
# Oral History of Pat Gelsinger, part 2

Interviewed by:
David C. Brock
Doug Fairbairn

Recorded March 21, 2019

Mountain View, CA

CHM Reference number: X8899.2019

**Fairbairn:** So the thing I'd like to start out with is just get your observations on Intel's ability to identify future trends and really capitalize on them before they had become obvious movements within the industry versus being able to jump on a bandwagon that has already started and out executing somebody who may have already established a lead there. So I'd be interested in your observations as to whether that was an accurate perception and to the extent that it was. What within the Intel culture or whatever might contribute to that?

**Gelsinger:** Yeah, I'll give three different observations on that. One is I actually give Intel a little bit more credit on the PC situation in the sense that, the Crush campaign as it was described at the time was go get design wins everywhere. We know we're not going be able to necessarily predict the future, so plant as many seeds as possible, right, and one of them just happened to be the PC, , the IBM design win. But maybe all these other designs actually came to fruition and became embedded in this and other things associated with it. So I think there was actually a deep realization that we're on to something big and we're not going to be able to predict necessarily what circuitous route it may or not play in the marketplace. So I actually give them some credit in that regard: That they sort of realized that to be the case. So I'll say, first observation, is that some of that-- that's the way things go. You know, could you predict the impact of an open computer architectural ecosystem at this radically different price point in the era of mini computers and so on? I think that was actually a pretty thoughtful way to go about market creation. The second observation I'll make is market creation is just bloody hard and being able to predict the future, and how many companies over time have successfully created major new industries? Very few, right? If you add those up over time-- obviously we give Apple credit, being able to under Steve's leadership create industry defining products but the ability for any one company to have multiple industry defining, creating product categories is very rare. And, in fact, one of those things as we'll get onto from the Intel story to the VMware story a little bit, but some of what I'm proud of at VMware is we're actually being able to successfully do organic as well as inorganic innovation which leads me to the third observation is that Intel's been bad at inorganic. Its ability to acquire disruptive technologies and be able to accelerate and bring them forward is actually something that the very strong culture of Intel—Boom, this is how you do it. This is how we do innovation. This is how we scale and manufacture, and so on, actually creates a whole type of the antibodies to bringing in these other disruptive forces. Because one of the

  

magics about Silicon Valley has always been, any idea has ten different startup companies that are trying to do it in many different ways and timing and so on, bringing that forward in the right way. So one of the things that I'd say -- and I learned this in the painful way of not being successful at Intel -- is how do you go bring a creative startup company and be able to curate it, without crushing it inside of a big major company? And that's something that Intel wasn't successful at, and I think that limited some of that ability to do it because some ideas you're gonna have but, if you invest in one or two teams to go do an idea you know that in an industry there's ten or twenty, and to be able to predict that you have the best one who has just the right combination of go to market, core technology, gets the product just right. You're usually betting against the odds so you really need to be able to do both organic and inorganic and Intel's strong culture also creates many of the antibodies that prevented them from being successful at inorganic acquisitions and making those successful as well. And as I participated in a number of that those experiences as well, I sort of learned many things about how not to do it also in that regard and in a very powerful culture like Intel it's really going against the cultural model that they have in place.

**Fairbairn:** So it's, you know, the problem you described in terms of doing new things within a well-established company is a problem that most, the vast majority of large companies have. So can you describe just one of those experiences and what you learned? Could it have been successful or was there just no way that, that was ever gonna work or could you have done it a better way and been more successful?

**Gelsinger:** Yeah, I remember my first acquisition that I was personally responsible for at Intel was Chips and Technologies, right when we acquired that. Being part of the chip-set company and I personally oversaw that acquisition, bringing it in, making it part of the Intel chip-set business at the time and it was a 2D graphics capability and in particular was what we were looking for. 3D was starting to emerge but 2D graphics, we didn't have that and, I'll say, personally looking back on it versus what I know now, if there are ten things to do about making an inorganic acquisition successful, I think I broke nine out of the ten rules. There just wasn't-- I mean, how do you make sure you keep that team well separated? You know, how do you make sure you have the leadership? How do you allow their vision to fit into yours? Forces many of those organizational barriers right off the table, minimize the G&A [general and administrative

  

expenses] burden that you're bringing on with the new entity. Many of those kind of things, if you're gonna bring the successful company into this larger structure that, if you crush the spirit, the vision, the innovation engine that you have inside of it, you don't get anything on the other side of the integration process. And you're seeing that play out on a number of occasions. It's sort of like wow, we can't do that. And like Intel with some of the communication companies: Boom, let's get them onto our process technology. Oh, my gosh. Now you're essentially telling them the reengineer onto the Intel process technology because part of the synergy case was get the fab synergy, right. Get the cost of wafers. Get that thin, now you're going to from essentially something that was culture for a communication process technology and forcing it onto a digital process. Hum, it doesn't quite fit. We don't have the libraries. A lot of reengineering. So immediately you're sort of putting yourself off the curve of that piece of the industry and, so boy, essentially you sort of miss, you quickly miss a generation. You miss a Wi-Fi generation. You miss one of the process improvement cycles that they might have gotten if they remained on TSMC but, oh, we can't be on TSMC. We're Intel, the mighty process company. So in every one of those hurdles that you put one of those companies through creates this barrier for success inside of a mighty powerful culture driven company like Intel. I've learned a lot in that process. Unfortunately, many of those weren't all that successful for Intel and the communication areas, one in particular that I think Intel threw itself at harder than most and for the most part it wasn't successful over a long period of time and billions of dollars invested in many of those acquisitions.

**Brock:** Okay, may I ask a follow-up question to that one? With these kind of characterizations of Intel I always think there's such an easy counter argument that could be made and I wanted to see if you think it's baloney or if it holds water, which is that you can just flip the whole thing around and say "Focus is important." By focusing on the Intel architecture, you know, in microprocessors of all varieties and with a kind of synergistic memory business, what you see is record revenue and profit over and over and over again to the present time. So in some ways that if you're looking at it in terms of revenue and profit Intel is following a highly successful strategy.

**Gelsinger:** Yeah and I think by the way, yeah, there's value to that and I think that's almost a cultural statement. Why have they succeeded? The maniac-ness of the Intel team, this execution

  

focus, this alignment. We sort of joked at different phases of Intel that it was the largest single cell organism on the planet. <laughs> Some aspects of that make it powerfully good, but then why are you spending billions of dollars on acquisitions? Why are you wasting your balance sheet? Obviously you've decided, as a company and as a board, that we're gonna pour billions of dollars into that acquisition. How do you justify the combination of saying we're gonna expand into this adjacent area, yet we're gonna focus on our core that produces such extraordinary results? Those are obviously incongruous, to demonstrations of leadership, and I think in it, you either need to learn what you can and can't do, or if you are gonna put billions into these acquisitions and put balance sheet dollars at play, then you better decide the rules of making them successful as well. So I think in that respect you're absolutely right. Why was Intel so successful? That single cell organism, that extraordinary execution passion. Disciplined engineering, data driven culture. This harshness of the execution in some very tough business areas. But then if you're gonna say those are our core principles, and we're never gonna deviate from those, then don't waste your money on balance sheet on these acquisitions.

**Brock:** Fair enough. Well, great. When we finished our last interview session with you we just had you, you were just finishing describing your efforts in the first decade of the 2000s to really establish some of what we were just talking about, Intel's dominance in sort of servers, the world of servers after your time as Intel's first CTO and I wanted to sort of take you back into that time and if you could talk about as your work on the server business was kind of reaching its natural endpoint, what were you thinking and talking about doing next?

**Gelsinger:** Yeah and the key as we were finishing that era, we laid out the tic toc model. We laid out the multicore strategy, the integrated memory controller where we launched what today became the Nehalem, Westmere family of products. Boy, this was starting to turn around and as I look back on it Intel's market share, price per socket, and margin, the server business became larger as a result of those products after that period of time than it was before the whole Opteron phase which I sort of say, "Wow." That's a turnaround that you just say "Done!" It doesn't get better than that from a technical, business, architectural result and to a great degree, given the massiveness of the datacenter business that came out of that for Intel, is I sort of say "Does it get any better than that?" Now I think it's close to $20 billion in revenue for the datacenter business with extraordinary margins, market share close to a hundred percent. Score, baby! But the other

  

thing that we had turned our attention to and I was turning my attention to, and I say this was the last of my list of things I wanted to get done when I was the head of the enterprise business for Intel, was graphics and the massive multicore. We really saw that the space that Nvidia with GPGPUs, CUDA-- that whole space, and we sort of said "This matters." Their silicon footprint was always sort of this view: Who has more transistors on here? Okay, memory, they're sort of commodity. Networking. We gotta go with networking and we started to build up the Intel networking business. But that graphics footprint: People were starting to use that for non-graphics purposes with CUDA and GPGPU. We didn't think of it through machine learning and AI as we would today but those throughput workloads were getting bigger. So when we started the-- it became known as the Larrabee Project. It was sort of the last big project that I was getting underway at Intel, and I knew that if workloads emerged that weren't on the Intel architecture, Intel lost. That project was underway. We had two purposes in it and one was high performance computing and one was graphics. It was sort of the two workloads that we were working on and, again, if we looked at it today we would have said over five years ahead of our time, in terms of getting a machine learning AI workload in place. It really wasn't seen quite yet, but it was a class of that whole I'll say throughput-oriented versus latency-oriented workloads that were really driving it, and that was sort of the last big thing I had underway. And when EMC offered me the job of going there as their president, well, I struggled because I had made a list of ten things I wanted to get done when I took the enterprise job. Took over the microprocessor development engine for Intel. And the last one left was this graphics-throughput workload one, and I knew that wasn't done. So I was sort of like, I never don't finish the job. I was one of those. I really struggled over leaving at the time because that one was undone, but nine out of ten were done. I was being wooed to consider coming over to EMC and I knew that was really important. It wasn't done but it was a couple of, three more years to get it done. So I decided to leave at that point. That one was undone and Intel killed it shortly after my departure, and that was hard, disappointing to see and in retrospect, Nvidia would not be the company it is today had that been pursued because the workloads would have stayed on the Intel architecture.

