



**IBM Tape History – Session 4:
LTO Virtual Company Panel:
John Teale, Ed Childers, Bruce Master, Brad Johns**

Moderated by:
Tom Gardner

Recorded: October 15, 2015
Tucson, AZ

CHM Reference number: X7684.2016

© 2015 Computer History Museum

Introduction

This is session one of five sessions held in Tucson, AZ, regarding IBM's tape storage history. The five sessions are:

1. Tape Media (CHM catalog number: 102737992)
2. Overview of tape products and product management (CHM catalog number: 102737994)
3. 3480 tape drive (CHM catalog number: 102738021)
4. Linear Tape Open (LTO) Consortium (CHM catalog number: 102738023)
5. Recovery of tapes damaged in Challenger disaster (CHM catalog number: 102738025).

See IBM Tape History Session 1: Media for an overview of IBM Tucson,

This session is about the Linear Tape Open consortium organized by HP, IBM and Seagate beginning 1997.

Interview

Tom Gardner: This is Tom Gardner, representing the Computer History Museum at our fourth IBM Tucson tape history recording session at the headquarters of West Press Inc., in Tucson, Arizona. Today's session will focus initially on LTO, an acronym which our first interviewee, John Teale, will explain.

John Teale: Good morning, I have gone through my biography in some of our other sessions¹, so I'm not going to spend too much time on it; just to remind you, I am a retired IBM engineer. I worked there for 31 years. All of my 31 years was in tape technology, the tape business, tape product development, participating broadly through the worldwide market, in other words, all of them. Today we're going to talk about LTO. LTO stands for Linear Tape Open.

I have three guests with me that we'll be introducing in a little while to talk about some different aspects of what I'll call the birth of LTO. LTO was announced in November of 1997. There was a joint announcement from three large companies: IBM, HP, and Seagate, of their intention to develop a new interchange recording standard for a tape drive, to be developed. It's important to distinguish between what LTO is and what it isn't. It is not a collaboration to produce a product by three companies. It was simply a collaboration to agree on how we would be able to interchange data with each other - the ability to go produce that product was an independent effort by each company. Let me say that LTO is a current, ongoing business. Technically, you could say that LTO is not history yet. However, the announcement of the consortia in 1997 is history, and that's the history that we're going to talk about today. We're not talking about the IBM LTO products. That would be another day, and another set of people.

So I'm going to set this up by talking about the what, the who, and the why. And in order to do that, I need to do a little backstory on what the IBM environment was like in 1997, what was the tape market like in

¹ See IBM Tape History – Session 2: Overview, beginning circa p. 27, recorded October 13, 2015, CHM catalog number 102737994

1997, and that's what developed the motivation for what later happened, and I think will make it very clear. So we'll start with the IBM environment in early 1997. We talked in some of our other sessions about how we had developed a bunch of technology in the early '90s, that we were intending to use to enable the follow-on to the 3480 and 3490 series of products. They would later become known as 3590 and 3570, but our efforts in 1992 were interrupted because the regime at the time decided that we were going to exit the tape business. That pile of technology was put on a shelf. Many of us redeployed to other activities. One of the activities I was redeployed to was buying other people's stuff, and qualifying it for other IBM servers, like quarter-inch cassette, four millimeter and eight millimeter drives and media. So that's how I met all of the players in the tape universe. Other people were redeployed, some people were simply told, you're on managed departure, and you can simply stay home and get paid until we decide when we're going to exit you. So it was a pretty tough time.

But we had a regime change, within a year or two. A fellow named Jim Vanderslice and his team came in, and reversed that decision, and said, "We're going to get back in the business, dust that pile of technology you had, and get it deployed." And he had a lot of ideas from his printer background on opportunities for tape that we didn't see for ourselves. So for example, in printers, they have the equivalent of what they call the razorblade model. Sell them the printer, or even give them the printer, and then make a lot of money selling them ink, and paper, and maintenance, and parts, and service. Because tape has a removable component, namely, the media cartridge, these are typically not produced by the drive manufacturers. There's a whole separate industry infrastructure that produces these. These can be thought of as the razorblades, if you will. And Jim Vanderslice envisioned a model where to the extent that drive maker's intellectual property is required to produce the media cartridge, then there was an opportunity to extract profit from the people that make these things, and sell these things. And we simply are the market-makers who deliver the drive, the capability, and the value proposition. Pretty exciting stuff. I didn't really understand much of what it meant at the time, not having ever experienced this type of business model.

But we had a problem. In early 1997, Barbara Grant and Kevin Reardon, who were our leaders at the time, told me that we're getting back in the technology business, and I said, "Well, that's going to be pretty tough to do, because we put all of our thin film head equipment in vans in the parking lot, and sold it on the used equipment market. All of the tape head developers and designers are either retired, dead, or working for competitors." And there's a little bit true, all three of those statements, unfortunately, it was true. And he said, "Well, we want to be back in the technology game, how do we get in it?" And I said, "The only way we're going to be in it is go to San Jose, and beg the hard disk people to help us out." Is it possible to make a tape head in a hard disk line? We didn't know the answer. They told me, they gave me a ticket to San Jose and said, "Good luck." I was in San Jose six months. It took me six months to hire the first tape head manager in San Jose. What I learned about the San Jose environment is that when you're asking people to consider an opportunity or a challenge, they've got four godfathers and three godmothers they have to check with that are helping them manage their career before they'll even talk to you. Tape was not viewed as an opportunity, it wasn't viewed as glamorous, it wasn't viewed as anything that would enhance the resume of a hard disk engineer. Kevin came out, and said, "Okay, you've got interviews lined up for me?" And I said, "No, Kevin, I don't have anybody who wants to talk to you about this job." Kevin, very creative guy, rebranded the job as a notch your belt, one year in and out, great

opportunity to enhance your resume. He sold it to the highest executives in San Jose at the time, in particular, Bob Scranton. And Scranton endorsed it, and pretty soon, we were getting the best and brightest candidates out of Almaden Research, applying for this job. First manager we hired was Shole Hasami, who did the technology transfer from Almaden to San Jose, on something called a Spin Valve Head. I don't even know what a Spin Valve Head is, but it sounded pretty fancy to me. Can I have some of that stuff? <Laughs> Long story short, we finally got some people going. We had a media partner at the time, and we ended up expanding our search for other media partners, because the media partner we had was about as obsolete as we were, in terms of reentering the technology game. And we had resolve, we had sponsorship, now what?

Now let's talk about, a little bit about the business environment at IBM, and then we'll go shift to the market as a whole. One of the reasons that the decision was even made in the first place that we might exit the tape business is because all of IBM's internal development efforts were aimed solely at what we called the enterprise market. In other words, mainframe attachment. So our business was growing with that market growth, which was low single digits. That was not a shiny-looking P&L to IBM, because the new regime of IBM were growth, maniacally focused on growing your business faster than the industry. And the only way to do that is to expand your industry participation. There's probably other ways to do it, maybe you've got some channel things you can fix, lot of ways to do it, but one way to do it was to basically design a tape drive that would serve a bigger market than just the enterprise market. It was called an open systems tape drive.

Unfortunately, we had a black eye. We had already done that with a product called 3570. That was a product that was specifically designed to hunt down and execute Exabyte wherever they lived. We got some good traction with that product, displacing Exabyte. Unfortunately, the target moved with the advent of something called Quantum DLT. With that, now I'll talk about the broader tape market. There was a huge open systems tape opportunity that opened up in the '90s, due to the pervasiveness of computing - the maturity of personal computing resulting in the introduction of servers that weren't just Enterprise servers. You had Intel servers, you had Unix servers, you had AIX servers, you had AS-400 servers, RISC servers, and all of a sudden there was a blossoming need for tape at a completely different cost point that would openly attach to a standard interface like SCSI or fiber channel to serve that market. And there were a plethora of companies that blossomed to fulfill that market need. For example, HP and Seagate borrowed the technology in a four millimeter digital audio product called DAT, Digital Audio Tape. And since it was a digital storage project aimed at music, it wasn't a big stretch of the imagination to realize that you could easily reorganize that as a more general storage device that would take all kinds of data, not just music data. Similarly, a company called Exabyte borrowed heavily from the eight millimeter digital video technology that was emerging for consumers, and saw the same opportunity. At a slightly different price point and a slightly different capacity point, but a complementary product that would serve other parts of that new market. There was a longtime player called Quarter-Inch Cassette. Originated back in the early 70s, but various flavors of that cassette had matured into Travans, and QICs, and there were a whole bunch of flavors of Quarter-Inch Cassette. A whole bunch of companies participating throughout California, Colorado and Tandberg Data in Oslo, Norway. And so there was a lot of open participation, and at the end of the day, Exabyte began to dominate that space, and that is why, in the early '90s, when we were doing that technology for the follow-on to 3480 that became 3590, we

also had this thing called 3570, a dual-reel smaller cartridge, open attach, to go after Exabyte, because everybody was envious that Exabyte was making all the money.

Then approximately in 1994, something happened in this industry, there was a bolt of lightning. A company called Quantum bought the storage business from a company called DEC, and primarily Quantum made the investment because they wanted to get into the disk business that DEC-- I don't know the details of the DEC hard disk, but apparently it was attractive to Quantum. And along about Christmastime, Quantum opened up the box of what they bought from DEC, and there was a diamond in the rough in there, there was a tape drive in there. A tape drive that no one had ever heard of, a proprietary tape drive that DEC had been selling only for their servers for many years, but it was never publicized, or marketed, or-- it was an accessory to a mainframe sale from DEC. They opened it up, and saw the potential for this diamond in the rough, which was a tape drive. It had a cartridge similar to 3480, little kind of box, single reel. It had a lot more capacity than what Exabyte was able to deliver with a dual reel cartridge in approximately the same space. And they saw the opportunity to take that to open system, rebrand it, and go hunting, Exabyte, with that weapon. It became known as DLT, digital linear tape, is how Quantum decided to market it. And they, in very rapid order-- and there's a market dynamic where cost was king, depending on which server, and all of a sudden over time, because of the Internet, capacity was becoming king. And Quantum ended up basically wiping out all of those other formats, made them obsolete. And one by one, slowly, all of those little companies that were doing four millimeter, eight millimeter, quarter inch, many of them went one way, many of them were acquired by other people. I know Seagate bought QIC companies, and HP bought one in Colorado.

But it was clear that Quantum redefined the game. It's all about capacity, deliver it as cheaply as possible, and the IBM mantra of absolute, reliable, bulletproof was not affordable in the market, and it was not a high requirement. So we missed. 3570 did exactly what we wanted it to do, it hunted exactly where we wanted to hunt, and then the market moved way over here. And it was all about capacity. And there was no way we could get enough capacity in a dual reel cartridge to ever compete with DLT. What are we going to do? So Kevin Reardon, who was our executive at the time, and I are drinking beer in his backyard. Kevin had an optical disk background, and he brought a perspective of the power of an interchange standard. I was certainly familiar with interchange standards. IBM kind of tended to give them away. We'd go to ANCI, we'd go to ECMA, we'd volunteer all of that information. When you do a real standard, through a standards body, you sign a piece of paper that says, "Thou shalt license reasonably." You're not allowed to get a standard accepted in the international community, and then hold everybody hostage who wants to practice. That's against the spirit of the international standards community.

