaps no category of products shows more promise for "true innovation" – both relevance and availability – than the handheld device category. At AMD, our Personal Connectivity Solutions Group is actively designing and developing a range of low power, high-performance processors for this important market.

Joining me now is Metrowerks CTO and Vice President of Linux Solutions, Bernardino Baratta.

BERARDINO BARATTA: It is exciting to be here to talk about something as amazing as our OpenPDA platform running on AMD's Alchemy set of processors.

DR. HECTOR RUIZ: What is OpenPDA?

BERARDINO BARATTA: OpenPDA is a complete platform that allows developers to create consumer electronics devices such as mobile computing devices (e.g. PDAs and smart phones). It brings together all the pieces you need to create these devices such as productivity applications, information management, and entertainment.

We are excited that we are showing it for the first time here at LinuxWorld on the AMD Alchemy processors.

DR. HECTOR RUIZ: It sounds like this is not your father's PDA anymore. I understand this is all on Linux systems.

BERARDINO BARATTA: Absolutely. OpenPDA uses embedded Linux at its core.

We went out into the open source and commercial IP communities and found all the pieces that developers need to create PDA solutions. We brought together things like Web browsers, Java VM, games, and PIM suites, as well as things like multimedia content. We brought all this to many developers to customize these solutions. Some of our customers are already using it for things like Sharp's Zaurus, which is a Linux-based PDA that started shipping last year. There is a Vercel ED1, which is a great low-cost Linux-based solution, as well as Infomart's Kai. There are many other licensees that unfortunately have not announced yet, but we are seeing some great strides with embedded Linux in the consumer space.

DR. HECTOR RUIZ: I imagine that when you partner your software with our low-power, high-performance Alchemy Au1100™ processor, it makes for a solution that I think customers are looking for.

BERARDINO BARATTA: Absolutely. The end customers especially are always looking for higher performance. At the same time, they are saying that their PDA or phone has to last longer, so they need long battery life. The Alchemy Au1100 processor really gives you that. You have 400 MHz speed and great performance as well as low power consumption, which leads to longer battery life.

Even though we are early on, we are seeing unoptimized video playback that far exceeds some of the other processors out there. We have exciting thoughts about how we can optimize this and we are working together to make sure that the video performance, for example, will be the best in the market.

DR. HECTOR RUIZ: If under an unoptimized situation we have what appears to be a performance advantage, I imagine when we optimize it, it is going to get even better and stronger. You have a demo for this in the show. Where can people see it?

BERARDINO BARATTA: You can come down to the AMD booth here at LinuxWorld. We have a station set up where we are showing off the hydrogen platform running on PDA. You can see Web services on it, PIM suite, and some playback of a bunch of cool videos.

DR. HECTOR RUIZ: Thank you for your time, Berardino. We look forward to seeing great things from AMD and Metrowerks.

There are a few people in this world who have a hard time saying the words "mission-critical" and "Linux" in the same sentence. There are even fewer who can figure out a way to include AMD in the discussion. I hope our next segment helps us put this to rest. You will never guess who is using AMD Athlon processor-powered Linux clusters.

It is the people who were there when the term "mission-critical" was invented: Boeing, the name synonymous with flight, whether cruising comfortably at 35,000 feet or carrying satellite payloads into orbit by way of their Delta IV rocket.

Here is something of which we in the Linux world can all be very proud.

[Boeing Video]

DR. HECTOR RUIZ: I think Boeing just put to bed any concerns over a platform not being reliable or stable. I cannot think of a stronger testimony to the stability and reliability of this technology than what you have just seen.

Think about the intensity imbedded in the statement: "In the rocket world, the first time we launch, it has to be fully functional. There is no time for an opportunity to find problems as you go." That is an understatement, but what a testimonial for Linux!

The rocket scientists responsible for one of the most tantalizing computing environments in the world selected Linux for their mission-critical version of true innovation.

All I can say is...look how far we have come together.

At this time, I would like to shift the discussion from the great things the Linux community has done to the great things we collectively plan to do.

To do so, I would like to introduce some of the brightest and most influential stars in the Linux community.

Here to discuss the future of Linux, Linux in the enterprise and what impact 64-bit computing will have across the Linux community – and perhaps computing in general – please welcome Executive Director of the Free Standards Group, Scott McNeil; President of Linux International, Jon "maddog" Hall; Giga Information Group Linux analyst Stacey Quandt; once again, SuSE Linux CEO, Richard Seibt; AMD's own Vice President, Server Business Segment for the Processor Group, and former CEO of Penguin Computing, Marty Seyer; and the manager of IBM's DB2 Strategic Technologies team, Boris Bialek.