**Fairbairn:** So you felt given the benefit of 20/20 hindsight you were on the right track at that point.

  

**Gelsinger:** Yeah because if you look back then Nvidia wasn't talking about machine learning and AI. They were talking about these throughput workloads, floating point performance, floating point ops, and all of that. Again some of these workloads were starting to emerge. It was more aimed towards HPC kind of workloads, et cetera. The graphics processor was getting generalized for the program we modeled. We were in the right space, and had Intel stayed with it the world of machine learning, AI, Nvidia would be a fourth the size they are today as a company because I think Intel really had a shot right in that space. I think that those new things-- and it sort of builds on the prior question as well. These are hard problems to predict the future and there's a lot of things going on at that point. It requires extraordinary commitment, top down support, determination, pounding through, not one rev of the technology, and boom it gets right, you know. Machine learning and AI, that whole space, we worked on it on the 80486 to add machine learning and, AI at the time. The age of Thinking Machines and LISP and PROLOG and all of that kind of stuff. We said we're gonna make the 486 a great AI chip. That was in 1986. I think of AI as a thirty year overnight success, where a lot of these technologies that curate over long periods of time, algorithms are being refined, and then all of a sudden the algorithms get good enough, the data gets big enough, and the overall compute capacity gets large enough, that oh, we can get interesting results and hey, you know, neural networks and, convoluted neural nets, those ideas were 20, 30 years old, but all of a sudden the algorithms that got good enough, the data big enough, and the overall compute capacity were instead of having crappy vision it's now better. Instead of having crappy chess, we're now better. We can outdo Go. All of a sudden, we start seeing these breakthroughs occur even though it took thirty years of building and refining those technologies. And that's exactly what happened to Nvidia and I give great credit to the Nvidia team, to Bill Dally, Jensen, et cetera, because they were on, hey, there's a class of workloads here that are really important and we're gonna keep hammering and focusing, you know, building the, software stack that allows it to occur, the whole CUDA architecture and they never deviated from that. They said we lost to Intel in the latency workloads, operating systems, et cetera, Intel's #1, but there's another class of workloads here that are really important and we're gonna win those. Intel gave up its shot to win those workloads at its today, I'd say its great demise. Even though you'd say, boy, that was an extra ten years of work. Yeah, it was ten years of work to go make that happen.

  

**Brock:** Do you think that it was in part that you weren't there to drive that forward that they decided to withdraw from that?

**Gelsinger:** Yeah, I'll say very directly that, boy you need that harsh, determined, passionate leader in that one and when they leave, hey, the next budget's come. We really can't afford that. Boy, it's a long time until it's profitable and oh no, we can't kill it. Yes, we can. No, we can't. I mean, these are hard tradeoffs in a big company and, budget cycles come and go. We all know the harsh cyclicality of the semiconductor industry. These aren't easy decisions at the time. It hurt me when it was killed soon after my departure but, I also realize in a big company those kind of things happen. You realize that to be the case and at the time, hey, the graphics program was struggling. It wasn't like it was an obvious; it wasn't a short chip shot onto the green at that point. There was a lot of work to be done yet and it was gonna be multiple generations to really refine it, building the software architecture to get it right, really needed a lot of top down leadership to go get it done so they decided to kill it after my departure. Disappointing to me and something in retrospect, we're looking back now almost ten years later and sort of say boy, the world could have been different.

**Brock:** Was it the case that people throughout your time at Intel, was it the case that people were approaching you with other ideas and opportunities as the folks from EMC?

**Gelsinger:** Yeah, of course, you have different opportunities, different job potential. As I view it, I've only really changed jobs once in my career. I'm a fairly loyal soul at that level and moving from EMC to VMware that was sort of an in the family move. I don't even consider that changing jobs. So essentially I'm a very boring guy. I changed jobs once in now 39 years of my career and to some degree that's sort of-- you get involved. You build relationships. You get passionate. You're always working on something today but also always starting the next project and the next project getting people rallied around it. You can't leave your team and your people. I remember the day. I'm sitting in one of my final ops reviews at Intel and I have the-- nobody in the room knows this, but I have the Intel offer to stay in hand and the EMC offer to leave in hand. My wife and I were about to go on sabbatical and I said I wasn't going to decide until I got back from sabbatical. I just needed time to clear my head. And I started looking around the room and it's like "Oh, there's Justin." I was at his son's wedding. He's coming to my daughter's

  

graduation. “Oh, there’s Steve.” Oh, he’s coming to my son’s graduation as well. All of a sudden tears start coming down my eyes. Because these-- after thirty years, these aren’t coworkers. These are people that you’ve done life with at that point and you get so “Where does work end and where does your personal begin?” After thirty years you can’t tell at that point so, it was a gruesomely hard decision. And obviously EMC was like “Hey, Pat, we want you to, you know, lead. We’re gonna coach you and help develop you to be our next CEO.” And Joe [Tucci – CEO] was extraordinarily committed to that, and board members reached out and it became a good life decision at the time but it was terribly hard. Because this was thirty years of life. I joke I went through puberty at Intel. I started when I was so young and I think if Andy Grove called me from the grave and said Pat, we need you back I think I’d still answer the call. These are just such powerful forces in your life that, boy, you just can’t imagine not being a part of their legacy.

**Brock:** It sounds like certainly a very big decision for you personally, and professionally obviously, but also from just your brief description of it, it sounds like EMC was making a big strategic decision to get you to join them. I wondered if you could just tell us a little bit about EMC as you had known it and why you think you became such a focus of their strategic thinking at that time? Why were they making such a big bet on you?

**Gelsinger:** Yeah, there were probably three or four things at play. One is they did sort of lose the head of products. Dave Donatelli, he had just left the company so they were looking for somebody who was a product guy. As a company EMC was known for its products and its sales but its product capacity had weakened over the years so it still had one of the most revered sales forces on the planet, but its product culture had diminished so they were looking for somebody who was product through and through, a strong technologist in that respect. I’d gotten to know Joe, the CEO there, because I was selling EMC to move from the Power[PC] architecture to the Intel architecture for their storage array, so we had gotten to know each other through that. I’ll say he had seen the “Passion of Pat.” “You’re gonna move your storage on to our architecture. Resistance is futile! We’ll bring you across the line.” So we’d gotten to know each other recently well. Obviously, he became reasonably impressed with me at the time to want to pursue me. Some of the board members were very aggressive. You know, the E in EMC was Dick Egan, right. His son, Jack, became pretty aggressive at pursuing me and some of the other board members as well. And my capacity to interact with the board at Intel was rather limited. Here

  

they were extremely aggressive at embracing me, particularly at that phase in my career where I said, “You know, boy, I really do aspire to be CEO some day.” And they became very committed that they’re going to develop you. There’s a bunch of things you don’t know yet, if you’re going be a CEO. They invested. One of them that was a lot of fun was they hired a personal corporate finance professor from Columbia to be my corporate finance tutor for a year. They clearly were investing in-- committing to help and build that skillset because being a product technology guy doesn’t give you the range of skills that you need. Then I had a little family pull, because our family was East Coast based and when I joined Intel I said, “I'm gonna be in the West Coast for two years.” Thirty years later, so they were super excited to see their black sheep coming home as well and we had a wonderful three years in Boston. We got down to my family in Pennsylvania. One of my sons and his wife came to Boston for university so they were in town and we would often get our family up, my family up or go down and see them so it was a really delightful three years. And then of course when I moved back to VMware in the West Coast it was like “Oh!”

**Brock:** You escaped again.

**Gelsinger:** Yeah, my wife was actually-- she was not happy about the move back to the Bay Area for VMware. I described I sort of took her clawing and scratching her across the nation right back to the Bay Area. She was enjoying our time on the East Coast.

**Brock:** From the outside looking at it, EMC I think at that time and certainly in more recent years with VMware and everything, it’s so identified with kind of like the datacenter and the Cloud and this sort of area so massive amounts of server chips involved. It seems like there’s a natural resonance there if they were getting on the Intel platform. With you comes a person who knows that exquisitely, so that is taken care of but it also seems like you must have been really living in that emerging, new kind of dominant area of computing, of the Cloud, of the world of datacenters while you were doing all this work at Intel. So was that part, did you see that as part of the connection of expertise and how you were going revive their product side of what they were doing?