Well, Kevin was familiar with something called an ad hoc standard, where you don't actually go through a formal process to establish the standard, you simply make it available to be licensed by anyone. You don't have to license as reasonably, and you don't have to adhere to the rules of ECMA, or the rules of ANCI, and I think there's one even above that, that I forget what it's called. So that kind of clicked. I was filling Kevin in on what I thought our technical capability was. I said, "You know, we can take that same engine from 3570, and port it into a single wheel cartridge, and very easily produce a competitive product, with DLT, without a whole lot of invention and risk. It'll be a whole lot of work, but it

won't be a ton of invention. But can IBM, all by itself, be successful with that? Because IBM has some corporate overhead. It's a little unusual, we're talking about a market where you really can't afford that kind of overhead to be successful. And furthermore, would anybody follow us? So the next brilliant idea Kevin had was Kevin and I had already traveled the world, trying to get partners in many other areas of our business when we were trying to reenter, because we had atrophied so much of our technology. We had been to Japan, looking for automation partners, deck partners, mechanical partners, head partners at San Jose. We even partnered with IBM Japan for chips and microcode. So we knew the drill, and we just decided we needed to visit different kinds of people. We needed to go visit people that might have an axe to grind with DLT, might have a need, and who also were market-makers like us, that could help us build the market.

So at some point, we knocked on HP's door, visited a fellow named Jim Browning, the executive from HP at the time. I don't honestly remember where we met Jim. I think it was a hotel room in San Francisco or something. I know it wasn't in Boise, because I've never been to Boise. And we're being a little coy, we don't really want to put all of our cards on the table. We're kind of fishing for, would HP be interested in doing anything with us? Kevin had a prior relationship with HP, because I believe we bought some optical libraries from HP, and rebranded them with our name. So that's why Kevin was comfortable going there. Jim Browning said, "Well, what do you got in mind? What do you think you could make?" And Kevin looked at me and kind of gave me the green light, and I said, "I don't think it would be too difficult to get 100 gigabytes in a cartridge about this big." And his eyes just lit up. There was immediate resonance, because unbeknownst to Kevin and I, HP had been partnering with Seagate for many years in the four millimeter market. HP was the leader, they would kind of establish what they were going to do— HP would service a large part of retail, and a little bit of OEM, and then they would recruit Seagate to go in and fill out all the rest of the OEM channel, and a very effective relationship between those two. So unbeknownst to Kevin and I, HP and Seagate already had a project cooking internally. It just happened to be right about 100 gigabytes. It just happened to be an architecture similar to what I was alluding to. And this immediately opened the door for more conversation. Jim asked us to go visit Seagate, because if HP was in, Seagate had to be in, and if IBM didn't accept that, then HP and Seagate would just continue doing what they were doing.

I didn't really know Seagate. I knew them from buying four millimeter stuff from them. So we went to Costa Mesa, I believe it is, and we met the Seagate executive at the time, a guy named Jesse Spear [ph?], and his executive technical guy was a guy named Leroy Thompson. And then their big technical guy was a guy named Bill Buchan, who sat down with them, and HP had already given them a send ahead on why we were there, and what we wanted to talk about. They were extremely receptive. They were also sensitive to the fact that unlike HP, we were a customer. So we were a customer to Seagate, where we were not a customer to HP, so let's just say it was not difficult to get Seagate to agree to let IBM join their party.

So the table's set, the time is right, we've got a mission, we know what we've got to do. Now, how do we do it? What do we do? Not as simple as you think, because we didn't really know what we were going to do. We realized very quickly that outright collaboration on the development of a family of products would involve taking the square root of the sum of squares of our tape market share, and would probably attract attention we didn't want from the federal government. So that was out. Because you might think on the surface that that was the plan, okay, I'll provide heads, and you go get a deck, and et cetera, et cetera, you do some chips, and we're all going to be happy, and we're all going to take it to market by our

channels. Another nice thing about the construct, by the way, of HP, Seagate, and IBM is that it appeared, without even having to have anything close to an illegal discussion, that we served different markets. In other words, we had channel compatibility-- we had more channel compatibility than channel conflict. That was serendipitous. That was something that we just kind of realized after the fact, that made proceeding more comfortable for all of us involved. So the executives got together-- the three execs were Jim Browning, Jesse Spear, and Kevin Reardon. They each brought their head businessperson-- because now this is not a technical discussion. Now this is a "How do elephants mate?" discussion. And they ended up coming out the door kind of where Kevin thought it would land, that we would work together to develop an interchange standard. Completely legal, completely clean, nice, neat, and then we would go compete in the space that we created. And that was the birth of what became known as LTO.

We had contracts to sign. There were still months of legal and business wrangling on a three-way contract between our companies that would enable the technical people to engage and get to work. There was lots of discussion about, what are we going to call this thing? And I'm going to let people like Brad Johns and Bruce Master give their version later, but I'm pretty sure LTO, first of all, it had to be three letters, like DLT. We really wanted the letter L in there bad, but we also wanted to differentiate our three letter from DLT, but adding an O for open, implying that DLT was a monopolistic, single-source proprietary product, and we could use the O for Open to help us ultimately market the category, we called it a category. And in fact, we were so excited and attracted to this that shortly after we finally did do the press release that these were our intentions, we went to Las Vegas and invited a whole bunch of people, companies that we thought might be interested in participating. People like Jesse Aweida came, Juan Rodriguez. In other words, perspective licensees of this thing when it comes to pass. Hundreds of people. Wined them and dined them in a big ballroom, and gave everybody a ticket to the O Show. O for open, O Show that was at the Bellagio -- pretty good digs. And we had a great time, and everybody got kind of plastered, and served its purpose. It was sort of our first public unveiling, subsequent to the announcement. Everybody got a ticket, and lo and behold, technical difficulties, O Show cancelled tonight. The businesspeople from the three companies were hysterically running around, attempting to retrieve all of the tickets that they had passed out, so they could get them refunded, we're talking \$10,000 or \$12,000 worth of tickets. And a lot of people, more than half, guilty as charged, walked over to the box office, I gave them my ticket, they gave me 125 bucks, and I went and put it in the slot machine. I think it's okay to say that.

So that was the birth, and now it's time to get busy. How do we structure this, how do we organize? Well, in someone's infinite wisdom, a great idea, we decided that we needed to eliminate partisan bickering, if possible, we needed a structure where when you were in that room, you weren't IBM, and you weren't HP, and you weren't Seagate. You were LTO. You left your IBM hat at the door, put on an LTO hat, you're one team. We decided to organize as what we called a virtual company, and we identified, I think, six or seven unique roles and responsibilities. Each company would plug a person into each of those roles and responsibilities, and very briefly, we had what we called the executive sponsorship team. I've already mentioned their names. That was one of the roles, they were kind of-- went to them for guidance, if we needed guidance, we went to them to resolve contention, if that was necessary. We had an appeals process. And in addition, each company appointed a technical leader. I was the technical leader from IBM. You'll meet Ed Childers sitting over there later, who is the current technical leader from IBM, and I was going to say that he's retired, but I would just be joking. Ed is still in the game very much, and he'll give you an update at the end. Each company had a marketing team

leader, to say, "How are we going to promote this category? How are we going to sell vaporware until it's real?" That kind of thing, the marketing thing. And Bruce Master will be joining us later as a former marketing rep to LTO, to tell you a little bit about what that was like. Each company had a businessperson. They were responsible for resolving a myriad of business issues: contracts, licensing, details of licensing, all kinds of stuff. We'll be introducing Brad Johns, who had one time was the business rep from IBM. Each team had a finance person. Those people were a little less visible in the process, because we did have a vision that at some point LTO would generate its own revenue through the licensing of the specification. But until then, somebody had to pay the bills for all the meetings, and all the things that had to happen before we even had something to put out there. And I believe Mary Ramsey was our first finance rep, and perhaps our only one, right up until she retired. I don't know if there even is a rep anymore. In addition, even more mysterious, each company donated a lawyer. Did I say that right? Lawyer. Because we had a lot of guidelines on what we could talk about, what we couldn't talk about, because like I said, there was sensitivity about three large members of this industry working together, and there was a lot of devil in the details of licensing agreements. And one of the things, I'll caution you right now, the four of us here are not even entirely certain what we can and can't reveal. We're going to use our judgment, and hopefully nobody will get in trouble. Because we've had a lot of experience talking about something we're not supposed to be talking about, and being okay with it. So the table's set, and the three technical leads got together. We knew what our assignment was.

Oh, I'll say something about rules of engagement. This was a very beautiful construct, this virtual company. And it came with rules of engagement, process of appeals. There were a lot of things that we anticipated would happen. One rule of engagement was that everything was going to be done by consensus. This was not a voting model. If any one of us dissented, then there was no agreement until we all agreed. Which in some ways, is a little bit less efficient way to get to the end result, but in other ways, it sure goes a long way toward keeping the peace. And had we had a voting model, it would have been Seagate and HP beating IBM every single time. Because Seagate would always look at HP and say, "What do you want us to do?" The other rule of engagement was that in the event of a technical disagreement, the guidance we got from the executives was that the tie breaker will be based on best of breed always. In other words, we're not going to make a technical agreement that maybe advantages one company disadvantages another unless it is the right selection.

And let me give you a tangible example of that concept at work. Very early on in LTO 1 we all agreed that we need to have a compression. That was standard in tape. For many, many years HP and Seagate had been shipping a compression scheme known as LZ2. They were very comfortable with it. It was in all their products. They had a very mature core in their chips, and they said, yep, LZ2. IBM for many, many years had been shipping LZ1. And we had a core, and we were comfortable. And so it's hard not to be a little bit parochial. If I agree to LZ2 then I got to get a whole bunch of logic guys scrambling because I don't even have that, and they're going to be disadvantaged regarding time to market. Now remember our schedules-- our race to market were completely independent, autonomous activities by the three companies. So agreeing to something new was going to hurt. Basically meant you had to run that much faster when you got home. And you had to endure the wrath of all the people that wanted to know why the idiot agreed to that. So you can-- you're starting to get flavor of what this was like. So how did we resolve LZ1 and LZ2 because both sides were entrenched? Well, we made a good old fashioned decision matrix with best of breed at the top, and we agreed on a definition of best of breed not only in terms of the efficiency of the compression algorithm against different types of data. There's the Calgary Corpus, and

different things do things differently. But also implementation we evaluated it. We looked at it. LZ1 won the day because LZ1 it turns out from an implementation perspective is extremely symmetrical where LZ2 is not symmetrical at all. And to HP and Seagate's credit they accepted it. They signed up. The model was working. The rules of engagement were working. And it was a beautiful example. Don't feel too sorry for HP and Seagate yet because when we got to Gen 2 they wanted to get rid of our crappy, little peak detection channel and replace it with PRML. And that was best of breed. I had to agree. So later on they won the day.