I would like to throw out a question to the group. Why has Linux become so important, not so much as an operating system, but as a focal point of a community?

JOHN 'MADDOG' HALL: Two things have caused Linux to become so focused in the community: the availability of the source code and the openness of the development community, which allows other people to join it. It has never been a closed community. It has always been welcoming to people who want to contribute, which is different from many other projects.

BORIS BIALEK: The fun factor is very important. When we did the first databases two to four years ago, the question was, "Is there a back-up software available?" When you see the speed today, the question is not, "What is available?" The question is, "Where is it leading Linux first?" Linux is now becoming the leading platform. When we did the first InfiniBand clusters, Linux was the platform of choice.

RICHARD SEIBT: I think it was very important that they can do what our customers want and demand. Nobody is holding up any development or innovation because of business phase or because the next phase is coming up in two months or three years or whatever. Every technology is available when it is ready. I think the folks out there are doing a great job delivering technology when it is available.

STACEY QUANDT: Another issue is the ability to innovate. With Linux, you can actually touch the code. Microsoft shared source is often "look but do not touch."

JOHN 'MADDOG' HALL: You can easily build on top of the things that people have had before you. You are not blocked from doing that.

DR. HECTOR RUIZ: Stacey, to follow up on this notion, what do you think the next big thing, if there is such a thing, will be that is going to hit the Linux community in the future?

STACEY QUANDT: Without a doubt, it is the release of the 2.6 kernel later this year, which will lead to more innovation, better performance, improvements in memory management, and threading. To a large degree, that will further narrow the gap between Unix and Linux and place more pressure on Microsoft to complete. In addition, the ecosystem for Linux becomes much broader. We will see more deployments with DB2 and SAP, moving the workload even further up the stack. That will place pressure on some ISPs who have been holdouts. There are a few in the CRM space, particularly Siebel, who have been holding out on having a port to Linux. We will see if that changes.

This places tremendous power with users and more ability to leverage their IT spending dollars.

MARTY SEYER: As that is rolled out, do you have any opinion on the reliability aspect of that and what the requirements are going to be? Is that following that next release?

STACEY QUANDT: Obviously the requirements will depend on the work load. We are going to see more application-specific benchmarks in particular to meet those requirements.

JOHN 'MADDOG' HALL: I am very excited about a re-emphasis of the desktop. For a long time, a lot of the distribution seemed to be concentrating more on the server area. I think SuSE having announced their open desktop product and Red Hat having come out with version 8 shows a reemphasis on that desktop. I think that is going to be the beginning. The years 2003 and 2004 are going to be the years of the Linux desktop.

SCOTT MCNEIL: It is interesting to see this as very much an evolutionary process: Linux moving from an edge-of-network technology to now being taken seriously as an enterprise platform solution. As Stacey mentioned, the next iteration of the kernel is going to be an important step in that.

Your question is always a good question when you have a whole bunch of open source developers, because everybody stands up and says, "Well, mine is the most important project of course." I am no exception with the Linux Standards Base (LSB), the interoperability standard that has certification around it and is being adopted by every major Linux distribution company. It should make it a lot easier for ISVs to see the benefit of Linux as their customers are now asking for it, moving from the edge of network technology into the enterprise.

They write to the LSB specification and they are able to run on every LSB-certified Linux distribution, which fortunately is every Linux distribution vendor out there today.

DR. HECTOR RUIZ: Maybe we can talk to Siebel. He is a U of I graduate, Midwestern honest guy.

RICHARD SEIBT: I fully agree how important the desktop is. The reason the desktop is coming up again is the very good experience all the last customers got using Linux on the server. Now, I would say the server is driving the desktop. This is probably different from the past when the desktop was driving the server. This has turned around and will give a new speed to further development on the desktop.

I believe that 64-bit capabilities with [AMD] Opteron will force ISPs to think about how they can use these capabilities on the desktop as well, because development cycles become shorter, the performance is much better and it now can be used on the desk as well. This will be the next move.

BORIS BIALEK: I would not stake that part completely to the desktop. The [AMD] Opteron system you have now is kick-butt when it comes to data centers. We are talking 64-bit database here. I am a database guy so in that space, when we look at what we run there, we run terabytes of database. We have upwards of 40 node clusters running. It was 1.5 terabyte, just for fun. We talk hundred terabyte sizes. The latest buzzword we came up with is "Lotabytes." That is millions of millions of something. I do not care how much, but it is big.

The key part here is the [AMD] Opteron plays very nicely there for us. 64-bit is a very important part of what drives the stuff. When Stacey says 2.6 kernel, that is all about exabyte file systems, new memory subsystems, new driver infrastructure and new components. That drives us nicely forward.