  

**Gelsinger:** I think part of it was clearly leveraging and moving up the stack as I recall it going from silicon to systems and when you've done your twelfth microprocessor, number thirteen isn't quite as thrilling right in that sense. You're just not learning as much at that level and I didn't feel like Intel was really developing me to be that next leader. So some of it was leadership development, capacity as the skills for CEO. Some of it was moving up the stack. I know how to do chips. I know how to do systems as well and that was something that Intel did, systems as a hobby. EMC did it as their life blood and so learning that next level. Also, essentially Intel didn't do any sales. It did allocation, and that's a little bit crass but to a great degree hey, you know, when you're in the silicon business, if you build a hot product you're just telling customers this is how many you can have and we'll negotiate the price. Essentially half of Intel's business was seven customers. To get to half of EMC's business I think you had to get to almost 2000 customers. So it was learning a lot more about sales and what that really meant so there was huge interest in that regard. Clearly I'll say silicon wasn't-- I was very much sensing that… silicon wasn't thrilling as it was for my first couple of decades at Intel. It was really the center, the nexus of innovation was happening at silicon where you could tell it was moving up and away. That systems and software becoming much more the nexus of innovation and where these key inflection points were happening in the industry and clearly Cloud was starting to emerge but wasn't very clear yet what that really was quite yet. But this idea of systems and software becoming more critical, distributed systems becoming more critical, I sensed that the core of innovation was moving north and doing the thirteenth microprocessor of my career wasn't gonna be the thing that necessarily was those marks on the wall that I think were having the industry impact that the first dozen had or creating USB or Wi-Fi or PCI or some of the other key standards, impacts, that we were creating. I was excited to move on to a faster learning curve as I moved up into the stack. Clearly the leadership and skill development as well and still being the ambassador of Intel and the x86 into this whole other system segment as well.

**Fairbairn:** I presume in joining EMC you must have had, you know, a very positive view as to what their future might be. Did you have a clear idea of how to get there? I mean did you have a plan or a vision in mind as you joined or was that something that evolved once you got there?

**Gelsinger:** You know, somewhat. I was incredibly naïve in what it meant. Hey, hot shot silicon guy. I can go make storage systems a whole lot better. I had this sort of naïve view and I

  

remember one of the first meetings I was in we were having a roadmap discussion and mostly it was them teaching me but I remember in this one conversation we got into it and I said, "No, that's not right. That's not how it works" and literally before the meeting ended we were down to the schematic diagrams of the storage processors using the x86 and I was explaining how they should be using the processor in this cross-connected, redundant fashion and was-- everybody that left the meeting was sort of like "Oh, we really have a product guy now." <laughs> But there was a lot of naïve-ness in my view but that's also where you grow and learn. As I say, you grow and learn in times of change or challenge. When you're being successful, you don't learn much. You might be pomp and arrogant and be, "Oh, I'm good." Well, you're not really good. You're mostly lucky in doing things that others have done in success, as they say. "Success has many fathers and failure's an orphan." Going to an East Coast enterprise software-led systems company, that was radical. So it was a huge time of growth for me, and I was just really eating it up. Having to learn systems. "What's a block storage versus a file storage, what is that talking about? What is this fiber channel thing anyway?" You're learning all these new things, but you're gaining a wealth of experiences also into that environment many of which are highly leverage-able as well. So it was a thrilling time for me personally to learn and the combination I'll say of my Intel experiences, of which many of those interfaced with VMware, and my EMC experiences where I had lots of ambassador roles to VMware on the behalf of EMC, really prepared me to come back and I think have a pretty exceptional run as the CEO of VMware. As I sort of like to say, "God doesn't waste experiences. You just don't know yet how they're gonna be useful in many cases."

**Fairbairn:** So you mentioned one thing and that is what the East Coast culture whatever-- was there a significant difference in terms of East Coast versus West Coast or certainly maybe Intel versus EMC?

**Gelsinger:** Yeah. You know, there were just a few little examples of that. I remember one time I was meeting with one of the board members, Jack Egan. I asked him, and this is the son of the founder, you know, "Jack, how am I doing? What do I need to work on?" And he says, "We're an East Coast company. You need to dress like an East Coast company." I dressed like an engineer because that's what I was. It's like "Okay, honey [wife Linda], we're going to Nordstrom tonight." It's like people showed up for meetings with ties, and there weren't

  

customers there. What's the matter with you? <laughs> Some of that formality. Another example was I remember one of the first ops reviews I was sitting in and so on like that Intel was just extraordinarily direct. So this one little topic comes up and as we were discussing a little bit and I say, "So what other data do we need on this topic to make a decision?" Everybody sort of says, "No, I think we have all the data." I said, "Why don't we decide?" "Well, we wanted to take it offline and discuss it a little bit." I said, "Why? Right now. Let's get it done." And they sort of said, "Well, you know, we wanted to discuss some of the organizational implications?" I said, "Well, let's figure it right now. Let's decide. We got all the data. You just said we did. Let's do it." Everybody's like Godzilla has arrived <laughs>, and this directness, this bluntness, this harshness about data. Oh, we step on a few toes. Who cares, right? We're gonna go fast and get things done and, boy, it was very much, much more, "Well, we gotta consider the people, the roles and teams and organizations." Intel in many respects, I'll say, was a super successful culture but it was also harsh. It had a very hard edge and many times people were crushed in the process of that without good reason. You remember the famous Intel phrase, "constructive confrontation," right? Just we forgot sometimes the constructive piece. It was just a hardnosed confrontation and a lot of those soft skills to me were very much lost at Intel where this is what I grew up in. So, hey, I was ready to mix it up with the best of them. I remember some of the interactions with the teams in Israel who were known for their bluntness and directness. They sort of said, "Pat, you're more Israeli than the Israelis are" because that's what I grew up in. I remember I went through puberty here and then bringing that into a culture that was much more worried about some of the people implications. Again, they're different. It's not one is better than the other. They're very different and out of that I think I became much, much more thoughtful in people skills and teams and relationships and how you build that. And as you come now to a software company where some of those harsh skills— "Yeah, we need some of the discipline and rigors associated with it but some of the, you know, you can't piss off one of your key software engineers and expect them to stay. In a fab environment you can't-- what are you gonna do? You're gonna go get this done and the culture was so very different in that sense. The cultural aspects of that were radically different. The East Coast culture, more of a sales culture, more of a people-centric versus product-centric culture, but I've learned a lot in that process as well and say it really rounded my leadership skills quite significantly.

  

**Brock:** When you made the jump to EMC in 2009, it had, if my research serves me well, in 2004 it had acquired VMware and in 2006 it acquired RSA Security, in addition to a variety of others. So those are some big names that it was bringing in and big and important operations that it was kind of bringing in. How were things organized and orchestrated? Were VMware and RSA relatively sort of like autonomous divisions, or was it very integrated? Could you describe that scene?

**Gelsinger**: Yeah, yeah, and EMC, I think, was correctly saying, "Hey, we've got to go acquire other things because the storage industry by itself isn't going to be the future. They really took this, I'll say, this era of the internet buildout, the Four Horsemen of the internet, and EMC's role in that, the huge profitability of its storage arrays and saying, "Boy, we've got to go acquire and build out some other assets." And that was very much Joe Tucci and RSA, Documentum, some of the backup products, Avamar, it was a pretty very quickly going and building out a set of different portfolio of assets that largely were enterprise-centric, and could leverage this big enterprise sales capacity. But they were off to have to a broader data center infrastructure set of assets to play. Now, when it acquired VMware, Joe Tucci-- and I think this will probably go down as one of his greatest moves of all time-- he bought VMware for, I think, $625 million. I think as of today, we're just at about a $75 billion market cap. It's like, wow, right? You just-- this is beyond exceptional, and again, a whole lot of people, a whole lot of work to go accomplish that. It's not like boom, it just magically happened. But wow, at that level. In 2007, they had spun VMware out. They had done the public offering for it, so VMware was largely separate. It was run separately, et cetera, and again, I had the role of the ambassador to manage the EMC-VMware relationship. So I was responsible for essentially the portfolio role that VMware had, but it was largely run separately by the VMware leadership team.

**Brock:** I get it. And this was-- EMC had an 80 percent stake?

**Gelsinger**: Exactly, exactly. They had a majority stake but they ran it separately and the separate board structure that was there. Everything else worked for me. RSA worked for me. All of the storage products worked for me. Documentum and the other assets that they-- every other product asset worked for me except VMware. And the core business though was storage. Everything else was sort of on the periphery. So job one was get the storage business back to

  

health And we had companies like NetApp that were sort of chewing into the storage business with some different ideas and, as I would say, NetApp was gaining market share. EMC was flattening, and over the course of my time as the leader there, we reversed those. So I felt really good that we got EMC back on a growth trajectory. We had done a few acquisitions, built out the portfolio -- Data Domain, Isilon -- but also just got core innovation back into it. By the time I exited EMC, we had gained on the order of 12, 15 points of storage market share, and sort of blunted NetApp's growth. And so that was sort of good, yeah? That was job one for me from the board, was "Turn this around. We can't let them continue to grow up." So job done, which did make for a very, very curious moment. The first day on the job at VMware was VMworld, the big VMware conference. I took over. Paul Maritz, the CEO before me, announced me as the new CEO. I took the stage, gave the keynote for it, and that night was the NetApp customer and partner event and I'm on stage with Tom Georgens, the CEO of NetApp, because NetApp was a huge partner of VMware. So here is a big NetApp event, Pat with NetApp, Tom, we were on the pitcher's mound at PG&E Park where the event was, at the big customer and partner event. And that picture gets tweeted out to the bloodthirsty, raw-meat-eating EMC sales force and I am now getting thousands of emails – "traitor," "turncoat" -- coming my way since yesterday I was enemy number one for NetApp and today I'm partner number one for NetApp. So it was a rather humorous moment there seeing how your friends quickly...

