Richard Heye

Spotlight
November 2003 • Vol.3 Issue 11

Richard Heye

Vice President/General Manager Of AMD

As vice president/general manager of the Microprocessor Business Unit for AMD, Heye oversees all marketing, platform engineering, infrastructure development, and program management functions for AMD's CPG (Computation Products Group).

CPU: How much of AMD is CPG and is CPG just microprocessors?

Richard Heye: AMD has three business units. CPG—which is supported by the microprocessor business unit, so from a business point of view, it's one and the same—is focused on x86 microprocessors for the desktop, server, workstation, and mobile market. We have a flash division, which sells flash memory, and we just spun that off into a whole new subsidiary. And we have another company, PCS. They are the ones currently who do extremely low-power MIPS microprocessors for handheld, and we just acquired Geode, which does x86 cores out of national. That's a brand-new acquisition, so we're just integrating that into the company right now. Far and away the lion's share of revenue comes from the microprocessor division.

CPU: From your job description (see above), you have a very full plate. How do these areas overlap, and are there advantages to having one person be responsible for all of them?

Richard Heye: Well, two things. One: What we actually have is a 'two in a box' structure, so I have a peer named Marty Seyer, and he and I are two in a box. Marty focuses most of his energies on the server side, and I focus more on the client, which would be desktop and mobile, and we kind of collaboratively work on the other groups. Given that my background is infrastructure and platform engineering—that's my area of expertise—and Marty is stronger on the marketing and sales side, so he'll focus a little more on that in terms of marketing programs and how we manage PR. So there are two of us, but from my perspective, the real strength in having it under one management umbrella is that you can really leverage all the strengths that we have. By that I mean, we can take full advantage of our partners, leverage that with our marketing teams to really make a strong story to our customers, and that's a very powerful thing. Because at the end of the day, AMD's strength really is leveraging the strength of our partners to make a complete system solution. To maximize organizational efficiency, having it under one management umbrella really allows maximum flexibility and maximum efficiency, as well.

CPU: What are the advantages of partnering with other companies as opposed to keeping all development in-house, and do you have a set formula or strategy for dealing with all your partners in the same fashion in coordinating product development and marketing?

Richard Heye: Let me go to 100,000 feet and then drill it down to specifics: The term we use at AMD—you may have heard it mentioned before—is the ‘virtual gorilla.' We refer that to the fact that, while we don't implement our own chipsets, motherboards, and all the necessary system components, by partnering we are able to build a very strong infrastructure to enable us to meet our customer needs.

So let's talk about how we do that. The first thing you have to do with every partner is you have to give them a legitimate business case why, working with AMD, they can ultimately be successful. Because no matter how much trust they have in you and how much they like you, at the end of the day, it's a business. And we have to continually show our partners that by working with AMD, they can make money. Given our success to date, it's a pretty good indication that we've been successful at not only making money ourselves, but also allowing our partners to make money by working with AMD. That's very, very critical.

What we do on a per-segment basis—because each one is different: server [segment] is different from workstations, which is different from desktop, which is different from mobile—and each one of those requires a different set of infrastructure—there is certainly some overlap, but also some differences—and what we need to do is a combination of marketing, define what's needed to meet the needs of our customers, and then the engineering team and the infrastructure team work with marketing to figure out technically what components we need. Because first marketing will give you the big picture and say ‘here's the capabilities they need' and you have to kind of break that down into specific chipsets and specific functions on chipsets, specific partners. We'll figure out which partners are best suited to meet the needs for end customers and then we'll work with them to ensure that they can technically be successful, that's what the engineering team does. At the same time, before they even go down that path, we have to show them that there is a business case where if they go down this engineering path and ultimately bring it to market, they can make money.