MARTY SEYER: I am a server guy so that is music to my ears. I want to comment on 64-bits on the desktop. Something the Linux community needs to think about is to imagine that in a couple years there will be 10- to 20-million 64-bit platforms out there on which you could deploy software on the client side. That changes the paradigm and the focus here. This community needs to plan on that. It is real. It is going to happen and we are going to make it happen.

JOHN 'MADDOG' HALL: One reason I got interested in Linux was that the open source nature allowed education and computer science research to be able to work with a real live operating system that was in major usage, yet also have full access to the source code and being able to share and collaborate. This was particularly so in the areas of large virtual memories and changing the algorithms themselves in order to give you much faster speed-ups. Not 10 percent or 15 percent, but 100 times or 200 times speed up. In previous tests we did in conjunction with Oracle, a processor with the same CPU speed, when they had a 32-bit database, would speed up 200 times in a six-way join with a 64-bit database on the same clock speed processor.

The speed enhancement was not the hardware per se. There was a change in algorithms they could use between a 32-bit address space and a 64-bit address space. This is a type of investigation and research I would love to see happen with a Linux operating system and open source.

DR. HECTOR RUIZ: Marty, I will underline what you said. We do expect, before year-end, to have literally a few million of the systems out there and next year tens of millions. The migration to 64 on the desktop is going to be a function of how well this community can support that from an application point of view.

Jon, beyond the technical merits, why is the Linux community important?

JON "MADDOG" HALL: The Linux community is important because Linux is the only operating system in major usage today that was started outside the United States. All the operating systems we use were either started within the United States or controlled within the United States, but the Linux kernel was started in Helsinki, Finland. Therefore, when you go to Germany, the Germans think of Linux as theirs. When you go to Spain, the Spanish think of Linux as theirs. If you go to France, the French think everything is theirs, but that is beside the point.

All these people are right because it belongs to all of them. It belongs to the community. When I see 14-year olds developing their own distributions that are

extremely successful; or a 13-year old setting up a whole trade show floor with all the networking and becoming systems administrator for the show; or a 7-year old arguing with systems engineers from a large computer organization on the best way to decompose algorithms to make it work well in a Beowulf, I have to think that this is the type of thing that we can do to allow our students to move forward all over the world. By the way, the 7-year old was right.

DR. HECTOR RUIZ: Do you all think that Linux will always be free?

SEVERAL: Yes.

BORIS BIALEK: That is the spirit that everybody wants.

Anything else would be a mistake. When your team came up and asked, "Do you want to work on the [AMD] Opteron platform?", there was no discussion which operating system it would be. The Linux team was crawling over the box in seconds. They had the stuff up and running in two days. They were pretty happy about it. Richard's team did a good job on the OS.

SCOTT MCNEIL: It is actually an interesting thing. The Free Standards Group is a community-based organization but we rely on corporate membership. Thank you very much, AMD is one of our corporate members. We rely on corporate membership to do a lot of the engineering. One of the core pieces that our members are looking to us for is to keep the standards free and based only on free and open source software. Like everyone else, they are looking to rid dependencies on any group or company for what is essentially a key piece of infrastructure and in many ways a depreciated technology.

Freedom in the software is not this weird esoteric thing that guys with big beards talk about. It is also important to corporations.

RICHARD SEIBT: I fully agree. I think that the open source community will be the driving force behind the next economic quantum leap to drive total cost of ownership. Compared to any other development, it is more innovative because there are no boundaries and nobody is controlling the innovation.

If I think about where the industry is today, I foresee that we will have a shift from investing in large projects more to projects based upon Linux driving the total cost of ownership. Partner standardization, clustering and these kinds of projects: this will drive the industry to the next quantum leap.

STACEY QUANDT: It is not just cost of ownership because it is often just line items. You need to look at the flexibility that Linux provides. The benefits, the costs and also the risks. You can mitigate your risk by using a Linux solution.

RICHARD SEIBT: I fully agree, but isn't flexibility part of total cost of ownership?

STACEY QUANDT: It depends what model you use.

DR. HECTOR RUIZ: What impact does 64-bit computing have to Linux? Are we collectively motivated to move the industry toward the 64-bit environment?

MARTY SEYER: I have a perspective on the high-performance market. If you look at that space, Linux on 32-bit has been quite successful there. There is a strong value proposition. If you look at adding 64-bit capability on top of that, it changes things dramatically. Several things motivate HPC customers. First, there is performance, then more performance, and then price performance. Adding 64-bit capability via x86-64 and 64-bit versions of Linux addresses those needs.