So this was a very dynamic process, but I thought it was a very well-orchestrated-- now, that's not to say that the meetings were easy. We literally had people getting mad and leaving the room. I think we had one guy cry. And maybe some of these guys will share some stories with you because there's all universal personalities involved and not to mention your natural parochial leaning. So back to the technical managers getting together, we realized that we needed-- that no one of us was smart enough to develop an entire interchange specification. It's a vast document with a lot of disciplines that need to be represented. So we formed five, I think, technical working groups that came to be known as TWGs. So these were additional members from each company that would meet on a subset of what we were trying to do. One of those TWGs was the cartridge media people because in an interchange standard for tape you define the cartridge physically in great detail. You don't dictate formulation or anything, but you say it's got to have these performance characteristics such as an SNR (signal to noise ration) requirement, etc. That's a whole art into itself, and Ed Childers who will join us later was the original person on the media cartridge TWG. The invention of a new cartridge was arguably the biggest amount of work to do, and I'll tell you why that's true in a minute. We had other TWGs. We had a logical TWG. These are the guys that are defining the digital details of how you write the data and all that jazz-- ECC. And we had a few other TWGs. I don't remember what they all were. I'll just let Ed talk more about that when he gets up here. So just make a note of that, Ed, to talk about TWGs.

Okay, let's talk about this first major decision we made. None of us were that concerned about the technology. Like I said, our capacity goal wasn't that big of a stretch from stuff that we were already shipping. For example, the areal density of our existing 3570 dual reel cartridge, if applied to a single reel cartridge like we were defining for LTO, achieved the capacity goal. Our aerial density of our 3570 tape was actually 10 percent higher than the original DLT which I think was in the 80 gigabyte range [DLTtape IV (1996) was 40 GB uncompressed and SDLTape I (1998) was 110 GB uncompressed]. I don't remember the exact number. Of course, IBM came to party with the 3480 cartridge, and said here's the answer. We'll give it to you. We'll give you the design. We won't charge you anything. It's free. We'll donate it. HP said no thanks. Not because they were being parochial, but because they were serving an entirely different market from IBM. And one of their internal requirements, which was not an IBM requirement, was a cartridge that could enable in the future a half-high drive. The 3480 cartridge doesn't enable a half-high drive for a couple of obvious reasons. One is it's a little thick, and we thinned it out. But that wasn't the big one. The big one was this leader block. And the engineering requirements that go into designing a way of threading the tape. The threader that we used on 3480 wouldn't fit in a full height version. IBM didn't put data drives in servers. They were peripherals. HP did put data drives in servers. So we had to acquiesce to that, and we certainly couldn't use a Quantum DLT cartridge.

Gardner: For the record the tape is actually sticking out of the cartridge. You pulled the leader block out. You can't get it back in now.

Teale: Yes. I did that yesterday to demonstrate a point, and there's a special tool you use to fix that, and I don't have one. IBM wasn't really planning to design a whole new cartridge for LTO. All of a sudden this was an element of technical unscheduled invention. It wasn't really anticipated when we had the big party and got all excited about doing something. So the evolution of this cartridge-- I'm going to let Ed give you some backstory on it. I've already spilled a little bit of your thunder, Ed. I indicated that one time the marketing guys cornered you guys to do a white paper on all the great things about it. And I remember Tom from HP saying well, it's really not that remarkable.

Gardner: So this is actually not a production cartridge, right? It was a gift to you from?

Teale: IBM Japan after we shipped the first product.

Gardner: And the characters in Japanese on the front?

Teale: That character stands for warrior. There was a joke about my name. Everybody would misspell my name and leave the silent e off the end. So I used to introduce myself by saying Teale with an e is a Scottish warrior. Teale without an e is a duck. So this is the kanji for warrior, and it was given to me as a thank you because it may sound funny that we're internally partnering but IBM Japan is a universe away from IBM Tucson in every possible respect. So it really was like partnering with an external entity. We ended up creating 40 or 50 jobs in Japan that otherwise wouldn't of existed, and they were extremely grateful to us for including them in the team.


Gardner: So which character is warrior?

Teale: That's one big kanji character, I think. I don't know. I don't read Japanese. We can take that as a to-do.

Gardner: I think these are two kanji characters

Teale: There might be. I don't read it. Most Japanese don't even read kanji. They have two other alphabets they prefer. Okay, where was I? Oh, okay, I gave an example of rules of engagement. I talked about the TWGs. I talked about the impact of the cartridge decision. We know we spoke about 3480 the other day, and we talked about the assumptions that went into 3480 that created all the pain. Well, the cartridge assumption was probably the biggest assumption that was made that created almost all the pain because now we're talking new mechanisms. You know what's involved in developing the mechanism. We're talking new threader. We're talking of hours and hours and hours of engineering debate to agree to a pin and more on that later. I think I pretty much set the stage here, and I think now you're probably tired of hearing from me, and I'm going to let people go. Who wants to go first? Ed's going last.

Gardner: I have a couple of follow up questions-- I believe you mentioned the O show was cancelled but didn't say why.



Teale: They said technical difficulties. Apparently when you put these big Cirque de Soleil shows together there's just a ton of technology involved in putting on the production if you've ever been to one. I finally did see the O Show many years ago, and it's quite a spectacle. That particular night we were simply told technical difficulties we'll refund your money if you find the tickets.

Gardner: And you mentioned LZ1 won out over LZ2 because it was symmetrical - Symmetrical means?

Teale: It's what we called elegance of implementation. So elegance of implementation was one of our decision criteria. So the difference in performance between the two compression schemes was not that big a deal. You know that when you pick a compression scheme that you have about three different sizes of a compression buffer that you can choose from, and we tended to pick the middle one. And a lot of details but end of the day everyone agreed that the implementation of LZ1 was considerably more elegant than the implementation of LZ2. And that was our best of breed decision. And that won the day.

Gardner: So I had interpreted symmetrical is meaning the encoding decoding challenges are more or less a symmetrical implementation.

Teale: That's what I meant. In fact, there's actually reuse of logic -- so it's super symmetrical. Where the LZ2 it's one set of logic to compress and a whole different set of logic to decompress. And it's a much larger chip image to do LZ2.

Gardner: Continuing the discussion of LTO with John Teale as the interviewer and now Brad Johns as an interviewee.

Teale: Okay, we're back and I said we had some special guests today. Brad Johns was a member of the LTO team perhaps multiple occasions. And I'm going to let Brad introduce himself and tell us a little bit about yourself.

Johns: Hi, I'm Brad Johns. I'll give you a little bit of my background. I graduated from the University of Arizona in 1977. I grew up in California. The reason I ended up in Tucson, actually, was I ran in the state of California 220 yard dash final in 1972 and I did well enough that I got a track scholarship to the University of Arizona. So I grew up outside Sacramento, actually in a little community called Fair Oaks. My dad worked for Aerojet General, and part of the Gemini program they were one of the subcontractors. So I got to the University of Arizona and along the way of getting my undergraduate degree there was an IBM mainframe in the bottom of a math building. And I had taken as much math-- I had enough math to get a minor if I could take one more math class. And I didn't want to take partial derivatives. I'd gotten through calculus, and the last thing I wanted to do was to jump into it. And they had this class on computer programming in the math department for three units on an old IBM mainframe. So I learned to code FORTRAN. Basically a brute force technique is how I would describe my programming expertise in the bottom of the math department at the University of Arizona. I graduated in '76, couldn't get a decent job because nobody in California at the time knew that the University of Arizona existed. It was what? Is that a junior college? What do they do? Because they weren't part of the PAC 10 at the time. They had their own conference. So I got admitted back to the U of A in the fall of '76 in the master's degree of business administration. And that was fine. I graduated in the winter of '77 because they gave me so

many credits for all the undergraduate work I'd already done. And IBM had decided to hire MBA's. So the luck of the draw was here I was in doing an IBM interview, and they said, well, have you ever done any programming? I said, well, as it turns out I had programmed on an IBM mainframe in FORTRAN. And they said, well, check one. They were hiring for sales, but we would like them to be able to do a little programming. And have you ever done any sales? And one of my summer jobs when I'd go back home to Sacramento was door to door sales selling those little street numbers that people paint on the curb. So I had done sales. So check two. And I had an MBA. So they said MBA's-- we're hiring MBA's this year. Guy's done sales, and he knows how to program. So IBM hired me in Phoenix as a sales trainee in the data processing division in 1978 when IBM was just hiring a lot of people in sales at the time. The mainframes were going to take over the world, and the data processing division was the mainframe division. And so I was lucky enough to-- for me it was hitting the motherlode to get hired by IBM in Phoenix coming from my background.

So I started off doing a lot of sales. I worked with a lot of large customers in Phoenix. Eventually they kind of turned you loose on the unsuspecting customer population after a couple of years of training. And I had some customers like Motorola which had a very large data center in Scottsdale at the time-- five mainframe computers all IBM disk, all IBM tape, 3420's were huge there. They were doing computer simulations on IBM mainframe computers, 3033's and 3081's at the time. And then I moved from-- after being successful there I moved to Los Angeles where I eventually ended up in sales management. And I had worked with some large aerospace companies-- both Rockwell and Lockheed. And then from there about the time John mentioned things got difficult for IBM in Tucson things got difficult for IBM in the field in Los Angeles too. And we downsized significantly. I ended up moving into a marketing and consulting role in what they called aerospace industry marketing at the time. And we reduced our headcount by 90 percent over a two year period. And I found myself doing consulting where I had been smart enough-- or lucky enough to have hired some really good engineers who were very familiar with CAD systems and engineering workflows, engineering process. And I was able to join the IBM consulting group as a business process reengineering which was the thing in the middle '90s and help large customers-- aerospace customers and automotive customers improve their engineering design and change processes. So I-- that was all great. It kept me gainfully employed. At that time in IBM if you could bill your hours you could keep your job. So I was able to do that, but I was on a plane, like most consultants, 60, 70 percent of the time. So that was kind of wearing on me.

Then I got this call from a friend of mine who was the director of sales and marketing for tape storage. Pete Toronto was the director. And he said, "Hey, would you be interested? I'm going to be hiring a marketing manager in Tucson." Well, I was tired of traveling every day of the week. As interesting as the problems were I had two daughters at home and it was very difficult to only be there 30 percent of the time. So I leapt at the opportunity. I'd actually started looking outside IBM because of the time living on a plane wasn't really going to work for me. So I showed up in Tucson in February of 1997.

[T1:0:50:00]

I actually showed up a week earlier than planned because my consulting engagements had ended, and I was anxious to get started at the new job. And so I called Pete and said, "Hey, can I come in a week early even though technically I'm still working as a consultant for another week?" So I ended up in Tucson in February of '97 which was, as you heard from John, was an interesting time.

The Consortium hadn't quite been announced yet. But all the executives had pretty much lined up behind it, and we were very entrepreneurial-- I guess would be the way I describe IBM Tucson at the

time. And so you didn't really say well, this isn't my job. So when they said they needed to develop a forecast for this thing called linear tape the forecaster just retired. And so I had the opportunity to do probably the first official business forecast that went through the IBM process simply because there was nobody else to do it. And that was for my first introduction to linear tape at the time was really working on that project trying to figure out how IBM was going to sell it, but also how were we going to do this thing called OEM sales which was new from my background. I came from the IBM brand perspective and-- but there was this whole other channel that Jim Vanderslice was very familiar with. IBM was very successful selling hard disk drives at the time to multiple customers for inclusion in their subsystems. And so we had this channel; we called it the OEM channel. But basically we were reselling HDD's to a lot of system providers and system houses. And the thought was well, we need to do this thing called OEM-- so all of a sudden-- at least our initial version said that the business will say well we got the IBM stuff. We know how to do those. But what are we going to do with this OEM channel. So we had to come up with some interesting and creative techniques to try to go after them as John mentioned. Trying to understand how successful were we going to be selling IBM technology to some of these major system houses like Sun or Compaq. There was a whole bevy of system houses at the time. And there was an entrenched competitor at the time called Digital Linear Tape which was pretty much in the-- what we called the open systems world which is basically large UNIX based systems as well as probably large Microsoft NT systems at the time Microsoft servers. How are we going to be successful in that market space? What were we going to provide? So it was fun. It was my first real touch of it in terms of developing the business case for it. And then I also had the opportunity to actually announce the IBM branded flavors because we were all wearing different hats in August of 2000 when we did the official announcement that was-- I had the IBM logo version. There were other announcements from other system providers that we had succeeded in selling our technology to. And that was kind of an interesting perspective because we had a new-- we were starting to see some transitions on the executive team at this point. We had a new business line executive called-- her name was Brenda Zawatski who was actually our executive at the time of the launch. So we had gone through-- we had gotten Barbara Grant, and then it had been to Barry Rudolph, right? And then it was Brenda.