I think that what's often very critical in this is there is a lot of trust between AMD and its partners. One thing we don't do is we don't compete with them. One big difference is in motherboards as an example. When AMD started to go and design its own infrastructure, we explicitly made a conscious business decision not to get into the motherboard business. One of the reasons we decided not to do that is we wanted to be able to really work with our partners and they should never feel threatened. They should never feel like we could take away their business. Because we didn't want to compete with Taiwan, we wanted to work with them collaboratively. And quite frankly, that was a major selling point because we came and said ‘Listen, we don't want to compete with you. We want to make you successful.' And in order to do that we had to give them technical support and show them financially how they could make money. That's kind of an ongoing story.

In our chipset business, for example, we do make chipsets but, unlike Intel, when we do a chipset, any unique intellectual property that's needed to communicate, for example, to our microprocessor, we give that IP royalty-free to anyone who wants it. Not only do we give them the rights to that IP, we give them the actual verlog [a hardware description language used to design and document electronic systems] and we give them test patterns and we say ‘Listen, we would like you to make a chip for AMD and we'll technically help you.'

So where's our value-add in the chipset business? Our value-add is typically the interface to our microprocessor. To the extent we have value-add, we work with them. To the extent they have better designs than we do, we're not going to tell NVIDIA or ATI or VIA or any of those guys how to do a graphics engine. They do that better than we do. We will say, ‘If you want to talk to the microprocessor and you want to use the HyperTransport bus, we'll give you as much technical support as you need to be successful in order to bring your product to market as quickly as possible.' And that is not the Intel model.

So, by having collaboration, by having good business cases, we have a very robust infrastructure. Case in point is: We're announcing our eighth-generation microprocessor, Athlon 64, and we're going to have chipsets from NVIDIA, VIA, SiS, and ULI (formerly ALI) at launch. And that's pretty

darn good. We're going to also have a wide variety of motherboards from all the major motherboard vendors. So again, that's a launch statement.

The reason they're doing that is that they have faith that AMD is going to be able to bring Athlon 64 to market and we're going to be able to drive the industry. Because at the end of the day, if they can't sell motherboards or chipsets, they're not going to do it. So for them to actually invest those kinds of dollars to produce those parts and those motherboards actually shows that they have confidence in AMD's ability to a) build a microprocessor and b) actually drive the market.

The other thing that gives us, by the way, and I think it is a major strength for AMD is we have a very rich and robust infrastructure. Look at Athlon today, our Athlon XP products, then, my goodness, you can find motherboards for under $40 and over $100. And you can find chipsets ranging from extremely low cost, stripped-down features to incredibly powerful graphic engines and integrated graphics cores. If you're a customer and you have a specific end system in mind, wow, you name it, we got it. But again, the ‘we' is the AMD virtual gorilla because those aren't our chipsets and they're not our motherboards, but they work with our microprocessor. I think it's great.

CPU: What infrastructure was in place before you arrived and how did you go about determining what was needed and when?

Richard Heye: Well, when I arrived, there was no infrastructure. Or I can phrase it differently: When I arrived, there was a beautiful infrastructure; they didn't even need me. Because when I arrived, they were just announcing K6, and K6 along with all the previous microprocessors that AMD had built and brought to market were all Intel-compatible interfaces. So I could take a K6, buy any motherboard in the world that worked with an Intel processor, [and] plug it in and it just worked. So to some extent for K6, I was just incremental headcount.

Now when Athlon came onboard, that was a challenge, because that was the first time in the history of AMD in the microprocessor division where the interface was no longer compatible with Intel's. So my challenge was to deliver an infrastructure to support that new microprocessor, the Athlon. So the first thing we did, quite frankly, was we did our own chipset because at that time we had no credibility. VIA, ALI, and SiS were the three major chipset vendors for that timeframe and there's no way they were going to go and embark on a brand-new chipset for AMD with no track record. That just wasn't in the cards. They were talking to us but we couldn't count on them to have a chipset in place to actually debug and bring our microprocessor to market. So we had to have our own chipset early on in the process just to enable us to test our own microprocessor. So we did our own chipset and we did our own reference design. We designed internally a standard Taiwanese-class motherboard that worked with the microprocessor.