**Brock:** Well, what would you say was your primary strategy in that turnaround of the storage business for EMC?

**Gelsinger**: Yeah. Part of it was we didn't have a good file strategy, so we acquired Isilon. Part of it was the product engine was sort of sputtering, so we had to bring innovation back into the core product machine of EMC. Too many products, somewhat scattered. They had acquired different product lines but hadn't really synthesized them together and then the other was Data Domain, the hot backup and archive solution and really exercising that very effectively. And by the way, we were commenting on my acquisitional learnings, and to me, that was sort of the first major acquisition at EMC that was under my leadership. So now I had all the bad learnings, some of EMC's better experiences on how they did it right, so Joe and the leadership team saying "No, no, do this, do that." So that combination of what not to do, some of their successes in Data Domain became a roaring success for EMC, extraordinarily successful, so how to make the sales

  

model, the innovation model, our work there, integrating that in an effective way. I really saw sort of my first test case for many of those learnings coming together into a very successful acquisition experience.

**Brock:** Was Veritas one of your competitors in that area?

**Gelsinger**: Yes. Yeah, yeah.

**Brock:** Okay. I did an oral history with Mark Leslie, so I understand a little bit of that scene.

**Gelsinger**: Yeah, yeah, so Veritas was Symantec, the NetBackup, then the NetWorker and some of the other companies there. NetApp was using its products as a store, as a backup device, as well. Data Domain, NetApp attempted to acquire them and EMC took it off the table, so that sort of built a little bit more of that grudge match, as well, between the two companies. There is a nice competitive dynamic there, and like I say, in my leadership time there, we were very successful to get EMC back on the right track.

**Brock:** And that's pretty quick, right? Three years or so you were at it, so...?

**Gelsinger**: Yeah, yeah. I certainly felt good that we increased the market share by about 10 to 15 points over that time period, increased profitability. We increased EMC's market cap by 60 percent over that 3 years, so I think all of the sudden the market started saying, “Oh, right, they're getting their mojo back,” as well, and again, with a very, very capable sales force. You start putting better products in their hands, good stuff happens, and as I said, a great learning experience for me, as well.

**Brock:** Well, tell us, if you would, about then the shift to VMware. What was your thinking? What was EMC's thinking as the major shareholder for VMware? How did that take place?

**Gelsinger**: Yeah, and it was really Paul Maritz had replaced Diane Greene as the CEO at VMware. And Diane was-- moved on, and there was clear tension because she was, here I'd call it seller's remorse instead of buyer's remorse. Soon after she sold the company to EMC, it's like “Boy, I really want all my independence but EMC now is the majority shareholder.” So there is

  

clear tension there. EMC did the public offering exactly like Joe committed, but there was clear tension. “Hmm, this is how we want you to operate in this model.” And very quickly, it sort was a parting of the ways after the public offering and then Paul Maritz was on the EMC team and Joe moved him to be the CEO, and that was from 2008 to 2012. And I'm leading the product efforts. Paul managed that and a founder transition is hard. If it goes well, it's hard, because Diane was the founder of VMware. Everybody there worked for her. Her husband Mendel was the CTO so he was the sort of the brains and so on, and they were a package. When Diane left, Mendel left, so you had the technical leadership as well as the CEO and inspirational leadership. Paul had an enormous task to stabilize the ship in that regard, but Paul was a visionary leader and VMware needed somebody who was more of the operational leader at that phase of its growth. And so it became obvious that “Hmm, somebody more of Pat's ilk fit but needed to still be very technical.” Paul, Joe, and I agreed also to form Pivotal at the time. We had some data assets, some developer assets, that weren't being, I'll say, well leveraged inside of the companies. We said we're going to form Pivotal and move Paul to be the CEO of that. We'll bring Pat over to be the CEO of VMware, so as I say, it was sort of in the family move and we'll make some product transitions on the EMC side and viola, I'm the CEO of VMware, moving over there and again, bringing some of that [Intel] organizational discipline, execution discipline, product and technology awareness, what are we going to focus on. And VMware's core product, the Hypervisor, the virtual machine, was flattening in the industry and one of the things I'm super proud of now is that VMware was less than $4 billion in revenue when I came in and growth rate was flattening. Now, we're just under $9 billion of revenue and growth rate is accelerating and for that it's like, “Okay, we got this machine humming.” We sort of grew well past our Act One, the virtual machine. We now have a rich portfolio of products, seeing our growth rate accelerate. It's been a great six years as CEO.

**Brock:** I guess, by this time, by, when you're thinking of making that move, 2012, the Cloud, if you will, is starting to become more visible, more public clouds, private clouds, and this almost idea of data centers almost becoming a commodity environment or, quite literally, a kind of a commodity environment. The turn from people's use of computers rather than being a physical server or a physical unit, now people are spinning up machines as they need them and things like

  

that. That's really starting to emerge at this time. I would just be interested to hear you talk about how that looked to you then, and how it's looking to you now.

**Gelsinger**: Yeah, and I think this whole notion of Cloud is such a profound notion because it essentially turns-- the virtual machine turned a server into a provider of workloads, where you didn't have to go configure hardware, right? Boom, you just spin up a virtual machine, and you didn't have to acquire it. You didn't have to provision it. Boom, it just sort of magically happened. Well, Cloud does that at scale for the entirety of an application, and you run it for me, and wow, that's powerful. And to a great degree, some people, a lot of the early positioning of Cloud was that it was cheaper. That was bogus, right? It was easier. I didn't have to go to that stupid, slow IT guy who is making me do some trouble ticket request for it. He may or may not get around to it for 90 days. So I can spin up my next application, I go to the Cloud. I put in my AMEX card. I have it after lunch. It was just easy. It was fast, and over time, it got better, as well. And in many cases, it was cost effective, right? So that easy button became the moniker of the Cloud and to some degree I would say the Cloud did to data centers what VMware did to servers. And they did that scale capacity with a business model that just changed the game entirely, and every Cloud uses virtual machines. So we're sort of an underlying ingredient, but VMware didn't put the whole solution together -- the APIs, the portals to make that easy at scale in a public cloud way -- because we were having so much success still fixing the compute servers in the data center. So to a great degree, I think, VMware had the opportunity to be the Amazon of the public cloud, IaaS models. It had sort of the right stuff and some of the right ideas, but we didn't go operationalize that business model like Andy Jassy and Amazon did. When I got to VMware, it was like “Boy, this cloud stuff is starting to take off.” This is 2012. We got to go, and we had a project internally that we turned into the VMware cloud and became vCloud Air. Boom, we're going to get that underway. So within about a year of me getting to it, we launched the vCloud Air product where we're going to compete with Amazon and Azure and so on. We went running out to go run our own Cloud and we're struggling to get that underway. We got a lot of other stuff going on. We don't have all the ingredients yet trying to build it, and as that's going on, as we all say, starting to ramp our own cloud. Amazon is just taking off. Their growth rate is accelerating in this period of time. They're already six, seven years ahead of us, and they're accelerating and that gap became larger and larger and as we're sort of sputtering to

  

get vCloud Air underway. We're just saying, “Hmm, right, are you going to be able to compete?” It takes a lot of capital to compete. We're a software company and plunking down data centers and building out global networks and all of those kind of things. Microsoft, they had the opportunity to leverage their investments in Office 365 and some of the other web properties, and they were already $60-$70 billion-dollar companies. So, boom, they could sort of slide an extra few billion of capital into it. I'm a $4 billion software company. I don't get to hide a couple of billion of capital to go build that out into our P&L, so all of a sudden people were questioning this. Can you guys compete? Are you going to pour the billions into it that's required to compete as Amazon has taken off and Azure has taken off. And Google? Well hey, search, they already have a global network. They have big data centers so they can go take-- how are you going to compete? And that led to what became the critical partnership where we exited the vCloud Air business and partnered with Amazon. And that was sort of the seminal shift in VMware strategy, where we became a fundamental partner with them, as I would say, bringing the leader in private cloud together with the leader in public cloud to deliver the hybrid cloud. And that's been an extraordinarily successful partnership that we forged over the last two and a half, three years since we announced that relationship. And now we're seeing great momentum for that. We've also partnered with IBM, Alibaba, have a wide set of people who use our software to run clouds now. It's really a software, SaaS, business model without owning the underlying hardware and capital for it. But this era of Cloud, if we think about it, if we just step back for a second, imagine you and I had a great idea for a new application 20 years ago, right? Oh, but it's, ‘Boy, it's going to take a lot of data centers. We're going to have to plop them around the world and build out a lot of capacity because this is going to be a great new app, big data requirements and so on.” That might have taken us two years to build that infrastructure, maybe $1 billion of capital. We have to-- oh, we have to put one in Singapore. We probably need to have one in Japan. We need one in Europe. We need one West Coast, East Coast, et cetera. Wow, $1 billion to go get that underway, and then we have to get the connectivity in place, et cetera, 2, 3, years. Today, I can do that this weekend with my AMEX card using the cloud, where I can have global resources made available at scale that I can begin to provision and run that. Years have become days, right? Billions have become millions to go get-- and that's the magic of cloud. I mean it just has so compressed and reduced both the cost and the time and the scale that you have available. We're literally now-- your AMEX card. Well, at least...