Teale: From a leadership perspective there were-- Bob Manis was involved. There were people who didn't want LTO done.

Johns: Yeah. There were some internal discussions because we were being very successful with our high-end mainframe attached product-- the 3590 at that time. You remind me of a very interesting discussion. We said, well, it seems kind of trite and-- but it was who needs a 100 gig cartridge. DLT was being fabulously successful with 30 or 40. We had a 20, I think, gigabyte drive on our mainframe. This 100 gig who's going to use it. We actually had sales people telling us we were developing a product that nobody really would want which was made our life interesting as we were doing the forecast.

Teale: There's a situation I had mentioned capacity became king, but I didn't say it was a justifiable assumption. It just became the ticket to admission.

Johns: Yeah. That was fun because during that timeframe-- so I arrived in '97. It was important enough we got incremental marketing budget to do market research. And we actually did a-- it took a couple of years because you can't do everything you want to do at the immediate. But the net result was we did in

fact validate the capacity was king because there was another format that was similar to IBM's dual reel 3570 called Accelis. We had to make a business decision. We couldn't do them simultaneously. We had to pick one or the other. So it probably validated what I think everybody knew, but sometimes that's what good market research does is that you actually had some numbers on a piece of paper that said yes, customers really do prefer the capacity offering versus a lower capacity offering which did have some other positive attributes such as very rapid access to the data. So that was a conscious decision. In fact, there was, I think, an expectation on the part of our partners within the LTO Consortium that we were going to do that format because we did all of the contractual work to actually provide that format in a license to someone if somebody wanted to build that. There was an Accelis format license available.

Teale: That's just a piece of clarity. I had mentioned the 3570-- the dual reel open product that was going to go where Exabyte went. We wanted to leverage this LTO Consortium to help us if it was viable to keep 3570 alive. It wasn't a little bit like it.

Johns: Uncannily.

Teale: Exactly. HP and Seagate were resistant. They didn't care about IBM's little problem selling a product that's kind of a non-starter from their point of view. But we kind of-- as part of our condition of coming to the party we accepted Seagate's, and now you have to accept a second flavor of this standard. So we had two flavors-- one Ultrium, one Accelis. These were foreign words to the audience. So Ultrium's the one that kind of survived and became what people think of as LTO. Accelis' first spec got done, but then it went away. It sounds like you're transitioning into LTO or did you want to finish?

Johns: There's some more things around that-- well, biography is done but I think it was kind of fun--

Teale: Were you retired?

Johns: Well, that's a few years later. But the point that was, I think, kind of interesting at that time is we-- by the time we did the announcement, like most projects, it cost you more and took longer than you thought. So there was a great deal of focus within IBM on the product announcement itself in 2000. And so we, like I mentioned, had done a tremendous amount of market research. And we were fairly confident that we had a very competitive offering in the market place. We had-- if I recall correctly surveyed over 500 customers. We had done focus groups. We had product concepts. We had talked to sales people who were familiar with the open systems space who were very excited about the offering. So we felt pretty good about it. But we still had this entrenched competitor called DLT that was looking at having to cede some ground in the marketplace for us to be successful. And they were very well aware of what we were doing because we were in the public eye. So the anticipation on DLT's standpoint is they started talking about something called Super DLT. They were going to have a new and improved product offering in the marketplace, and it was going to be backwards compatible with their existing DLT products. And, of course, they saw our specs which we had publicly made available that we were going to have a 100 gigabyte cartridge, and their spec was 110. So they had a little more capacity. They had backward compatibility. They had a huge install base of customers. So I recall this one dinner with-- I won't name the analyst. But we had a very nice dinner because he wanted to meet with his various industry consultants and brief them on what you're doing. So hopefully they'll say something nice about you when

you actually do the product announcements. And we're there with Brenda Zawatski who is the current executive and Bob Maness who was the business guy at the time and this consultant. And I don't know John, if you were at this meeting -- I'm not sure who the technical person-- but it was a big meal, a nice restaurant, lots of wine. And the consultant said, you know, you guys you're not going to be really widely accepted. I hate to break the news to you but given all the advantages of the DLT format and now Super DLT. And they've got compatibility, and they're specs look better than yours. You guys are just-- you'll be lucky if LTO gets 10 percent of the market in total. So this is as about a month before we launched the product. This is our good luck and too bad to tell you candidly that I think you guys are not going to be very successful in the marketplace. Of course, I didn't share with him that we had answered that question to our own satisfaction because we'd done all of this market research. And I had spent hundreds of thousands of dollars to come up with a completely different answer. But I wasn't going to give him my answer in that conversation.

Teale: Now, was this an LTO meeting or an IBM meeting?

Johns: This was an IBM meeting, but he was talking about the LTO format in general.

Teale: I just wanted to clarify that because you're going kind of down an IBM centric route, here which is fine.

Johns: Which was true in 2000.

Teale: I'm curious as to how you felt the first time you walked in a room and sat down with HP and Seagate. What was that like because that certainly wasn't our normal way of doing things until that happened.

Johns: First time-- I came in as a marketing person -- as John described. We set up these different virtual teams. So at first, for about one year I was the marketing person. I was being pulled different ways within IBM, which I'll talk about that, but then a couple of years later I was the business lead. So there really were two different introductions. One was to the marketing guys-- I can't remember all of the players I met, but a couple of meetings there, but it was in the middle probably between Gen1 and Gen2.

Teale: Well, we're going to have Bruce talk about the marketing teams, so maybe you could just share some of the business difficulties particularly in the context of all three companies, not just in the internal IBM business case.

Johns: So then a couple of years later about-- it was just before Gen3. It was the year we announced Gen3. So in 2004, I took the business lead role for IBM, which I did for a little over four years at that point then, and then, John asked a good question. It's at this point the virtual company concept was pretty well entrenched, and the people that I was interfacing with from HP and from, at the time it was Certance, because Seagate had spun that operation off. They had been in those roles for quite some time, and John and the IBM teams had been in roles, so there was this virtual team concept that was pretty well entrenched. The biggest challenge was understanding we're in a consensus model. You had to reach 100 percent agreement, which on a business side-- I know there were technical challenges. Sometimes

on the business side, there were also one or the other of us would have a completely different perspective based on our own preferences in terms how we were approaching the marketplace and what our business strategies were. So we had some interesting things, but I think overall, in hindsight, that was one of the key aspects of being successful and being able to keep this collaboration consortium working is the fact that we didn't have voting blocks that we-- we actually forced agreement, and we had a pretty well-defined escalation path at this point, which I found myself using pretty quickly as a business lead, which was if there was a disagreement with any of the sub-teams, the business team was the next point of escalation. So if there was an issue within finance or if a legal issue, the first place of trying to get a resolution was to come to the business team, and we would try to work through the issue and reach agreement.

Teale: The business leader really worked across all of the other groups. You're the only one that did do that.

Johns: It was a great education because there is no real training for how these things work. It really is coop-ition because you would very much have to work together in the room, and you're trying to do the best thing for the LTO format. Then when you left the room and you went back to your real job, you tried to aggressively attack the other players in the marketplace with your offerings versus their offerings. So it was an interesting dichotomy. We all had it, but we were mature enough and we worked together well enough to do that. So we were the first point of escalation, but we would find ourselves at loggerheads on specific issues occasionally usually as a result of changes on the executive teams within one of the companies or change of ownership like the movement from Seagate to Certance involved a certain amount of discussion. It brought in a new cast of players who weren't familiar with this virtual company concept and found this very foreign in terms of how they were going to..

Teale: It was funny that the new people that would enter after we had established the concept always came in with a very parochial point of view, a new boss is in town, it's going to be my way or trail ways, and we would beat the stuffing out of them until they got with the program but that's human nature. You're the new exec. You think you're the new boss, and you find out there is no boss in LTO.

Johns: It was a fascinating experience because, to John's point, when you transition, there would be typically-- the new executive or new whoever it is on whatever team would have a very parochial point of view. It says, "Well, this is what's best for my company, so this is what needs to happen on the LTO consortium." And they were very upset when what they thought was important to have happen was failing to happen and that would result in us going to the executive committee, which was the next point of escalation. If we couldn't resolve it at the business, we'd go to the executive committee, and we'd go to them to reach a resolution and sometimes that didn't happen on the first meeting or the second meeting, and I won't name the specific names, but there were occasions I was in some escalation meetings where we all flew into a hotel in Chicago at the Hilton right there at the airport, so we didn't have to waste any time, and we had meetings and finally the new dissenting executive would just give up in exasperation to the two ones that have been there before..

Teale: Seniority became important.

Johns: It did in a sense that they heard about us needing to collaborate and work together and reach a consensus, but until they actually had to compromise something that they thought was really important, they didn't really understand what..

Teale: It might be enlightening to share maybe some specifics of a couple of escalation examples if you can think of any that you're comfortable with. One that I could throw out is the perpetual inability to agree on the marketing budget. HP constantly wanted to reduce it. IBM constantly wanted to spend it all, and I don't know if that's a good example for you or not.

Johns: Well that's actually a really good example because different companies approached the specifics in terms of the marketing budget in different fashions. Some of them ran it through individual business units. Others kept it at a higher level, so it was noise, and they didn't worry too much about it.

Teale: And just for clarity-- the marketing budget we're talking about has nothing to do with any of the companies as individuals. This is a pot of money that is in the virtual company -- people like Bruce would use to promote the category of LTO. It wasn't promoting a company. It was for the category. So you would think that that would not be a very parochial discussion, but as Brad was about to say, because the accounting was done differently at the different companies-- so for example, we tried to keep the virtual company's assets invisible to IBM. We claimed that we needed a firewall, and I didn't want my executive to know that there was this big pot of money out there. IBM just cashed the royalty checks and that's all they knew about the workings of LTO. HP had integrated LTO into their P&L almost and so to them LTO expense was real, where at IBM it wasn't real. What I don't know is how you guys ever resolved it. I think there was another funny rule of engagement that if we had precedent set and if we couldn't agree to change it, then the precedent stood. So every year it was disagreed to but every year it sort of moved forward.

Johns: Well, it did. It did end up in several executive escalations, and the logic that we would use to try to resolve that is, what's the right thing for the format. We do need to do a certain amount of category marketing. We need to go to some industry shows and have a presence at things like SNIA or other industry-specific shows where you could talk about LTO as a category and then all of the vendors would be there, including the media vendors. The media licensees would be in that booth, so they could talk about their offerings and talk about the category overall. Generally, we would end up maybe not spending as much money as we would ask for, but the marketing teams figured that out, so they would always ask for a little more than they really needed to do. So even though you came down from the original number, it was still enough to do at least things that were viewed essential.