The first challenge we faced was going to Taiwan and saying ‘Listen, we would like you to take this reference design, do what you do best—make necessary modifications to meet your specific needs—and ship that board to work with AMD Athlon.' It took a lot of trips for me personally to Taiwan to make that happen. Or I really should say ‘we'; it's really quite a group effort. In parallel we were building an organization in Austin, Texas, to be able to actually have the technical strength to support an infrastructure. Being in parallel, it was very clear that we needed to have a Taiwanese lab as well and we didn't have one at the time, so we had to build a lab in Taiwan.

Initially, all this stuff was going on in parallel. We talked to the motherboard vendors, and [their] initial response was ‘We'll think about it.' What you do is you present a business case, in our case we actually had to show them a working motherboard and say ‘Hey, this is for real; we're not making this up. We have a technically viable part and it actually works really well.' Athlon was a fine, fine part. To the motherboard vendor's credit, they actually said ‘OK, we'll do one motherboard for you.' It was sort of a test case from their point of view. When I say ‘give them credit' you have to understand it was not to their advantage to do an AMD motherboard, because they already had a whole line of Intel motherboards—that was the major part of the market—and obviously they didn't want to gratuitously annoy Intel because you don't want to annoy a major

vendor, and yet they actually did it. And two things happened: They started making money on it and they started growing our market share. Because Intel does their own motherboards and we don't, they were able to grow their share pretty quick. The wonderful thing about infrastructure in business is if there's a market and you can meet the market needs, you can make money. And that's what's happening.

Over time, you start building up momentum. Initially, it was really hard to kick start that momentum in the beginning. I can remember literally week in and week out tracking exactly the number of motherboards that were produced in Taiwan, down to the single digit, and just tracking it and getting that infrastructure ramping up and ramping up and ramping up. Now we don't do that. We track it, but we track it at a much higher level because it's pretty self-sustaining right now. In the old days it was trying to get one motherboard. Now some of these top vendors have three, four, five, six motherboards for AMD in the works, using different chipsets, going after different targeted markets and segments. So initially it was very hard; in the end it worked out really well.

The other key thing I learned early on is that you always have to be honest to get their trust. In any big engineering projects, you have good days and bad days. On the bad days, you just tell them ‘Hey, we got these problems but we're working through them. Stay with us and as soon as we fix the problems, we'll pass it on to you and we'll keep going.' I think that sort of really open relationship with third-party vendors actually got us a lot of respect. Over the years, we've been very up front in terms of what our requirements are or the status of our product, and it's worked very well. The other thing we've done over time is we've built this fairly large Taiwanese lab where now we have a lot of support for the motherboard vendors locally. Because the reality is that if you are producing a motherboard in Taiwan, you want pretty quick access to technical support. You don't want to wait 16 time zones later to call Austin, Texas, you want to be able to just pick up the phone and talk to someone in your own language and in your own time zone. That's what we did, and that's helped quite a bit too.

Another interesting challenge which wasn't uniquely obvious is [that] AMD has a really good sales force. The problem was that a lot of these motherboard guys don't need a sales guy. If they're buying a third-party chipset, why would a salesperson talk to them? They won't get a commission out of that motherboard; there's no AMD product on there. So what we did to ensure continuity [is] we have account managers in Taiwan, and their job is to work with the major motherboard vendors and chipset vendors. And again, it's working really well.

The proof in all this is that in the history since Athlon shipped, you can search for all the stories you want, and you won't find a story that says the AMD infrastructure is melting down, [that] it has quality problems. We had availability the first two or three quarters of the Athlon ramp because we were growing and we had fits and starts for a while, but once we got over that knothole of figuring out how to work with Taiwan, set up processes and procedures, got the lab in place, it's been working really well. I'm not saying it's easy, but it went from an art to a science. We built the account managers, we built the right technical teams, we understood their requirements, they understand us better and it's like any partnership, as you mature over time it gets better and better and more and more efficient. That's were we are on infrastructure right now; it's running pretty good.