  

**Brock:** Somebody's AMEX card.

**Gelsinger**: Some-- Michael Dell, Michael Dell's AMEX card.

**Brock:** A better one than mine.

**Gelsinger**: Right, Michael Dell's AMEX card. Right. I can build the world's largest supercomputer, assemble it over a day or two and then decompose it over a day. It's just is so magical, the computing scale that now is available to any application developer on the planet, and I call this, that we're entering what I call the era of the superpowers. And the superpowers, we used to refer to them as nation-states. Now they're technologies, and I describe it, the four superpowers of today are cloud, mobility, AI, and IoT. Cloud, like we just described, unlimited scale. Literally, I can have any capacity of computing you want. How many million servers do you want? Extraordinary scale. Mobility, we're now over half of the planet is connected, right? I can reach billions of customers over that mobile network and on the mobile devices anywhere on the planet, unlimited reach. AI, I bring intelligence to everything. It's just I have the scale. I have the data. I have the algorithms. I can bring intelligence to everything. And IoT, you get to bridge the digital and physical world as never before. Wow, and each one is making each other better. I mean, it's this extraordinary era that we're in, and I'll say that's sort of the fun of being the CEO of VMware now. We're sitting in the nexus of many of those, and if we go back to earlier in the conversation, as I said, it felt the nexus of innovation was leaving the silicon layer, and moving up and now, to me, those four in conjunction, that's where the nexus of innovation is occurring today. Where cloud, mobility, AI, and IoT are really not just powerhouses by themselves, but they're causing each other to get faster. Because if I have big cloud, I can have more mobility connected. More mobility, I connect more data, which makes my AI better. I can then execute it over my IoT with more introspection and inflection points which gives me more use for my cloud. These become these reinforcing. And I believe that, as I like to describe it, today is the fastest day of tech innovation of your life, today is the slowest day of tech innovation of the rest of your life.

**Fairbairn:** Because it all expands.

  

**Brock:** And it's interesting that you were just saying a few minutes ago that initially maybe, when people were making the argument about cloud, that it would be cheaper, that that was not the case, that it was faster, but it was still plenty expensive.

**Gelsinger**: Fast and easy.

**Brock:** But now that has changed. Now it is just cheaper.

**Gelsinger**: And in many cases, a well-run private cloud, your data center, and I call it the three laws, why will people run, given you now have this unlimited scale of public cloud, why would you ever run a data center yourself? And I will assert the three laws: laws of physics, laws of economics, and laws of the land. Laws of physics, "Hey, if I need low latency, I can't be round-tripping to the Cloud and back." So if I need to respond in 50 milliseconds, if that's a shadow or a grandma on my driving, my street, I can't give a 100-millisecond round-trip to the Cloud. I need this, so latency. Laws of economics. If I'm doing a surveillance application, how many pictures of the cat do I need to send to the Cloud? You need intelligence locally. So maybe learning I do in the Cloud, but inference I got to do locally. I can't be-- and bandwidth still costs, right? Storage still costs, so laws of economics. But laws of the land, where certain applications, certain datasets, people will dictate them into certain proveniences of execution. And those reasons, as I would say, and with the emergence of edge, of IoT, I think we'll see a swing. And one of the other things that I've observed is, over the history of computing, we've been oscillating between centralized and decentralized and distributed over the history of computing. Cloud has been a force of centralization. I'll run your data centers for you. Well, I believe IoT will be a force of decentralization, where it will push more computing and capacity back to the edge and then maybe 5G and low latency will help push it back, and we're going to see that pendulum move back and forth. And again, it used to all just be a mainframe. And then we could distribute minicomputers and then minicomputers became centralized and we needed distributed PCs and one after the other has had different effects of centralization and decentralization. And clearly it's not one versus the other. And that's why we would argue that the future is the hybrid cloud, and this ability to operate between private resources and public resources and bind them together in flexible and interchangeable fashions where this workload, it's running pretty good. But I need lower latency for it. I can move it to the edge. The network is getting fast enough. I can move it

  

back to the Cloud. Oh, the law changed in Malaysia. I have to move it on-premises. Oh, they're now allowing me to centralize that in UK but Germany still requires it local. It's not one versus the other, and it's not going to be a static picture as technologies emerge, different laws, regulations occur.

**Fairbairn:** Let me ask you a question about, just for clarification on the laws of physics, the latency time, and so forth. You said the private cloud has the advantage. Is that because of the control of the workload, or it's certainly not location?

**Gelsinger**: Well, part of it is location, as well, right? I can push the, put the compute and storage capacity locally or closer to where the inference or the use of that is actually occurring. If the robotic arm needs 20-millisecond response time, hey, that mini data center is going to sit right in the factory.

**Fairbairn:** Okay, so you're talking about smaller scale things that you need locally versus...

**Gelsinger**: Could be, but even there, imagine an aircraft carrier. There is a mini data center on the front of the boat and the back of the boat. It only has satellite connectivity, so I got a fricking thin pipe. I mean, these are pretty meaningful compute capacity. Imagine an oil rig. There's-- and in an IoT world, that oil rig is going to become much more instrumented with dramatically more data and compute occurring there, and I got a teeny-tiny pipe back comparatively to the amount of telemetry, analytics, et cetera, that's occurring. So it really isn't just one or the other in that sense and latency, bandwidth, laws, regulatory requirements, one of my, one of our customers, one of the big bank CIOs, he says, "Every application that doesn't touch money or customers, I can move to the Cloud, and if I find one of those, I'll shoot the person working on it." Because he says, "Everything I do touches money or customers. If I'm doing anything else, tell me what it is. I want to get rid of that guy" in that sense. And they're viewing the regulatory requirements. They're so high on one of the cannot-fail financial institutions that, “Boy, can I trust the Cloud guy to do it?” And again, so some industries will be very slow to move to public cloud enablement. Maybe they do some test and dev there, but there are operations requirements, resilience requirements, regulatory. That same CIO once, he told me, he said, "Hey Pat, today is a great day. I only have two regulators here today." So many different regulators. It's similar in

  

health care, government applications. And even the mighty Amazon public cloud, they've built the local instance of the public cloud for government. That's what the FedRAMP GovCloud is for Amazon, right? It's essentially their own private cloud on the other side of essentially the government's firewall that's locally operated inside of the government by government-certified personnel, only government workloads, meeting all the regulatory requirements, so in that sense, it's not one versus the other. And that's why we say the answer is a hybrid future, because that really gives you the ability. Maybe this app is running. It's tied to some physical IP addresses. It uses this physical store that I have. It still is tethered to some mainframe. It's going to be there forever. The latency of moving that to the Cloud, the bandwidth requirements, right, oh boy, it's going to be forever. So we really see it as a hybrid world is not a waystation to the future. It is the future.

**Fairbairn:** How do you see this or what is the split in your revenue for private versus public cloud?

**Gelsinger**: Yeah, so I'll give three different perspectives on that. Our revenue is over 90 percent private cloud. That's our heritage. We're building up the public cloud revenues as we grow as a company. We're now at 10 percent of our revenues at that level, and I say we are disproportional in private cloud, the private data center, because that's where our heritage has been and we're a little bit late, as I already said, in building up some of our public cloud offerings. The industry metric is now about 75-25, right? We're about 25 percent of revenues and workloads are now public cloud based, right, today versus 75 on-premises and if you measure that by workloads, to me, that's a pretty good moniker, right, because when you're in the cloud, they're operating it. When you're private, it's not just the hardware and the software but also the personnel costs, of running it for you. So workloads, I think, is a pretty good way to measure those, and today it's like 75-25, 70-30. I expect it gets to be 50-50, and some would project that it gets 70 percent public cloud, but then I believe that edge becomes this force going in the other direction. So my best view of the future is it gets to be 50-50 but today there is something on the order of 170-180 million workloads as measured by OS instances, private, and public today and we expect that becomes well over a billion over the next decade because computing continues to expand. The applications continue to expand and then key new technologies like containers and Kubernetes will help to expand, and microservices will increase the number of workloads as we go forward.

  

**Brock:** Is it-- do you see a similarity between-- and thinking of a little bit of the microprocessor story, and thinking about hybrid clouds where I could almost think of it about networked microprocessors so sometimes where the-- if you have a private cloud, it's the same structure as the public cloud but it's just a matter of scale. It's server racks with lots of microprocessors so that, in a way, these organization of computing into these clouds, whether they're at different sizes, in different-- run as privately, private clouds or large public clouds, it's almost like the cloud has become the unit of compute where the microprocessor had been before. Is there anything to that? I mean because I hear...