Gardner: Can I ask a question about how this worked? Were the funds given to the virtual company, and then the virtual company spent it or was there an annual agreement that the partners would spend so much in their budgets, and then how did that change when royalties started coming into the virtual company if it changed at all?

Johns: Well, as I recall, we would agree we were going to spend a certain amount of money. A lot of that would go to third-party agencies that were actually contracted by the LTO consortium. So we would have

people help us set up for trade shows. We would have people help with our websites. We'd have white papers that we would have developed to talk about our technology.

Teale: To my recollection, we had our own bank account in the Cayman's if that's accurate.

Gardner: We're just going to ask, we with your IBM hat on or we as the LTO consortium?

Johns: As an LTO.

Teale: I think there was an LTO consortium bank account and that's where everything flowed through.

Johns: There was a consortium bank account controlled by the finance team.

Teale: Finance team, manage team. That was pretty cloak and dagger, as you can imagine.

Gardner: Actually, it sounds pretty open in the sense that apparently there was an annual budget presented to the three parents

Teale: It was really internal. I presented no LTO budget ever internally. The fact that HP's accounting was the way their accounting was was their problem, but I firewalled. So these guys could make any decision they want within the LTO consortia with their own money, and they never had to ask for a dime from the parent.

Gardner: But IBM must have initially contributed a lump sum on an annual basis.

Teale: Well, all three companies subsidized it to get it started, but once the royalties started flowing, there was additional resources. Now don't forget there was a cost of the membership. We don't even know how much detail of the business structure we could discuss, but I will tell you that I had to buy a license to get a copy of my own specification to go do my own product. I can tell you it wasn't cheap. HP had to buy one. Seagate had to buy one. All the media licensees had to buy one. So there were pots of money. Every time we evolved the specification, there would be a new big pot of money. It wasn't royalty driven. Yes, I did get my drive license money out of my own budget back home, but once it was into LTO, I didn't manage it. LTO guys had full reign.

Gardner: So it sounds like before the royalty stream established, the three companies contributed lump sums into the consortium..

Teale: There was a little bit of taking turns picking up the check for a while, and I don't even know how people got reimbursed for that. We would take turns paying for the meeting rooms out of our company budget. All the other expense was viewed as normal business travel...

Johns: It evolved over time. It started off each of us was responsible for our own expenses, but as the program matured, the LTO consortium actually had ability to write our own checks and to enter into our own books. So the business team or finance team developed a budget and reviewed that with the executives and that's where we would get into these discussions where one company runs their budget

through their division marketing budget and other companies chose to keep it off at a higher level, and it was relevant to all the other budgets.

Teale: Just to calibrate you a little bit. First of all, anyone can jump on the internet and say, "I want to buy a mechanism license " or "I want to buy a media license." And I think we even have a third tier of license called, I just want to read your spec, and it was lot cheaper.

Johns: Basically.

Teale: Was that what it's called?

Johns: Yes.

Teale: And you're going to give me a check, and I'll tell you right now that check has quite a few zeroes in it, but the basic license, I think, we made that pretty affordable because there were people who were infrastructure people who needed format specification. They made components that the drive makers needed, and they wanted an early sneak preview of what the next level of specification was going to be to protect their piece of the business. So they'd come in and buy just a document. I just want the document. If you're a media guy, you got a lot more stuff because there's verification that Ed will talk about, a lot more stuff. The point is, there was an ongoing source of income independent of the royalty stream from selling licenses, and we actually argued that our attendance at LTO meetings, they were customer meetings. The media people who bought licenses were my customers. They're paying me to deliver something to them. All right. That's a little bit of a stretch but that was a nice story, and we stuck to it and that's how we got a lot of our travel approved when nobody else could travel because it was the big C, customer travel. That money provided an annual budget and just to calibrate you, it is on the order of three to six hundred thousand dollars for everything, and Bruce had a piece of it, and meetings had a piece of it, and promotions had a piece of it. It wasn't millions, but it wasn't thousands.

Johns: It was enough that it could be problematic.

Teale: It wasn't enough to attract attention.

Gardner: I would bet these license fees are probably on the internet.

Teale: No. No, the reason that I was being a little mealy mouth with you is because Tom will say, "I want the license." And the first thing he'll get is an NDA to sign. Then, you might get something interesting that you're not allowed to disclose. So us sitting here in this room, we don't even know-- I know the exact amount of the mechanism license, but I don't know that that's public, and I'm not at liberty to say what it is because that's an LTO, virtual confidential thing.

Gardner: Can I give you an assignment while you're editing the text of this? Go to the LTO site and see what's public, and add it into this transcript?

Teale: I'm just saying I don't know. If it's public, I'd be happy to talk about it, but remember I said we're all going to use our own judgement and say what we're comfortable saying, and I'm not comfortable

personally saying what those things cost, but I just wanted you to know that this virtual company had some real infrastructure independent of any of the parents.

Gardner: Understood and thank you for the explanation.

[Editor's note: As of March 2016 the license fees were not disclosed on the LTO Consortium's [licensing webpage](#). Three classes of licenses were offered: Tape Cartridge Licenses, Tape Mechanism Licenses and Format Specification Document Participant. License prices are not posted.]

Teale: And that was what Brad was involved in managing. I didn't have visibility to any escalations that were purely driven off of business concerns and marketing.

Johns: Well, one that comes to mind and I won't name all the players, but when the transition from Seagate to Certance, we had new set of executives. Certance was funded by Venture Capital during that time, so they had a different perspective than everybody else did on the consortium. So a lot of the players from Seagate came over with that, so actually the infrastructure that we worked with in the virtual company that people that came over knew it, but the new executives that were running Certance, this was all new to them. There were a lot of cultural practices this whole idea of having unanimous consent and doing the best for the format was all new. So we did have some escalations at that point and then once again when Certance was acquired by Quantum, we went through it again

Teale: Well that one was really painful -- now we're fast forwarding. DLT had waved the white flag. Finally, they said, "It's dead." And I don't know if you remember the Quantum CEO or not, but he was a former VP from Microsoft. I can't remember his name.

Johns: Rick Belluzzo.

Teale: He was a little bit of a crazy guy, but having Quantum buy an LTO company and get the royalties associated with it was painful for everyone involved because I was saying royalties don't go with a transfer of ownership, because I knew for a fact that Quantum bought Certance to get access to LTO. I also knew for a fact that Certance was no longer really developing a product. They were riding HP's back. They were using HP's supply chain. So arguably, they weren't really contributing anymore-- and by the way, unlike the license fees that were nominal, not too big, not too small, again I can't reveal the exact number, but media royalties were huge, millions of cartridges, think even add a couple of bucks a cartridge, which was not the number but gives you scale. So Quantum to be able to just buy a revenue stream that they don't have to do anything to generate, chapped my lips. At any rate, that was escalated and arm wrestled and argued and ultimately..

Johns: It ultimately was resolved, and it was interesting, but we had the same challenges. By that point, we were into the third working on the fourth generation of the product specification for LTO. What was interesting is we were viewed as a success within IBM at this point, but the IBM executive team on the business side kept changing. So what we found is we had this subculture of people that actually understood the LTO consortium and how it worked, and it was viewed as successful. So there was actually a time where I found myself in the unique position at IBM where if you said, "Hey, I need to go do this. I need a check to buy something." there was nobody would disagree with it.

Teale: There was a time when we walked on water as the model worked so effectively, and there were other times when they didn't want it to.

Johns: And the executive turnover was such that every year or two we were getting new people, and they just knew this thing was working; it was successful. So if you said, "hey, I need to go to Las Vegas to meet with the licensees to have an LTO discussion," it was how quickly can I sign it to get it approved because it was working

Teale: Let me show you how effectively this was firewalled and how the extent to which LTO consortia activity was not visible to IBM standard process. I'm going to say the word "Brewers". I'm not going to explain it. I'm going to let Ed explain Brewers later. But there were IP people in IBM who were accountable, of course, for going out and getting IP revenue and reporting it back, and they had their own measurements. You got a little flavor of that from Dan Winarski yesterday. And the storage IP person would go into the big wigs and present a piece of paper, and there was line item on there, and all it said was Brewers and a number. By the way, it happened to be the biggest number on the piece of paper by far. That person didn't even know what the Brewers was. I kid you not. Because I remember being in some of those meetings and they're reporting it, and not only that, the number was so big nobody even cared, and next chart. But that's how firewalled it was in IBM. It was not as firewalled in HP, which led to some of these disputes on budgets and things.

Johns: So I think those were the main highlights.

Teale: That's good. I wanted to go back and touch on a ground rule because the business team was the point focus for when we couldn't agree, when the children are fighting. We had a rule about disclosure of confidential information. We're not allowed to do it. So if I have a IBM piece of what we'll call secret sauce that HP and Seagate aren't privy to that allow me to be comfortable with proposing something in the specification that they're not comfortable with, we actually anticipated that, and we had a ground rule for it that said, and the lawyers helped us find it, that said it is okay to disclose confidential information if it is required to get agreement. I know in Gen2, we had a huge dispute over measly 1 dB of SNR in the media spec. I wanted to just reuse Gen1 media; it was good enough. HP wanted to improve the specification because they felt they couldn't achieve the Gen2 operating point without improving the media little bit. At some point, I had to disclose to HP why IBM thought we didn't need to improve the media specification. They wanted 2 db. And I showed them some of our secret channel sauce stuff that we do a little bit differently from the way they were doing it. They said "ah ha," but we can't do that, and so we ended up compromising at 1 dB, and my media guy initially just sorted his tape, but later on they did improve their tape, but that's another example of how we might handle a specific type of dispute.

Gardner: That would have been handled in the technical TWG?

Teale: Well that was something where we couldn't agree. So what's the process? What's the ground rule? The ground rule is that I can get a one-time get out of jail free, and I can disclose my secret sauce, but in general, we were not allowed to share any confidential information during these discussions. These discussions had some very rigid ground rules around what we can discuss, what we can't discuss, and so it's an elegant concept. We probably set it up so that it was probably the least efficient way to

arrive at the specification, but it turned out that I think it did result in the goal of a best in breed specification that as hence hasn't been demonstrated to have some legs.

Gardner: I was just asking this Gen1, Gen2, the compromise was ultimately resolved at the TWG level.

Teale: We changed the SNR requirement by a dB that I remember.

Gardner: By evoking this get out of free jail card?

Teale: Well, no, the get out of jail didn't work. I proudly showed HP why IBM felt that we didn't need to improve the media, and they said, "Well that's very cool, and you're a real smart guy, but we don't do that."

Gardner: And so then you compromise..

Teale: So we compromised.

Gardner: At a lower number of dB increase.

Teale: The point of that example was not about dBs and SNR and media. The point of the example was the way confidential information was handled, and I don't know Ed. Did we have a separate confidential classification that was LTO confidential or how did we label our documents?

Childers: It was all TPC confidential. We had an NDA agreement between three parties.

Teale: Because we couldn't talk about individual company confidential, yet we were generating confidential documents in the form of the spec, in the form of the business agreements, in the form of various things, and I had forgotten we had a classification called TPC confidential. I don't remember what TPC stood for.

Johns: Technology Providing Companies.