CPU: As you were building this, how far ahead did you look in terms of relationships and allocating resources? Was it one product at a time or were you looking to build long-term teams and relationships?

Richard Heye: From a relationship point of view, you really have to be long term. When working with Taiwan (and I think actually that's another reason why we have their respect), we've always said ‘This is a long-term relationship. There are going to be days and weeks where you don't think what AMD's doing is the right thing to do.' Sometimes you do have to make a hard decision to work with one vendor over another, but when it's really clear that we are really there for the long haul. Through thick and thin, we're going to be working with them. That it's not just a one-product deal or a one-motherboard deal or a certain product cycle on the microprocessors but starting with Athlon, then Athlon XP, now to Athlon 64 and Athlon FX, and eventually we'll do a processor

called K9 and K10 and K11, as we go out generations through 2010 whatever, the expectation from our partners is that we are going to keep on showing them our roadmaps and they are going to be working with us. And that's my expectation, as well. So, on the relationship side, it's always, always, always been long term in my point of view. In fact, one thing I think is also useful is that AMD, and myself personally, I'm pretty tenacious. When you build these things, pretty much from scratch, you just really have to show that you're always going to be there. You may not get what you want on the first round, but that's OK. [You'll] probably get it on the second or third time, but you just keep on being there, show them the business model, show them the technical support, and keep working with them and eventually, it all works out. So certainly it's a long-term thing

Now getting down more to a pragmatic thing, when you talk about motherboards, for example, [they] are very tactical, and you are looking at a six-month horizon. When you deal in chipsets, by definition—we're talking to chipset vendors now about features that won't show up until 2005. So when you talk about anything based in silicon, you better have a two-, three-, four-year horizon, otherwise you'll miss the boat on some stuff. Motherboards tend to be a little more tactical because the design cycles are shorter, the validation cycles are shorter, and quite frankly, the lifespan of the motherboard itself is shorter. So, on a per-product basis, it depends on the specifics.

CPU: How do you go about allocating engineering resources to your partners?

Richard Heye: It depends on whether we're talking about chips or boards. What we have found is that some partners really want a lot of technical support in that they really view AMD as sort of their engineering arm, and we'll provide that to them. There are some companies that say ‘You know, we're better engineers than you [are]. Just give us the reference design. We'll take what we want, but in reality we're going to do it ourselves.' And quite frankly, that's great. Because I'll be honest with you, there are some motherboard vendors out there that are just unbelievably good. One of the myths that people had in the United States was that if it wasn't designed in the U.S., it wasn't going to be that good, and that's just not true. Some of the best designs I've ever seen have come out of Taiwan, not out of the U.S. So if they do it themselves, that's great.

What we do is we are the backstop. We do validation testing. A way to think about it is that there are three phases: design, implementation, and validation. The design side really varies [depending] on the partner. On the implementation side, that is always their business because they do the manufacturing, so how they implement is pretty much their business. We will do design reviews of both their schematic and layout as a board. We will do design reviews of some of their I/O cells on a chip as it deals with our microprocessor. On the validation side, we test everything. What we have found is that no matter how good any design is, be it AMD's or a third party's, the first design will have problems with it. I've never seen a design work perfectly the first time.