**Gelsinger**: Yeah, I believe there is a lot to that, and I'd say the Cloud is now the data center, where, boom, you have these at scale. You have such massive resources available. They're readily available, easy to use. I've always viewed that computing follows the Gas Law. It fills the available space. It always keeps expanding and it's really more a statement where computing, how much computing do you want to do? Well, as much as I can afford. How many more simulations do you do, before you send a chip to the fab? Well, as many as I can get done before I think I've exhausted it, but there are always more tests you can run. How many more analysis of your radiology results do you want, running through the AI algorithms to determine if it's carcinomic or not? Well, as many as you can afford, right? Please give me the best results you can, and the list goes on and on and on. So to me, computing has always wanted to fill the available space where the available space is often more limited by economics than anything else. And if I make the unit cost of computing lower, and the ability to reach the data -- so data storage, the network. I got to be able to connect to it. It's the whole set of computing resources. Every time there is a dramatic decrease in that, you open up new opportunities for computing. If we use the AI example, hidden Markov models, convoluted neural nets, et cetera, those ideas were around, all of a sudden that got economical, and Cloud made it economical and all of the sudden Cloud made datasets large enough that I could use learning algorithms that before were infeasible, now became feasible as well. So that combination of compute capacity and datasets, allowed AI to start demonstrating meaningful breakthroughs and now it's sort of like, “Wow, how much computing do you need for AI?” Well, the learning algorithms, it's almost unlimited, right? Really, if you give me another thousand GPUs in my GPU farm, I'll use them all. Many of the hardest problems in computing have always demonstrated this characteristic, whether it's

  

weather prediction, whether it's predictive modeling, whether it's computational fluid dynamics, these are n-complexity algorithms that, boy, you can just keep throwing computing at them. You can come up with a different algorithmic breakthrough. It just keeps expanding, and I think that's one of the things that Cloud has done. It's making not just the accessibility but also the cost characteristics very predictable for it, and you can look forward. I can tell you today what 1000 cores will cost you, with a terabyte of memory next year, with high accuracy. You can access it. You can predict where it's available. You can start experimenting with those. Then why are people putting GPU farms into the cloud? Oh, it has a lower economic cost for certain workloads to run them through GPGPUs. Oh, Jensen's a happy guy, as that all “Boy, I need more of those” and then people are saying, well, some of the core AI algorithms, maybe I commit them to FPGAs, because I'll get another order of magnitude improvement in the cost of economics and if that happens, what's going to happen? I'm going to do more.

**Fairbairn:** Even more.

**Gelsinger**: Right? Because now I can start making more aggressive use of some of those learning algorithms. Maybe it's going to be inference algorithms. Boy, I can't afford to put inference at the edge yet. Hmm, so I might have to compress those and do some algorithmic breakthroughs. Maybe I have to start putting more fixed function into some of my compute capacity so I can get it to the edge because I really, if I can do inference at the edge, wow, I can start open up new machine learning, new vision algorithms, new detection algorithms at the edge. But I need to get them maybe two orders of magnitude cheaper or lower power and boom, I can go make it happen. And every time you do that, oh, all sorts of new applications emerge, when you do it, and that's really the beauty of technology, sort of over and over again. Because remember we started just doing-- all of computing started out by computing missile, bomb paths and being able to compute census. It was just, you know, I just need to add. So we've just been doing calculators at scale for our entire existence.

**Brock:** Well, it's interesting because the notion of a hybrid cloud goes inside of a data center with-- it's not just, it's not a monoculture inside there. You have the GPUs and FPGAs, the standard kind of server backbone. That, I think, is fascinating but also listening to you talk about the economics of the Cloud, if you will. It has so many features that are similar, that are familiar

  

to you from sort of the economics of the silicon transistor. As the price goes down, you experience over decades this hugely elastic market, and it's the-- maybe it's the same phenomenon at different registers, levels of abstraction, but it's so similar.

**Gelsinger:** It continues to sort of replicate the same model of, "Hey, if I make computing a lot cheaper, I get to go to new algorithms" and the beauty of the Cloud is you benefit from Moore's Law because you're putting the latest processors in there with more cores, clock rates, et cetera, bigger memory footprints but you also get to benefit from the distributed nature, because I now can assemble 1000 cores for a few hours. Then, I can disperse them, as well, so you're not-- don't just get the I'll say the scale-up characteristics. You also get the scale-out characteristics of the Cloud, so you actually get two dimensionalities working for you when you go to the Cloud. Clearly, you get it through this easy-to-access API portal now, as clouds are competing with each other on raw economics. You got all these people figuring out, "Oh, I can tie them together better if I make the networks faster." "Oh, I can tie them together better if I use this NIC offload." "I can tie them better together if I change my distribution algorithm." It's a one-by-one, you got all this engineering capacity saying, "How can I really drive that competitive advantage of making more available more rapidly?" Imagine if you and I, we were in a corporate data center before, and somebody came running to us and says, "I got a great idea. I just need you to stitch a thousand GPUs into your data center and you can start running these machine learning algorithms. I think you're going to be able to predict consumer behavior in your marketing programs more effectively than ever before." Now, imagine, tomorrow you said, "Oh, it's a pretty good idea, showing some promise. I'm going to go to the CEO tomorrow and ask him for an extra hundred million of capital so I can put up the GPU farm so that I can do new marketing insights." How do you think that's going to go?

**Brock:** Probably a hard conversation.

**Pat Gelsinger:** Probably a hard conversation. But now I get to go to the Cloud and I say, "Hmm, let me try that experiment." I'll rent the thousand GPUs and combine it with the thousand cores I have. I'll run the experiments over the weekend. I'll produce some insights and I'll walk into the CEOs office and say, "You know what, I spent a few thousand bucks on my Amazon bill this weekend. I apologize, take it out of my budget or my hide or so on, like that. Here's the bill for

  

ten thousand bucks. But here's the insight that I got from running those algorithms over the weekend. Let's give this a try to market this way." Oh, my gosh. I've changed the economic model from a $100 million for capital to overrunning my budget by 10 thousand bucks this weekend and giving you something in a few days. That is just changing our ability to leverage compute capacity. That's really the thrilling aspect of, I'll call, the superpower of the Cloud.

**Brock:** Could you tell us a little bit about forming that relationship with Amazon Web Services and some of the other public cloud providers that you mentioned? This, I would imagine, is a big deal. If I'm thinking of it the right way, and please correct me if I'm not, it's almost choosing the operating system for your Cloud.

**Pat Gelsinger:** Yeah.

**Brock:** So that's a big choice. Can you talk about how you've approached that?

**Pat Gelsinger:** The Amazon partnership was one of those seminal, the shot heard around the world, changing the cloud industry as we partnered with them. As I described it, five years ago, I stood on stage and said if you used Amazon you were stupid. Andy Jassy stood on stage and said if you run your own data center, you're stupid. Now we're on stage doing bro hugs and announcing a joint future together. It really has been one of those fairly significant changes for both companies. Amazon does not have this long and deep reputation for partnering. They're chewing up and destroying and disrupting industries. So, as I also described it, we had the last and final board meeting on seven different occasions to approve the strategy. Because it was one of those, "Hmm, can I really bet on this partnership? How's that going to go? What's the future here? What's the give and the get? How sustainable is that relationship?" You're going to be in it for the long term. You're going to throw hundreds of engineers at this partnership and go build that out. It's worked out to be extremely positive for both companies. There was just a big information article, the information just was released recently on the bet that the two companies is starting to produce dividends for both companies now. We really are seeing that momentum starting to emerge. The idea of the hybrid cloud, as we already discussed, is a very merited, sustainable strategy well into the future.

  

**Brock:** I'm also thinking about in the move to VMware and thinking about the sort of people building the technologies in VMware versus the sort of people who are building the technologies when you were with Intel. Is it very different working in the software context versus working in the silicon context? Could you talk about the similarities and the differences?

**Pat Gelsinger:** Yeah, software versus silicon was a pretty radical shift in understanding and how those engineers work. In a silicon project, major microprocessor, you have to get your satisfaction on five-year increments. For four years and 364 days, you're in the salt mines. Then you have one day of glory when you announce the chip. Then you go back to the salt mine for five years, and you build these teams. A big microprocessor project now might be six or seven hundred engineers. So you ramp it up over time. You got test people and fab people and package design and all these kind of things. Clearly, by the time you release a new microprocessor, you might have had two, three thousand engineers touching it. Five years, thousands of engineers, oh, my gosh. This was some of the largest scale engineering. One time I had the general, I think it was a two-star general, who was in charge of the Joint Strike Fighter program for the military, come and talk to us about techniques used for large-scale engineering. Because you just don't find many projects that are that scale. Five years, thousands of engineers, name all the projects at that level. Boy, okay, we got airplanes, what else? There just aren't that many at that kind of scale. Now a software project, it might be four engineers as the nucleus. Then they get a prototype running in three, four months, particularly very much in an agile, cloud-driven, microservices world. The timescales are so dramatically different. For thousands of engineers, five years, five engineers for four months. The type of talent is very different. Hey, on a silicon project, you know you have a few, I'll call, prima-donnas, the core architecture people at the center of it. But you need lots of execution machines. People around with many different skillsets type of things. A core software team is very small in comparison. Again, it's not one but, boy, as long as you got those four or five that can work together, okay, go for it. Over here, it's like you're trying to create this distraction-free, focused environment for them. Over here, everything's distraction. Have you worked on the packaging team? What about the test team? Have you lined up the process thing? You're always working across all these different sides. So it's quite different at that level. But there's also a lot of, I'll say, core engineering disciplines that serve you well at that level, where you really just dig deep into the technology, understanding

  

how the systems work. Obviously, big products, it isn't four people forever, projects get bigger over time, you have all sorts of other interconnectedness. The core for your team today is probably five or six hundred people, but it's modularized into many components as well. I'll just say, good engineering skills and disciplines work across any industry. They really do. At a place like Intel, you learn good, core engineering disciplines and skills and how to program-manage big projects and teams and organizations. I just have great respect for all the learnings that I got in my decades at Intel.