Teale: My bigger point was just the construction of the virtual company and the rules of engagement and the appeals processes was all, I thought, very elegant and enabled this thing without breaking the law.

Johns: Well that was one of the key aspects I think in the creation that the other dimension of this, I think, was we debated this over time was we had third-party testing to make sure you were in compliance, that you complied to the spec, and we had a company-- where were they? Torrance?

Teale: They're in Torrance, California. Ed is going to talk about the compliance in a few minutes, but they're an example of a contract with a service provider, in this case, that was a contract between the service provider and the virtual company, that contract completely invisible inside the three companies just existed out there. So this virtual company really did operate like a virtual company. It had its own bank account. It could initiate contracts. It was lots of things it could do.

Johns: And it provided us a mechanism to say an independent third party validated that the various licensees had in fact complied with the specification, and not that any of the companies probably couldn't have done that themselves, but this cleared any perception that it was anything less than really adhering to the specification that there was no real way to get around it.

Teale: It was a very powerful marketing tool, which is a good segue to marketing.

Gardner: One last question. What mainframe did you work on?

Johns: When I joined IBM, the current portfolio of mainframe products were the 3000 series, so it was the 3033s, 3032s, 3031s, and one data point I still remember on those machines in 1978, one megabyte of memory, one megabyte cost a million dollars. So there were sales reps who made their quota by going from a four megabyte 3033 to an eight megabyte 3033.

Gardner: Actually, I have a second follow up for both you guys. You've used the term Gen1, Gen2, Gen3, I'm assuming those were the code names that lead to LTO, LTO2, LTO3?

Teale: Well, , there was correlation.

Johns: This goes back to the early days when we first kicking around well how are we going to go to market with this thing called LTO? One of the aspects that was really differentiating in the marketplace is providing customer's visibility into the planned road map for the technology.

Teale: And once again that's a marketing tool, so I'd like to get us over...

Johns: But that was something..

Teale: Because if anything, that's something Bruce wanted because that's something DLT didn't have a roadmap.

Gardner: Continuing the LTO discussion, John Teale, is now joined with Bruce Master, John.

Teale: Great. Bruce was our marketing rep. I'll let him introduce himself in a minute. I wanted to just quickly capture something that didn't come out in the business discussion that turned out to be of great interest. I mentioned that licensees paid for a package, and I mentioned royalties. And you might wonder what is the basis of the royalty. And it is not an enumeration of patents that were owned by the three mother companies that we were allowing a licensee, whether it's a drive maker or a media maker, access to. We didn't enumerate it. We simply said we don't know what we're giving you, but if we own it you can have it for the price of the license and the royalty. Unfortunately, in the beginning the three companies didn't own everything we were going to need. And I'll give you a funny anecdote. There's a company out there. I forget their name. Maybe somebody will remember them. They have a patent on one bit-- one bit that says, "The following information has been compressed." Anybody remember the name of that company?

Master: Stack.

Teale: Stack. That wasn't our IP yet. We wanted compression and the specification, and Stack brought it to our attention that we were potentially using their IP. Actually, I don't think Stack did. I think we realized it ourselves and said well, how do we go get the rights to this Stack IP? They're lousy, little, pesky one-bit - not patent. Without waking them up on how big this could become someday. And beyond that it goes into a black hole and gets cloak and dagger. But the way we tended to handle one offs regarding confidential information or other people's IP was one company would get assigned to go fix it. And we typically sent the guy with the shortest pockets. So it tended to be Seagate representing Seagate wanting access to the IP because then LTO could use it. At some point though we did pay off Stack. They didn't know then what they probably know now-- how big this was going to be. I think there were a couple of other examples. I don't remember them. But I wanted to complete the discussion around confidential information. How it was handled, rights-- because this was all above board. We were absolutely paranoid about not breaking the law which was my point. Okay, and we can talk more later in the Q and A if that's of interest because there's a whole backstory on how the royalties got divided. And I won't go there right now.

Gardner: Introducing Bruce Master. He was long-time marketing team lead from IBM to the LTO Consortium. Tell us a little bit about yourself, Bruce.

Master: I started with IBM back in 1977. I'm retired now after 35 1/2 years. And I started in sales. I'm a graduate of Arizona State University in business advertising. Before IBM I was a manager of some GAP stores in the Phoenix area. And at that time IBM was hiring pretty much anyone that could present themselves well. Had a business attitude but could also complete successfully the aptitude test that they would give you. So I had business background with my education And I managed some very high-volume retail stores. And fortunately passed that aptitude test. I started in sales. Unlike Brad, who spoke before me in the data processing division I was in the general systems division which was small to medium size computers. And the first year as a trainee you would spend about four months total off and on in Atlanta, Georgia getting trained on all kinds of concepts-- business concepts for accounts receivable, billing, sales analysis, and more as well as computer concepts. About three-quarters of the way through that year you make a decision to go into the sales area or into the systems engineering area. I went sales. And started the following year on quota and selling computers and did that for several years. I was assigned to Tucson. I moved from Phoenix to Tucson. Spent my whole career after that in Tucson. In Tucson we had the development lab. At that time it was a development and manufacturing lab. And I had an opportunity to go to that site and get some experience in plant life, if you will. After about three years in sales I went to the manufacturing and development site and started in recruiting. I went from recruiting to education, and then to community relations. I enjoyed recruiting, enjoyed education and community relations. It was really a lot of fun. I was there about five at the plant site, when I decided I wanted to add technical experience to my career. I went and interviewed back in the sales field for a systems engineering job with the branch manager. He really wanted me to become a sales rep again. So we made a deal. He said, "All right, I'll let you be a systems engineer if you come back in sales for the remainder of this year." Which was about four months. "I'm going to give you a six months quota. You have four months to work this territory and be the sales rep. And at the end of that time then you can decide if you want to stay in sales or go into systems engineering." I made the six month quota in four months, and at the end of it even though I had made some pretty good money in commissions I really wanted to add that technical aspect

to my resume so I went into systems engineering. It turned out that I was actually in a systems engineering slash sales role because I made a lot of the presentations to clients as well as installing systems throughout that territory. I did that for about five years. At the end of that an opportunity came up again back in the plant environment, for a marketing position. They decided in the manufacturing development lab they needed to do some of their own marketing from that facility. I was intrigued.

Master: Back then I called up the plant site and I asked a contact if they needed a marketing person. I was told that about this new product called a 3494 tape library. This was in late '93, My contact said, "We actually do have a position -- we're going to need to do some marketing of this new tape library. Interested?" This was a Friday. He said, "Why don't you come in on Monday and interview with the manager."

Master: So I planned to interview. even though I knew very little about a 3494 tape library. So that weekend I went on the web and learned a bit more about it.

Teale: We had the web in '93? A little bit?

Master: We had a little bit of info at the IBM website then. And I put together a presentation on the 3494 and a plan to market it. At the interview the manager explained what he envisioned, what he wanted to achieve. And then I showed him my laptop presentation. I said, "Here's some ideas I have. I think they coincide with your objectives and some additional opportunities." He went through it and said, "This is really good. I really like this." I said, "Will you hire me?" I guess he never had somebody be so bold, I think he remembered that I was coming from a sales environment. I was going for the throat, so to speak, trying to close the deal. And he said, "Yes, I will." We closed the deal right then. It opened up a wonderful marketing career for me. I was in marketing The for the last 15 years of my career with IBM—I learned a lot and worked with some great people, these gentlemen here included.

Teale: Was the 3494 the big, yellow robot or was that our first real library? I don't remember.

Master: . The 3494 was the smaller version of automation which grew from 1 to 16 frames.

Teale: Okay.

Master: And the last 11 years of those 15 years was with the LTO Consortium. So I joined the LTO Consortium around Generation 2 of the LTO products. Then I retired from IBM in 2013 after 35 and a half years-- I'll get back to the whole LTO Consortium part of my career and marketing with the LTO Consortium in a moment.

About two and a half years ago I retired from IBM and since retiring I am active in three roles. One does not have anything to do with technology or storage but two do. So one of those roles is I teach golf. I've been playing golf for a couple of decades. The plan that I'm on, and many of us are on, is that when you leave they give you a little bucket of money that you can use to get retrained in a new career. So I immediately called up the redemption center and said "Hey, can I go on a golf vacation and use some of this money for a golf vacation?" They said, "No. But what some have done is get certified to be a golf instructor." I thought, "Really? That sounds like fun." So I did some more research on the web, flew off to

Chicago about one week later, and I got certified as a United States Golf Teacher's Federation golf instructor. Came back and I've been teaching a lot of golf since then. I've had about 50 golf students and have conducted about 20 clinics. I love working with golfers new and experienced. There are two other things that I am doing related to tape technology that just sort of came my way which I'm really having a lot of fun with: One is managing business shows for the LTO Consortium. One or two shows a year managing an LTO Consortium booth at the business shows. And the other one is writing the blog for the LTO Consortium at the LTO.org website. I've written about 14 blogs, and it's really been a lot of fun. On various subjects but they all revolve around storage and being efficient with storage and, of course, it has a large element of LTO technology involved with it.

Teale: Not to be too nose-y but do you get compensated for any of that?

Master: I do. I do get compensated

Teale: Kind of like the Karl Brueckner type of arrangement or something. Karl was hired by LTO as a contractor to interface with the independent compliance company.

Master: The business shows are actually managed by a third-party company which John mentioned early on. There are various third-parties that help the Consortium get done what they need to get done whether it's through the marketing team or the business team.

Gardner: Are you getting paid to advertise Ultrium right now?

Master: I am not paid to do advertising at all.

Gardner: I'm focusing on your shirt there.

Master: No. But I do have a lot of these golf shirts or polo shorts--

Gardner: It's called a free shirt is what it's called.

Master: The team-- and I'll talk more about the actual activities of the marketing team in a minute. But one of the things that they do is go to business shows and have a booth which I mentioned where they can demonstrate the products and talk about the technology. They like to dress alike as a team so people can recognize the representatives that are at the booth that can respond to questions and demonstrate the technology. So they all wear the same shirts, and this is one of those that I wore when I used to be on the team.

Teale: And Ultrium is your company identity?

Master: Right. It's an LTO technology brand not a company thing.

Gardner: And Ultrium is related to LTO as?

Master: Oh, John explained a little earlier, but I'll elaborate a little bit. There originally were two formats of the linear tape open technology. One of those formats was Accelis that never made it through fruition to

become a product from companies, but it was a specification. Second specification was Ultrium which you can see on this shirt-- Ultrium LTO.

Gardner: I did a zoom in on the shirt when I asked the question.

Master: Ultrium is the format of the cartridge. So it's the shape of the cartridge with the single spool and how the data is written

Gardner: So that is an Ultrium--

Master: This tape cartridge is Ultrium.

Gardner: -- and that is what an Accelis would have looked like?

Master: Yes. It would look very similar to this Accelis, but that didn't make it to market because customer demand kept demanding the Ultrium high-capacity cartridge.

Gardner: Do you have children?

Master: I do. I have one son. U of A grad at the Eller College of Business, and he's 27 1/2 and a great guy. He's got his own home now and a great job. So he's doing really well. Very happy about that. I'm proud of him.

Gardner: Did you make the 100 percent club as a salesman?

Master: I did. I had a couple of 100 percent clubs, and a couple of-- systems engineering symposia. The sales part of my career was really a lot of fun. Systems engineering was quite an experience. I really enjoyed that too.