So we have extensive validation labs because the idea is that we need to have a high-quality infrastructure out there. One way we help ensure that is we do very extensive testing on chipsets and then a fairly thorough testing on the motherboards. We will spend a lot of time with one motherboard and really ring out the chipset. Once we know the chipset works, then we can do less thorough testing to just focus on motherboard issues rather than chipset issues. It's just more efficient to do it that way. We test every motherboard and put it on a Web page, and that's worked extremely well. One of the interesting comments I'll make, though I won't mention any names: The better the motherboard vendor, the more they like the fact that we test their parts. From their point of view, it's to their advantage to have high-quality motherboards out there with us. So they view us as a great partner and they go ‘Great, we do a lot of testing internally but AMD, the fact you're willing to test our board, great! Any bugs you find we want to hear about it as soon as you find them.' The motherboard guys that actually have lower quality don't like the fact that we test them because they'd just as soon get something out that quick, and they'll worry about quality later. That's clearly not our approach. Certainly we work with them, as well. We find bugs, we recommend fixes. They'll fix them and eventually get on the Web page and it's a great motherboard. So in the end, there will be a continuous supply of high-quality motherboards, but it's funny how different partners react to the same requirements.

One thing you learn pretty quickly with partners is that every company is almost like a human being in that every human being is different. Even though we talk about a generic engineering person or marketing person or managing person or PR person, in reality, each individual has their own quirks and idiosyncrasies. That's the same with every company. I think what makes us so strong is that our account managers respond to the various companies' idiosyncrasies. It's not a good-bad thing; it's just a different thing. [It's] the key, just like when dealing with customers, [and] it's no different with partners. Every partner worries about something a little bit different and you need to address their respective concerns. In the beginning, we were just learning that. In the beginning, we didn't pick up on all the idiosyncrasies, but now we do. It was a learning process through Athlon, and I think we are in great shape. The other observation I make is that we could never have done Athlon 64 if we hadn't done Athlon first. By that I mean, we are driving a huge new software infrastructure with Athlon 64, and there is no way we could do a new hardware and a software infrastructure at the same time and learn it. The fact that we've done quite a good job working with our partners to get the hardware infrastructure in place for Athlon 64 enables us now to focus a lot of energy on the software infrastructure, which is a very new thing for AMD, for 64-bits on the client [side].

CPU: How important is seeding software developers with 64-bit hardware, and how does that impact your relationships?

Richard Heye: AMD historically hasn't really worked all that much with [software] developers. With K6, we introduced the 3DNow! instruction set so we had some [limited] experience working with folks writing graphics drivers and optimizing some inner loops for graphics-centric software. Very limited. Our interaction with Microsoft was pretty much working with them to tune DirectX for our 3DNow! instruction set. When we actually did Athlon, [it] added virtually no instructions, so we were kind of spoiled in that Athlon was such a good processor, anything that ran on 32-bits ran better on Athlon so we didn't have to do any porting or changing of software because it just worked really well. So then we come along and say ‘We have an idea, let's do 64-bits.' And that changes the whole situation because now you need operating systems, developer environments, the whole works.

The first step is to make sure you have all the tools in place because you can't write a line of code without compilers, assemblers, linkers, loaders, and all those constituent components. So we worked with the open-source community to get a good Linux developer environment going, and we worked with Microsoft to get a Windows developer environment going. Given all the success stories you've read in the press to date on how quickly some of these porting efforts have gone, it's a testament to two things. One is that our developer suite is actually pretty darn good. It also is a testament to the fact that our entire 64-bit strategy is working. One of the reasons why these ports go so quickly is that you find, especially with these really big applications like DB2, that there is a lot of ancillary software that they've built around it to build up the development environment (utilities, special libraries, scripts, all that kind of stuff) generated over the years to build these huge software programs. What is nice is that you don't have to port those over to 64-bits. They just run great on 32-bits. So the whole story just works really well. When DB2 was bringing stuff up, a lot of these utilities just work in 32-bit mode; they didn't have to port them, they didn't have to emulate them, and that really helped. So the first step was getting the tools.

We worked with the Linux community, initially on SuSe but now working with all the players including Red Hat, to develop 64-bit Linux operating systems, and certainly Microsoft has been extremely supportive in the development of a 64-bit OS for Athlon 64 and Opteron. So we had the operating systems.