**Brock:** I wanted to ask you if you could talk a little bit about, in more recent years, the whole relationship between EMC and VMware and Dell. I have to admit that, in preparing for the interview, I was trying to parse out all of those moves and it was very complicated. I'm quite certain that I didn't understand it properly. So I thought maybe you could help us understand that.

**Pat Gelsinger:** Yeah, so in 2012 I became CEO of VMware. The storage company, EMC, and a few of the other businesses then, was starting to flatline. People were looking at, okay, where's the next phase of EMC? So as '15 is coming around and Joe Tucci is coming to the end of his tenure as CEO. It's like, "What's next?" EMC started to look at its options and, essentially, a handful of options emerged. One was spinoff VMware and give it to the shareholders. It's going off pretty nicely. Another was, we should merge with HP, maybe merge with Cisco, maybe merge with Dell. So a lot of turmoil, both of CEO transition as well as company transition. It was a very tumultuous time. I described it as the worst year of my life. A variety of personal things going on as well. I'd broken my foot so I'm on a knee scooter, an invalid for a year. It was my right foot, I couldn't drive so I'm rolling around everywhere. I had a son who had cancer, so he was going through treatments. Thank the Lord, he's healthy now. We had two weddings that summer. I remember I was walking my daughter down the aisle with my broken foot. As I describe it, this eye was tears of joy and this one was tears of pain. The Dell EMC merger gets announced. The VMware stock cuts in half through this process. It's being described that I'm being fired and I'm going to emerge as the CEO of Dell. So everything, rumors I'm being spit out, we're going to be spun out, I'm going to be fired. I was, in fact, sitting next to Michael Dell at an event at the time and one of the rumors comes out that I'm about to be fired. So I took my phone and I showed it to Michael, "Something, I don't know?" All this tumult is going on. The

  

stock price, as CEO, it's like your scorecard for everybody to see. We're at half. We got a joint venture that collapsed. So just huge tumult in this period of time.

**Fairbairn:** What year was that?

**Pat Gelsinger:** So this is '15, '16. Out of that, the merger gets announced, it gets consummated. We, essentially, laid out a path to leverage Dell, to remain independent and to really see the vision that we had been working on for several years start to be accelerated. We hit a low of $43 a share, not that I remember the exact number, but $43.58 to be exact. Today, we're trading at $188 a share. So over a two-and-a-half year period, we've seen an extraordinary acceleration of the business. We saw accelerated growth rate in revenue. Obviously, the valuation, the new business areas, Dell as a partner is accelerating our growth rate as they're highly motivated to sell and bundle our products into their solutions. So I like the Churchill quote, "When you're in hell, keep going." It's one of those things where you learn resilience in those periods of time. One night, I came home, in the middle of that, and my wife looks at me and says, "You have become unlivable." She was right, I had. I was wild like a spring or a coiled snake. It was a very rough time. But now we're on the other side of that and seeing extraordinary success. But you also learn, I'll say again, a lot of this I look back on my Intel development and training where there's cyclicality in the semiconductor industry. You have these harsh times and how to deal with having to lay off and restructure. None of those learnings were wasted in the process in getting the company through to the other side. Now we're enjoying extraordinary success as a public company, but with a majority shareholder named Dell. So independent yet seeing great benefit to the interdependence.

**Fairbairn** What percentage does Dell own?

**Pat Gelsinger:** Eighty percent. So they have taken ownership of the EMC portion.

**Fairbairn** It remained at that 80 percent?

**Pat Gelsinger:** Yeah.

**Fairbairn** The other 20 percent is publicly held?

  

**Pat Gelsinger:** Publicly traded, I hope you're a shareholder. You would've been well-rewarded with one of the highest total-

**Fairbairn** I was just thinking that I would have been.

**Pat Gelsinger:** Yeah, one of the highest total shareholder returns over the last five years. We had a major special dividend last year. Between dividends and stock appreciation, we're now at 50 percent per year over the last four years or so. It doesn't get much better than that. And more to come.

**Brock:** Unless there's an aspect that we've missed about the VMware story that you think we should capture, that I've missed, I did want to talk to you a little bit about your life outside of work before we end the interview. Is there something that I missed?

**Pat Gelsinger:** I think we've hit the big ones.

**Brock:** Okay, well, I wondered if you could, in preparing, I just saw some just amazing snippets about your life outside of work. Just very active, like climbing Kilimanjaro, I think I saw, and also, I believe you're very involved with the creation of a school somewhere. So I'd love to hear about those other activities and interests that you're finding time to pursue.

**Pat Gelsinger:** Yeah, very early in my Intel days, I became a Christian. That faith perspective has been central to character. Wrote a book*, The Juggling Act: Balancing Faith, Family and Work*[1] and really integrating those together. So this idea that each of us has a higher purpose well beyond our day jobs has always just permeated me, my family. I can show up for every day and, hey, I work for God but I'm happy to get a paycheck from VMware or Intel or whoever it was. So lots of charitable activities. My wife and I committed many years ago that we'd give an increasing percentage of our gross income to philanthropy. So we made that commitment. We were giving away 10 percent. Now we're at close to 50 percent. It was 10 percent of nothing when I was a technician at Intel and now it's 50 percent of a CEO's salary. So it's a big number

[1] [Interviewee's note] Gelsinger, P. (2008). The Juggling Act: Bringing balance to your faith, family, and work. United States: David C Cook.

  

now. One of the examples was the work in Africa, in Nairobi. The Kilimanjaro climb was a fundraiser for the schools in Nairobi. When we first started to work there, it was a couple hundred kids. About 200 kids were in the schools in Nairobi. Now it's over 17 thousand kids are in 22 different schools in Nairobi. The statistics are exceptional. Again, a lot of these are born of AIDS parents, extreme poverty, well below a dollar a day environments. The mortality rate was 30 percent for the slum kids. The kids in school, it's less than five percent. Twenty-five percent of these kids wouldn't be alive. Because it's not just a school, it's also healthcare, it's spiritual training, life training and jobs training. The average schools in Nairobi, 20 percent go to college, most of those in to state schools. Forty percent of the kids from the schools go to college. So these are slum kids that are now having two acts, the efficacy into beyond high school education. One of those is in his second year at Stanford on an international scholarship. A slum kid, born of two AIDS parents, raised by his grandmother in extreme poverty, in some of the darkest slums, is now in his second year of Stanford. Just extraordinary. So the fundraiser, the Kilimanjaro climb, was a fundraiser specifically for building a girl's boarding high school. Girls are particularly susceptible to tribal patterns. An 11, 12-year-old girl, being raised by an uncle or a grandparent or something like that, a tribal leader, is worth maybe four cows, maybe three cows and two camels. Okay, great. She becomes the third, fourth, fifth wife. Being able to keep girls, in particular, in school was the purpose of the fundraising climb. So we raised about $300 thousand. Our goal was $175 thousand. I never like to meet a goal; it's always beat a goal. So now there's a girls high school that's going to open this Fall as a boarding school just outside of Nairobi to keep girls in school. Quite excited about that. But it's also we're looking to expand, my wife and I, we are funding building out STEM education into all the 22 schools. So being able to have science, technology, computing labs, etcetera, into those environments and really, I'll say, building, I call it building a city out of a slum. You're building structure and infrastructure and capacity to see these kids emerge and looking at extending that to other countries, not just Kenya and Nairobi but also some of the adjacent countries in the area. That, and a variety of other philanthropies. But we've had extraordinary impact way beyond what it might mean to be the CEO of a great software company.

**Brock:** Would you say that that system of schools has been one of the areas you've concentrated on the most with your kind of philanthropy and service?

  

**Pat Gelsinger:** As I say, we're very involved with the schools in Nairobi, visit it often, invest in those. That's one of the big ones. We also worked to create a church-planting organization that's now a national church-planting organization. When we started to partner with them, they just did work in California doing a couple, three plants per year. Now, this year, they're going to do about 170 church plants, but also do ones in South America. We're working with Compassion, I call it a community in a box, where Compassion wants to sponsor kids, but they'll only do it where a church is in the community to provide an infrastructure for the kids that are being sponsored. So our church-planting organization is partnering with Compassion to plant the church, sponsor kids and transform communities. A number of those in South America, Bolivia, Peru, Brazil, Mexico, Colombia are now participating in that model. So US church plants, as well as South American sister church plants as well. So that's another one. We've also been super involved with the William Jessup University, a Christian liberal arts university, where I was on the board for many years for that, really helping them to go from a little school in San Jose to now closing in on two thousand students at William Jessup over in the Sacramento, Rockland area. So super involved with that one. We've also been very involved with the Louise Palau Association, worldwide evangelism and ministry. So that's another one we've been very involved in. Then the final one, and coming back to the Bay Area, we felt that there was a higher purpose to come to the Bay. The Bay Area, four characteristics for the Bay, it arguably is the most influential area on earth. It is the richest area on earth, as measured by per capita income. It is one of the least churched areas in the nation, one of the lowest rates of people attending church or church participation of any form. It is one of the least philanthropic areas of the nation.