Gardner: Any fun stories from 100 percent club keeping it short?

Master: I don't know if this is really a funny story. back then when you went to the 100 percent club there was not only a lot of fun with all of your peers in the industry but they would have some really great entertainment. One of those experiences was Kenny Loggins. And it was in a theater in the round. So every seat was really a good seat, and he really put on a show. I wasn't a huge Kenny Loggins fan then but I became one. And he really put on a great show, talked a lot to the audience and had tailored jokes that dealt with sales and IBM. It was quite entertaining .

Gardner: Is it true that system engineers really ought to be compensated because they frequently are more responsible for the sale than the salesman is?

Master: Well, I've seen both sides. Sales-- the pressure's really on you in the sales environment to make a sale. We would ring, the bell when we made a sale. Back then they actually did have a bell in the bullpen. You probably remember that, Brad. And it was a lot of fun but the next day you're asked, what did you do for me today? So there is a bit of pressure but it's a good motivating pressure. In systems engineering there was the pressure of actually meeting deadlines to get products installed, and make

sure that they're installed properly and the customer's really happy. Customers trusted you a lot in systems engineering. Whether you were a field engineer actually repairing products or in systems engineering installing and making them work. So because they trusted you so much they would ask you a lot of questions that were really more sales oriented questions, but they wanted to hear it from you about will this product do this or that. I found myself doing some of the salesy type things in systems engineering because of that. And I enjoyed that too. Customers were great and really enjoyable to work with.. Writing this blog has really been a lot of fun. I wasn't sure I wanted to do it. There was a bid that came out from the Consortium looking for someone to write their blog for them. That brings us up to date as far as some aspects of my career and what I'm doing in my retirement

Teale: Is this LTO.org?

Master: Yes, www.lto.org, same website The blog is actually called [BLOGBYTES](#) if you go to LTO.org, click on the news drop down button it'll show BLOGBYTES. There was a bid that came out from the Consortium looking for someone to write their blog for them, but I wasn't sure I really wanted to do it. At first I thought do I really want to put my head back in that world again because my head right now is in the golf world. But I thought it does sound challenging. My wife strongly encouraged me. She said, "Do you realize this is the kind of job you can do from anywhere? All you need is a connection to the internet and to a word processor and just start keying away." And it has worked out that way. I find myself sitting in a Starbuck's, and I'll be typing away working on the same blog, editing it, looking at it. And then I might go up to a vacation up to Sedona, and I'll be sitting in the room. I'll get up early, and I'll start working on that blog again. So it's really been a lot of fun to do that.

Gardner: Are you making minimum wages? An hourly rate when you put all that time in?

Master: It's more by the blog.

Gardner: Turning back now into LTO I think John has some provocative things he'd like to discuss with you.

Teale: You did a large variety of activities, and just kind of a broad brush what types of things. There was publishing stuff. There was other stuff. There were the business shows and all the goodies, golf balls, t-shirt, etc.

Master: Exactly. That's right there are golf balls.

Teale: We had a tee time this afternoon until you derailed us, Tom.

Gardner: I did cancel that.

Teale: We were going to run this train on time, and we were going to be on the golf course in about 20 minutes.

Master: Quite right. The ultimate goal of the marketing team was to perform activities that would invoke demand generation. I entered the LTO Consortium on the marketing team around Gen 2. And the goal

then was not only to create or to influence demand generation but more basic than that to increase the awareness of LTO technology and to get brand name recognition since the technology was still in its infancy.

Teale: If I could interject real quick. Just because he says Gen 2 doesn't mean that LTO had any mindshare yet. In fact, the Gen 1 products from all three companies were a little on the lame side because of the-- there was a lot of engineering involved. And there was a lot of pressure to get to market. We were really racing with each other. Gen 1 did not have a whole lot of penetration into DLT-- barely even got their attention.

Master: If I may take a liberty on the word lame. I think what John means is that the actual penetration or acceptance by customers wasn't real large yet. And our products were great but the penetration wasn't.

Teale: I'd go so far as to say that the marketing team that was the original marketing team really didn't do much. Believe it or not it was a guy named Dave Graves who subsequently retired and left, and he wasn't even a marketing guy. So Bruce really-- even though he says Jim too, was the ground floor of the real serious push.

Master: Still pretty much in its infancy. Back then we did a lot of things carrying on with what the first team had started, but expanded it to get more brand awareness. We promoted that this is reliable technology with a goal to penetrate the marketplace. So we did a lot of press interviews, a lot of press releases. We did a lot of analyst interviews to educate the analyst community which was very influential in the storage technology marketplace. We wanted to get both press and analysts [like Gartner and IDC] to write stories and white papers about the technology. Hopefully their clients would read about that information. We would publicize the articles on the website, at business shows. We were continually trying to get the brand awareness and mindshare for LTO technology. We expanded making the website even bigger, with pushes to get clients to the website. To do that we would advertise the website. We started an advertising campaign as part of our budget. We would go to business shows and start collecting names so we could have those in a database for email blasts. And then the email blast would promote the white papers, and the news articles that were published. And it just started to grow larger and larger until we had tens of thousands of names in this database of customers and potential customers that we could market to. The marketing activities expanded even more to where it is today with press releases, white papers, business shows, analysts and press relations, the website, management communications, business conferences and expos around the world-- this is a worldwide marketing effort not just the states but in Europe and the Asia Pacific region. It now includes social media marketing. I mentioned the blog, but the LTO Consortium currently has a Twitter account and a LinkedIn account and actively markets with those as well. And apparently the blog is going quite well. I was just looking at some stats on the blog page. And I was happy to see that hundreds of views of those blog pages have taken place for each individual blog. In fact, they even ranked them the top 10 blogs, and it's interesting for me to see which ones they were reading the most out of things I've written.

Gardner: What was it? Do you have what they've read the most? What do they like?

Master: The LTFS blogs were at the top of the readers list, which by the way is the logo that is on this sleeve [points to T shirt]. Linear Tape File System is another part of the technology that was developed originally by IBM then adopted by the Consortium. Another one high on the list was a blog about how LTO technology is very low cost although highly reliable storage. I'm currently writing one which will be fun to see how the audience reacts to since we're near Halloween. This one will feature what I'm calling data goblins that can cyber-attack organizations and steal information and steal IP information and cause havoc in an organization; it will be interesting to see how this one is accepted. So lots of activities as you mentioned, John. Some things that are more behind the scenes that the marketing team does is to organize meetings with all of the licensees which includes not only the three technology provider companies but also the media licensees. At those meetings we present what the marketing team has been doing to help promote the technology. It is also an opportunity for the technical and business teams to present information to the whole group-- to all of the licensees. The meetings are held about three or four times a year.

Gardner: We are after all a history museum, and so a lot of our interest is in marketing challenges and campaigns background to what you would say Gen 2 or Gen 3.

Master: Sure. The early on challenges?

Gardner: Yes

Master: Very good question. Again, the very initial part of the marketing campaigns were to increase brand awareness. But very quickly we found resistance from the largest part of the market which was owned by the DLT products. So it was a big challenge to educate potential customers on what LTO Ultrium technology is, and how it is differentiated, the features and benefits of the technology versus the leader in the marketplace at that time which again was DLT. That was a big challenge, and we had to really put forth some very positive comparisons without disparaging other products out there. Positive comparisons to show the differences both not only on a technical level but also on a usage level and an economics level.

Teale: A couple of big levers you had was the openness because that was something that you could talk about without saying anything about that monopolist over there.

Master: Definitely.

Teale: And the existence of the roadmap having credibility.

Master: And that there were multiple providers of the technology was a really big benefit, and a lot of clients eyes would widen when they heard that-- the openness related to it. And, by the way, you can obtain LTO products from a variety of providers, and they're all interchangeable and compatible.

Teale: Customer choice. And you can switch any time you want.

Master: Yes. You can take a cartridge from Hewlett Packard and put that in a Quantum drive or an IBM drive and it works.. So that was a very attractive attribute of the technology.

Gardner: As opposed to the Super DLT and DLT which had less availability and less interchangeability?

Master: Correct. There was one provider and once you chose that you were locked into that one provider.

Teale: Give you an example of some of the images we had. I don't know if it was marketed. We certainly did have a modest media royalty, and we justified that in part too because Bruce was doing what he was doing. He's doing this on behalf of everyone. Everyone benefits by growing the LTO piece of the pie. However, the DLT media royalty was onerous because they were single sole proprietors, sole owner of the IP. They were disadvantaged from a media cost point of view initially. I don't know that they ever changed it or changed their business model but I know that they're-- the media companies that tried to participate with DLT really chafed at the unreasonableness of the--

Gardner: For what it's worth, Quantum licensed second and third sources beginning 1998.
[Editor's note: Source: Quantum 10-K, March 31, 2000]

Teale: At what cost though? See we don't know.

Gardner: It's always the issue.

Teale: The rumor has that it was-- there was a number of things that ultimately conspired to help Bruce win the day, but those were kind of the basic tenants.

Master: It really was very much a team effort. So in each team-- and John mentioned these teams before. There's a legal team, an executive team, business, technical, marketing, etcetera. We work as a team. The representative from HP, the representative from Quantum and myself would work closely together to develop the marketing plans, tactics and initiatives. We would present the plan initially to the business team. Then after working with the business team we would present the plan to the executive team to get the plans approved and then for the next twelve months we would execute the plans.

Teale: Another fun thing that the marketing team did is they worked to keep the message fresh and relevant and up to date. It wasn't just okay, here's the brochure we printed five years ago. So it was really fun bringing Bruce back to talk to the development team. It was really good for the development team's morale to get visibility to the message of the day. For example, I remember when we put encryption in LTO. You guys had a lot of fun with that. You had some funny themed commercials. I don't know if you ever followed through with those or not.

Master: Well, it was a lot of fun to do the advertising part.

Teale: The guy going through the security checkpoint or something like that. I don't remember the details. My point is as the world changed all of a sudden there was a sensitivity to tapes falling off a truck or tapes getting stolen-- boom, encryption-- boom, whole new thing to go market.

Master: I'll describe that ad. it showed a business man holding onto his briefcase going through the scanner at an airport. And he wanted to protect this information.

Teale: He's lying on the conveyer belt to go through with it.

Master: Right, to go with it in the scanner. It said you don't need to do that all you need to do is encrypt it, and that got into the encryption aspect of the technology for data security..

Teale: Did that ever see the light of day?

Master: Oh, yes, that was a great advertisement. The initial ad as I think back on it now-- was a young girl roller skating that had a pillow tied to her backside. And the tagline was "everybody needs good backup." That was a pretty good one. I was not involved in this one; it happened with Gen 1, but it carried on, I think, a little bit into Gen 2. We had lots of ads. At least one main theme each year and sometimes it changed mid-year. That was a very creative and fun part of the marketing team responsibilities.

Gardner: Some more memorable themes?

Master: Yes. We had one with a catch line, "are you running out of storage space?" --. This one showed a fellow that was standing in his shorts and high-top socks that had a garter holding them up like the old days. He was trying to force a huge bundle of clothing into a commercial sized washing machine with much of it falling out he couldn't get it all in there. So that was the visual of "need more space." It talked about how much capacity an LTO cartridge can hold.

Gardner: That was against DLT at that time?