Then it was kind of funny. It's almost the same story again with software as it was with the hardware. Imagine going to Oracle or DB2 the very first time, and you're AMD. You have zero history in servers, you're a chipset vendor, you have no major OEM support whatsoever, and you go and say ‘Hey, we think you ought to port your database to Opteron.' (Back then we'd use the codename Hammer.) As usual, it requires multiple visits. First you have to show them technically why you think it's a good idea. Then you have to show them working hardware and you have to prove that the thing really is real and that the scalability and the performance that we said on

paper would happen actually did happen. Then they start seeing the tools show up and the OSes show up and start seeing some momentum. But I still see it as a courageous move for the Oracles and the DB2s to move with us when they did because at the time, we had zero market. We weren't shipping anything. I think that speaks to the strength of our strategy that they were able to work with us, and I also think it speaks to their strength that they are willing to take some risks. I personally think those risks are going to pay off, but you work with partners and there is a lot of give and take, and these guys, they're going out on a limb. Five years from now we're all going to say it was the obvious thing to do; of course they were going to do it. But for the time? Opteron could have been a complete bust from their point of view because we weren't in the market.

If you go there with a solid strategy, you have to have the technical wherewithal to deliver. You can't just show them white papers. You've got to show them real hardware, but once you can establish the fact that you have a credible story, you have credible hardware and you paint a strategy of how you are going to start getting market share, even though you don't have any [market share], you start winning partners over. Then you start to snowball. Once you get a few big guys, you get some more big guys and it starts moving from there.

Same thing on the client side. A lot of people are thinking ‘Well, why the heck do you need 64-bits on the client?'. But we're finding that we have some game guys right now starting to do 64-bit games. We have digital content creation folks that have some stuff [they want to do] in terms of transcoding. One of my strengths is that I'm not a detail guy, I'm more a big picture guy. I've started to get a sense of how the infrastructure and ecosystem is going, and nine months ago if you had asked me how the client side on 64-bits was going, I'd say ‘Hmmm, I'm not sure. We're talking to people and they're listening.' But the way I judge success initially is: Are they spending engineering dollars for the port? Because that's when I know they are committed. So nine months ago I would have said ‘Wait and see.' Now it's [different]; if you call around to your friends, there is a buzz out there and people are starting to allocate engineering dollars and they are starting to look and evaluate. I'm not trying to tell you that tomorrow or by September 23rd [AMD's Athlon 64 launch date], we are going to have a bunch of apps out there. But I am telling you that by September 23rd, you are going to have a lot of people seriously interested in doing some preliminary porting and playing around with it, and that is the necessary first step to get you to the 64-bit client story, which is going to be pretty strong in 2004, and in 2005 it will really start taking off.

With infrastructure, it is [all about] tenacity: If you stick to it, if you believe in your strategy, if you believe in your story, you'll ultimately get the guys onboard. If you're right. If you're wrong from our point of view and you lose, well then the partners were smarter than you were and they didn't choose to work with you. But I will tell you right now, and I won't mention names, but there was one ISV that we just could not get them onto the Opteron bandwagon and our 64-bit bandwagon. Two weeks ago, they called us up, ‘We want to be there at launch with you.' And it's not like all of a sudden, our comment had just found the magic sauce. What happened was that they were looking around and seeing that their competitors were doing ports for Athlon 64 and more importantly, they must have perceived that their customers would buy those products, which means that they would lose a sale. So what that one company determined was the fact that they were looking around and seeing their competitors were getting on the AMD bandwagon.