**Brock:** I didn't realize that last one.

**Pat Gelsinger:** So you have the richest that are not giving, the influential that are not based on a faith perspective in any way. So we started an organization that I'm the chairman of called TBC, Transforming the Bay with Christ. That has three missions. One is to unify the Christian leadership of the Bay Area. Second is to amplify works of service to the Bay. Third is to multiply the churches of the Bay. We started TBC and now that's in its fifth year, I think, at this point, to be this influence for bringing transformation to the Bay Area and having great success. So those are the big five ministries that we're involved in. But we're very philanthropic and one of the things I'm super proud of, as CEO of VMware, is that VMware really has developed a

  

reputation, as a company, of being philanthropic, giving back, being a champion for tech as a force for good. We have what we call our citizen philanthropy program as well that we're trying to make all of our employees citizen philanthropists. Being the geeky culture we are, every year we gives multiples of pi [3.14] dollars. So last year it was a hundred pi that we gave to employees that they could go invest in the charity of their choice. We'll match hours with 10 pi dollars. So every hour, up to a certain level, will match that you can contribute dollars to that for every hour that you contribute to that philanthropy. So it could be Habitat for Humanity, it could be tutoring in your local school, it could be working. We do what we call "good gigs." We'll assemble teams to go participate and give them time to go. One of them, a couple of years ago, we connected up, I think it was, 700 schools in the jungles of Malaysia became networked. That was a good gig of a team of 50 people that we sent there to go build out that. So lots of these things. So my own values of philanthropy, giving back, investing in causes that are higher than yourself, are well represented in the company. Many of those were there before I got there, many of these programs, but we've really been able to accentuate and extend them in significant ways.

**Brock:** That's fantastic.

**Fairbairn:** Tremendous, how did you get involved with Africa? Some of these others I can understand.

**Pat Gelsinger:** One of the people at the church that we went to was a missionary there. So we became somewhat loosely associated with it. Then we took our kids and visited the missionary friends in Africa. Then our hearts were stolen. I'll say, once your heart's been moved to something like that, like a different picture, seeing where I'm lecturing and teaching to a thousand boys in their school uniform. These kids, a school uniform, you might consider that in the US, "Oh, that's so old school," for there, that might be the only nice set of clothes they have. It's a picture of dignity for them to be able to do that. Boy, kids screwing around in school here, "Why do I need to do this?" Over there, they realize that that might be the only opportunity they have to leave a life of squalor, their commitment to studying. They'll just come, I got pictures of my wife where maybe a hundred little kids are just crowded around her just wanting to touch her at that level. You've been in a few of those experiences, you sort of say, "Yeah, what I do for my

  

day job is good but what it enables me to do in my spare job is truly transformational and touching lives for eternity."

**Fairbairn** That's wonderful. Bringing us back from the philanthropic side to the more practical Bay Area thing, and that is the story of VMware and the Stanford Industrial Park. What has been VMware's experience in the Industrial Park? Has being there been of any value impact or it was just a place to land? What are your observations having been there for a period of time in terms of the impact of the Industrial Park on companies like yourself?

**Pat Gelsinger:** Yeah, some of those, I'll say, decisions were made and some of those were happening without me. So there are probably a few others that can give better perspectives on that than I might be able to. But VMware, your typical startup company, is trying to slavishly find anywhere to meet. Its initial location was actually what was called the "Cheese House." It was a couple of rooms on top of a restaurant, a little sandwich shop, sitting in the Embarcadero Mall. So that was its first location. Then, essentially, its first permanent location was a few buildings at the end of the Stanford Business Park triangle where it's located now. So that become the first place that it really landed as its true home of scale. There it's been this phenomenal relationship. One is there's been a strong affinity to stay close to Stanford, be nearby. I call it, I just want to build a thoroughfare between us and Stanford. I want our people going over there lecturing, interaction. I want every bright, capable student to come our way. We do different research programs, fellowships, etcetera, where we really see that ongoing vibrancy of ideas, innovation, talent, people in both directions going back and forth.

**Fairbairn** So the physical proximity is important?

**Pat Gelsinger:** Absolutely. To me, if you're more than a bicycle ride away you're not effective. So being, literally, a graduate student can be finishing finals on Friday and on campus on Monday. Boy, if they finish finals on Tuesday morning let's get them Tuesday afternoon. Yeah, you really want that vibrancy of relationship. The Business Park has just worked out pretty fabulously for us because it started as a couple of buildings and then it became six buildings. Then we took over some more property from SAP. Then Genentech Roche. We had the opportunity to take over the rest of the campus and it was a decision that was made just as I was

  

becoming CEO. We renegotiated all the leases on the property there, from the Business Park, and established ourselves. I think we still have, I think, about 30 years left on our lease with the Stanford Business Park. So my view is if we can negotiate a hundred-year lease we go do it. Because we really find that that proximity to Stanford, the location that we have, is just exceptional. If you come onto our campus, I think the Fortune "Best Place to Work" article described us as a Zen-like setting. I've always said, when I drive on Stanford campus, or go for a walk on Stanford campus, I feel better. My blood pressure goes down. My ability to think goes up, the interactions that you get with the energy, the creativity. We want the VMware campus to be identical. People can show up there, like, "Aww, I can relax here. Come and eat my peanuts, come and walk around the campus, our Redwood Walk, or our Alleyway, or areas that we have for community to come together. Come and see my turtle pond." There's all these kind of things where people just feel, "Oh, this is a place I can work. I can feel good. I can innovate here." There are great people. That's really that other connection. If you're working on great problems, you have great people, you can attract great talent, you treat them well and you have a very value-centric culture, people will be here a long time. That's what we've been able to create. Even as a worldwide company now, we're five thousand people in Bangalore site and our sites around the world, a thousand people in Atlanta, probably 500 in each of Seattle, Boston, we're closing in on a thousand in Sofia, Bulgaria. There's no doubt that Palo Alto is the center of VMware. It is, it has been and it always will be in that campus setting.

**Fairbairn** How many people do you have there on the campus?

**Pat Gelsinger:** We're closing in on six thousand. We just took a building across the street as well, so I plan on keep building up the site. So every time something adjacent becomes available we'll just keep spreading our tentacles out.

**Fairbairn** Is there a feeling on the Park that they don't want anybody like yourselves to get too big? Like, they say, "You're got enough," they want some diversity or whatever within the thing? What percentage of the total Business Park do you occupy?

**Pat Gelsinger:** I don't even know. I'm not even sure what percentage we do represent. We've never had any pushback, because I think they have viewed us as great partners, great tenants. We

  

also have a great relationship with the City of Palo Alto. I think everybody views it, "This is good, let's do more of this." One example, one more little story on this that I absolutely love, I was walking by the office of my head of HR. Sitting there, waiting for my head of HR, was my head of sustainability. So I go plopping in and sit down, looked at her. I said, "So we set a goal of being carbon neutral by 2020." She says, "Yeah." I say, "What would it take to get that done a lot faster?" She says, "Well, what do you mean?" I say, "Let's get it done two years sooner." She's, like, "Ahh." Niccola is having this, the CEO is telling me how can I go a lot faster, for the core of what she believes in with our sustainability program. So this interaction lasts two minutes. I say, "Come back with a plan. I want you to come back and break the envelope, come up with some creative thinking." Because we've said one of our corporate mottos is innovate in everything. Not just in the R&D and the products. So out of that came the program that we're now well underway in implementing called a community microgrid, where we're building out a large solar capability, large battery capability and the ability to be both generating more of our power locally, storing it locally, but also becoming part of the grid for Palo Alto, which is its own mini-utility for Palo Alto as well. So we can both receive as well as feed back in. The objective of working with the City of Palo Alto is that we'll become a sustainable emergency response center as well, where emergency happens, earthquake, something like that, boom. VMware's campus is not only- because you know you have power. We're building more redundant communications capacity and to build out some basic medical capabilities as well. We'll become an emergency response center for the City of Palo Alto. So some of those kind of things. As you're doing projects like that, where people say, "Wow, the campus is beautiful. It's representing our community values. We're the source of innovation and we're doing things that just make the community better." Like, okay, yeah, we like you guys hanging around.

**Fairbairn** Is there any interaction with any of the other companies in the Industrial Park?

**Pat Gelsinger:** Somewhat, but that's not the center. The relationship with Stanford is the center.

**Fairbairn** Okay.

**Brock:** Would you like to invite Pat?

  

**Fairbairn** Yeah, so we have a final request. That is, as you've probably seen on the wall downstairs, we have this program where we ask innovators and entrepreneurs and venture capitalists like yourself to provide us one word of advice that you would give to an entrepreneur starting off. You write the word on the card, you sign your name to it, and we take your picture. So you own that, that becomes your word. Yes, there are duplicates and so forth, but everybody has their own flavor and taste and so forth on it.

**Pat Gelsinger:** Can I give a two-word answer?

**Fairbairn** You can give a hyphenated.

**Pat Gelsinger:** Yeah, we'll hyphenate it if we need to.

**Fairbairn** Yeah, you can give a two-word answer. Normally I would say no but we've had a few.

**Pat Gelsinger:** Right, because I'll just say, the core of who I've been has always just been "hard work." I'll say, what I lack in intelligence, I make up in hard work. One of my favorite little, you know who Steve Prefontaine is? He was the Olympic middle-distance runner from Oregon, where we lived for many years. He says, "I may lose but the other guy will bleed to beat me." There's this grit aspect. When you've been born and raised on a farm, as I say, when I came to Intel, it was, like, "Hmm, no horses kicking me. No cows biting me. I'm not covered in sweaty hay dust at the end of the day. This is close to heaven." I just outworked everybody. Why would I leave? I like to work. Showing up for milking time was before 5:00 and if you quit before the sun went down you were pretty lazy. So hard work.

**Fairbairn** That's great. That's perfect.

END OF THE INTERVIEW

  