You have the larger memory address ability. The typical applications are rendering and simulation, which are memory hog applications, so they benefit directly. If you address the memory latencies between the CPU and memory, which x86-64 is doing, that improves the overall environment.

There are a number of things that the core hardware can do itself to improve the environment.

From an AMD perspective, with x86-64 enabling both 32-bit and 64-bit operations, you can offer up some interesting equations for ROI. You can start out with a 32-bit OS with a 32-bit app. You can upgrade the OS and still run 32-bit apps. At the customer's preferred pace, you can upgrade both the OS and the application to full 64. The dynamics on that and on the return on the investment are phenomenal to me.

BORIS BIALEK: That is the same thing for us. We run 32-bit and we still run the 32-bit database product. Customers buy the product today on existing 32-bit platforms and are happy with it. They migrate to new hardware platforms and do not want to re-engage and retest databases, so they can stick to 32-bit. If they want to move over, for us it is seamless. The customers bring the backup, bring the 64-bit app, and they are done.

64-bit is an absolute imminent factor to make Linux into the high-carat class enterprise systems. That is the right way to go.

DR. HECTOR RUIZ: There are some pretty smart guys at AMD, you know.

JON 'MADDOG' HALL: I would like to give people an idea of how large 64-bits really is. If you were to take and fill up a 1 gigabyte disc every second of the day every day of the year, you could do that for the next 586 years. The concept is being able to take any number of discs out there that are filled and simply map them into the virtual address base of my application without having to worry about whether I am at the end of my addressing capabilities.

That opens up whole new ways of doing programming.

DR. HECTOR RUIZ: Given that attractiveness, it seems like there are things we could do collectively to market this capability. Richard, do you have any comments on how we might be able to do that?

RICHARD SEIBT: I think it is a joint approach. As soon as possible, we have to do a "marketing program" to provide all ISPs, whether large or small, with the capabilities that exist today. Make sure that they know about the Linux development environment. If [AMD] Opteron is available, they should get an update with the capabilities of 64-bit, the SuSE Linux enterprise and the system development kits. Then they already know how to use it and what advantages this might have. We have done some education on it and everybody is ready and has thought about how to use 64-bit.

MARTY SEYER: To date, what has been done in terms of marketing 64-bit is that customers tend to view it as a penalty if you go the 64-bit route. We need to reposition that so you take your 32-bit infrastructure and migrate to 64-bit relatively painlessly and on your own schedule. We need to change the dynamics and the perception out there and obviously, we are going to do that this year. But 64-bit should be an easy advantage to achieve, not a penalty for getting what you need.

DR. HECTOR RUIZ: I want to thank you for being here.

I would like you to leave today remembering a few key points from this presentation. AMD and the Linux community are similar on many fundamental levels. We both value open competition as well as collaboration and know that by working together, underdogs

can make difference. Even maddogs can make a difference.

AMD is dedicated to the Linux community and to supporting true innovation – that is why we design processors that work for you rather than make you work for them.

We believe that our upcoming AMD Opteron and AMD Athlon 64 processors are going to be the best tools you could ask for to make the next leap in programming. Our 64-bit strategy is the only one with the business model to achieve a pervasive footprint in the industry. In fact, we are confident that our 64-bit processors will fundamentally change the economics of 64-bit computing.

This year, our processors will be supported by full, complete 64-bit Linux distributions, including those from Red Hat, SuSE Linux, and United Linux. Beta versions of x86-64 Linux are already available and in use today, and our x86-64 patches are already open sourced. Every one of you can go out and build a Linux kernel that supports the AMD Opteron today. And I hope you will! We are committed to releasing the details of our new technology to the public early so that we can reap the benefits of open software joint development.

The future is not built by those who react, but by those who act. Learning, science, technology, and especially business are not advanced by those who fear free trade or the open exchange of information. They are advanced by those who see free trade as the cornerstone of economic and business growth and the open exchange of ideas as the foundation of international understanding, cooperation, economic stability and a rising quality of life.

The Linux community is an amazing example of the power of collaborative effort to work for the benefit of all.

The success of Linux as an open system adapting to and partnering with the commercial computer industry shows how the needs of business and the free exchange of information can live in concert for the betterment of both.

AMD supports – now, and always – not just the growth of Linux, but the spirit of Linux as well.

We are committed to continuing to play well with others and to continuing our collaboration with you. We look forward to driving the next great wave of true innovation with each and every one of you.

I would like to thank Scott, Stacey, Richard, Marty, Boris and maddog for sharing their thoughts and ideas with us today. Thank you for your attention and have a great time at LinuxWorld!

Thank you.