At a certain point, the ecosystem picks up momentum on its own. My job is that I'm the guy who drops the snowball on top of the mountain. I have to give it a little push to get it going, but once it gets going down that mountain, then AMD just has to massage it, help it, feed it, nurture it, but it really takes on a life of its own. When that happens, it is just awesome. That's really exciting because then you get some real innovation going on that you didn't even think about, especially on the software side. I really do think, maybe [in] 2004 but certainly [in] 2005, we are going to see some really cool applications out there that you and I just never even think about. That is when it really gets exciting. That's when the ecosystem just kind of takes hold. We've built the base technology, not the end solutions, so it's really gratifying to see that this baseline technology really did enable a whole new set of applications that just couldn't be done on 32-bit machines. It's going to be great. Tenacity and time. When articles are written in October of this year going ‘Hey, there aren't 64-bit applications out there. AMD's not delivering the stuff. It's not working.' I'll just smile

and I'll just tape it to my wall. We're not going to change a bit on our strategy. We are going to keep on going, we're going to keep on plugging away and we'll turn the industry a lot sooner than you think.

CPU: At what point was it determined that the launch strategy for Athlon 64 would be "32-bit functional and 64-bit ready" and how much of that is AMD rolling with the atmosphere?

Richard Heye: Actually, that's not rolling with the atmosphere; that actually was planned. Because you couldn't get a 64-bit infrastructure without shipping 64-bit capable machines into the market. It's kind of a catch-22. Let's say we said we were going to wait until we get a whole bunch of 64-bit applications and wait for the OS, wait for all these applications. Well the application guys are going to be sitting around going ‘Are you really going to ship? What's the story?' And eventually you'd get this really nasty, negative feedback where people are going ‘Well, AMD's not shipping. I'm not going to write my app.' And it'd spiral in a negative way. If you wait for these applications, it could take awhile.

The part of the strategy that I think is really powerful is that we're going to go out with a strong 32-bit machine. Because we have a strong 32-bit machine, we'll get that installed base, we'll get the momentum going of driving a lot of 64-bit capable machines out there. And a lot of those machines will actually be developer machines. I think people are going to buy them as developers. I think there are going to be a lot of Unix developers buying these PCs. I think there will be a lot of PC guys buying these PCs to develop 64-bit software. So the way you phrased your question, it kind of sounds like we capitulated and said ‘Well, we can't really do 64-bit infrastructure so let's do 32-bits now.' And, seriously, that was never the thought process. The thought process has always been from day one ‘Let's get a really nice 32-bit machine out there. Let's really start driving a large 64-bit-capable infrastructure out there, and in parallel, let's get the ecosystem going on software.'

There is no better motivator for software than hardware. Once you start working with ISVs, you can never satiate their need for more and more hardware. Quite frankly, shipping this stuff and saying ‘If you want to buy 50 or 100 machines, go down to a CompUSA and buy a hundred 64-bit-capable machines.' That's really powerful. What it does is solidifies your message. So that was actually day-one strategy.

When you want to bring a brand-new capability to market, you're always going to have a chicken-and-egg problem, which is that the technology you bring to market won't have the installed base of applications and solutions. So someone has to blink. Someone has to say ‘Now is the time to start this ball rolling.' And if I think that 64-bits is going to be really important to clients in 2004 and 2005, man, I better get it out in 2003. You've got to start the ball at some point. So someone ultimately has to blink and say this is important, and I'm willing to bank my business on it. What AMD was able to do is [while] we're banking our business on 64-bits to some extent, at the same time our short-term business proposition is not 64-bits, it's 32-bits, and that's our strategy. We are saying ‘We're going to blink. We're going to go first. We're going to put our foot in the water and say this is important.' The journey of a thousand miles starts with one step. September 23rd is our first step. We've got one thousand miles less one step to go.

The major take away in all of this is that AMD's strength in large portion is derived from the strength of our partners. At the end of the day, AMD delivers great technology and great microprocessors to the market. But in the end, people don't buy microprocessors; they buy systems and solutions. And that's all due to the success of our partnerships. To the extent that myself and my team are able to deliver that, it's kind of a fun thing to do.

Enter A Subject (key words or a phrase):

ALL Words (‘digital’ AND ‘photography’)

ANY Words (‘digital’ OR ‘photography’)

Exact Match ('digital photography'- all words MUST appear together)

Search

Return to Previous Page