Teale: No. No, not against DLT. I think as Bruce stated that we were just trying to establish ourselves.

Master: We never mentioned DLT in an ad.

Teale: We didn't even want to acknowledge that they existed.

Gardner: But you were positioning yourself against DLT positively?

Master: Definitely. DLT was a great product and the key competition back then, but LTO had a big advantage as an open standard, and that's what we were really trying to promote-- and, of course, the capacities that were achieving with each generation and the speeds of the drive too. So that was initial objective-- to get awareness of the brand. LTO technology garnered more and more market share from zero to let's say mid-way through Gen 4 above 50 percent of the tape drive market shipment share. Eventually the competition became another storage technology which was disk technology and removable disk drives. That became a new challenge. Customers instead of choosing just the category of tape were choosing disk technology. So it wasn't LTO technology versus another tape technology. It was all tape as a category versus disk.

Teale: So an example of that is what they call virtual appliances. It became extremely prevalent in the industry for a period of time. I don't know if it still is where you could take a bunch of disks, put them in a box and put a tape personality on that box so that it responds like a tape drive to existing applications, etcetera. But it looks like a fast one.

Master: Because most of those applications were written to tape.

Teale: So they're just putting a personality on the front end that says I'm a tape drive talk to me like I'm a tape library, and I'm going to act like one. And then they were able over time to achieve cost points that were starting to threaten us. And there was the emergence of something called data deduplication which was a huge enablement to make these little virtual-- just a bunch of disks that would be competitive with a real tape library where they could remove-- they can insert pointers in places-- in place of large-- so let's say people are backing up. And they're backing up the same Microsoft presentation file, and it's evolving. And it's evolving maybe two sentences at a time. Why are you saving the whole damn thing every time you are protecting your information? Why don't you just take the big common chunks out put a pointer there? And so that was-- believe it or not that was as a 15x capacity improvement for a bunch of disks. You say well, why didn't we do it in tape? You can't. The reason you can't is because when you're trying to reassemble the data you got a tape in there, you're reading it and it says okay, go get that chunk over there and that might be on a different tape cartridge. So anyway I go get that chunk and insert it. And then you got to put it back and come-- the point is you can do it, but it's ugly. So that was an example of an emerging threat. When I left the business I actually bought a company that does that.

Gardner: Oh is that right?

Teale: Diligent. That was my acquisition before I left the business. Sorry, I don't know what today's environment looks like.

Gardner: How did you positively market your technology as opposed to data deduplicated disks?

Master: The concept of storing information on disk is a really good one to get fast access to that data. However, during the first 30 days only 5 to 10 percent of it is accessed again. At least 80 to 90 percent is never looked at again. And disk is inherently more costly than tape, typically by an order of magnitude. Therefore, the stagnant data should be moved to tape to reduce costs.

[Editor's note: As of late 2016 tape media has a about a 40% price per terabyte advantage over disk drives². A tape subsystem (drives plus automation) has a substantially higher initial cost than a disk subsystem (RAID chassis) so that tape is only less expensive than disk at some large storage capacity. Operating costs of tape are also generally lower and also contribute to tape's price advantage at large storage capacities.]

Teale: It's dollars per gigabyte of total storage cost.

Master: Right.

Teale: It's called Total cost of operation, TCO. Also, disks spin, there have been efforts made to turn off disks that are not in demand and restart them, but that leads to a whole bunch of reliability issues with most disks. And one thing that Bruce made great hay of was the greenness to tape.

² Based on Internet pricing of 6 TB LTO 7 tape at about \$115 versus 6 TB HDD at about \$190.

Teale: Because when tape is at rest, it uses zero watts.

Master: To store all that information, the 80 to 90 percent that's not accessed again on a more costly real estate... disk technology is a waste of that technology. Disk [HDD or SSD] is important when you want to get the data fast but it's not economic in the mid-range marketplace for large amounts of data not frequently accessed.

Teale: So in the high end, we addressed that issue very, very neatly, but for some reason we've never addressed it as neatly in the low end, and that we have what we call a disk cached tape library. So it is a real tape library. There is a real disk cache. But the disk cache is presenting a tape personality to the host. And what we do is we watched that happen.

Master: You're jumping ahead.

Teale: Oh, okay.

Master: <Laughs> But that is where it's going to evolve.

Teale: And that is the question, "Well, what's been talking you so long? They've been doing that in the enterprise since the '70s." The point is that you can reclaim disk space. You can identify dead data, and move it onto tape, and then reclaim the disk space. Because the other thing Bruce hasn't mentioned yet is scalability.

Master: Yes, but I'm getting to those.

Teale: All right, I'll leave you alone. I'll go take a break.

Master: Having data on that more expensive real estate is a waste of that real estate. But you do need it, to get to the information fast. So that's one part of it, making use of that valuable real estate, both real estates, is the economic part of it. A strategic blend of disk and tape is preferable. The second part of it is that information on disk is at higher risk of data loss than it is on tape, because tape is inherently an offline technology. If you take a cartridge out of a tape drive, it's no longer connected, it has no more electrical connection to the system. And if it doesn't have an electrical connection to the system, then viruses, sabotage, disgruntled employees cannot get at it. Whereas disk is always connected to the system, so a virus, sabotage, disgruntled employee, somebody who knows the passwords can get at it, can destroy it, or copy it and steal it, especially if it's valuable information, like IP, or customer addresses, social security numbers, et cetera. And there's a lot of that. In fact, that blog that I mentioned earlier, talks about that. Having some form of storage that is offline, segregated from system attachment is a necessity. The advent of removable disk played a role in that, but removable disk had less than desirable reliability. Just moving it around can cause the disk to be inaccessible. Tape technology fills the role for offline storage quite well.

Teale: I'll tell you a funny story, especially if you're doing it with a consumer technology, and you're probably talking to some degree about consumer disk. Facebook, when it got too big, and the concept of big data, the term was kind of invented for these social media companies that are basically keeping track

of all my bowel movements, it seems like. Facebook thought-- they'd never heard of tape, they were a bunch of kids. Does tape still exist, really? They put together a kluge library, with rewritable Blu-ray DVD as their big data repository. Big refrigerator, just buy the DVDs at Office Depot. So this is media that's not in cartridge, it's not protected in any way. One of my consulting things, I mentioned some venture capitalists, so I was giving a venture capitalist some advice on investing in a Facebook developed consumer based storage. technology. So I gave him long lectures on why you don't want to do that, why you don't want to base any real business-critical mission to a consumer technology, and there's a bunch of reasons for it. It's not hard, and basically it wasn't what he wanted to hear, and he kind of, in a very polite way, called me an idiot, and thanked me for my time. Facebook actually published a picture of their brainchild, that they've reinvented after hundreds of people already invented it. Guess what? Facebook came begging to IBM for tape about a year after that.

Master: Well, it's an interesting story. I'll talk about another one that shows the sabotage that can take place on disk technology, Google, and this is their story, their VP of their information technology talked about this, that [in 2011] they had a system error. So this wasn't sabotage, but it was just a systems error, that caused them to lose [at least 20,000] Gmail accounts^{3,4}. And this glitch affected the backup, also. The backup process was disk to disk backup. So not only was the primary copy of the Gmail accounts destroyed by this systems error, but it affected the backup, it just replicated itself to the backup. Fortunately, using LTO tape, they had additional backups that weren't affected by the systems glitch, and they were able to recover everything in a few days [6:00 pm Feb 27 thru 10:55 pm Mar 1]⁵. So what we would do, we would find out these stories to help publicize these value points, or benefits of using tape. Doesn't mean you don't use disk. You do need disk, for those very fast inquiries. But for the original copy of all the data, put it on tape and remove it from disk, free up that disk space for new data, for the economics, the security of offline tape storage, and because tape has a very green footprint. A cartridge on a shelf draws no power, so it's totally green. Even a tape drive, when idle, draws very little power. Even when it's active, with a cartridge in it, the tape drive uses far less the amount of power that's used from disk technology. Some studies have shown that tape is 200 to 300 times less energy-consuming than mid-range disk technology. That alone can be enough to justify a complete purchase of tape technology, versus the similar amount of petabytes in disk technology.

Master: Just the green part of it. Just the green savings alone will justify a purchase in a complete tape system.

Gardner: Just as an ex-hard disk drive guy, this is fascinating, and I'd like to go on for hours, but John won't let me. So if we could back up, I think you mentioned somewhere around Gen 4 or 5, LTO had achieved about a 40 percent penetration.

³ [“How did Google lose, and find, all those e-mails?”](#) CNN, May 1, 2011

⁴ [“Gmail back soon for everyone,”](#) Official Gmail Blog, Feb 27, 2011

⁵ Ibid.

Master: Between gens three and four, it was over 50 percent penetration, and if memory serves me right, against other tape technologies in that midrange tape category, by Gen 5, the LTO market share was over 90 percent. This is according to analysts' data.

Gardner: Of course. And one of the features you suggested was in part, responsible for that success was the open technology, leading to lower prices and multiple choices. The benefits to the customer are obviously-- lower prices, opportunity to move. What are some of the other features and benefits that you used to drive that improvement in market share, that you recall thinking of? And getting back to the campaigns -- the features and benefits of LTO that drove that incredible success, and market improvement.

Master: Sure. Well, I mentioned the economics, that was key, because that's the first thing on the checklist of a CFO.

Teale: Established road map credibility over time, that wasn't there on Gen 2.

Master: LTO technology was the first published tape roadmap. Currently, it goes up to ten generations. And the LTO consortia announced the specs for Gen 7. Having that roadmap two or three generations out really gives customers confidence that this is an ongoing technology that can be relied on for many years to come. , There's typically about two or three years between generations.

Teale: Also, the roadmap has some compatibility interchange requirements between generations that were probably a little better than some of the competitors.

Master: That's right, the paradigm is still holding true that a Gen 6 drive today can read and write a Gen 6 cartridge, read and write a Gen 5 cartridge, and read a Gen 4 cartridge.

Teale: So let's say somebody wants to move to the new technology, let's say it's faster, for example, just happens to be faster, but they're not ready to rebuy 100,000 cartridges, they want to keep milking that investment until they have to buy new cartridges. This protects them, and holds their hand all the way through that decisions process. So the drive technology, which is capital investment, media technology tends to be expense, so you've got your machines, and you've got your paper, right? Those can proceed in different financial cycles.

Master: So yes, this compatibility between generations is another key selling point, to answer your question.

Teale: I believe super-DLT, when they announced their intentions to do it, I'm pretty sure they were going to break compatibility with the prior DLT⁶, because the technology made them. I don't think that they were

⁶ Quantum promised backwards compatibility to DLT but their first SDLT drives that shipped in 2001 did not initially have it. Subsequently backward compatibility was added. Source: "Tape Drive Technology Comparison," Spectra Logic, November 2001, <http://brandyfetzner.com/websites/lynx/TapedriveTechnologyComparison.pdf>

real vocal about it, but I think we knew about it, and I think we used the information. That's a big deal, when you ask customers to buy new media and new technology.

Gardner: Going back in tape time to reel to reel, that was always a requirement, I think from the 727, perhaps, right through to the 3420. A big part of tape culture.

Teale: The difference being that the first 40 years of tape accomplished two orders of magnitude of moving the football, and then then the last 15 years of tape I was in was then four more orders of magnitude. So I mean, being able to read a seven-track 1957 tape wasn't as big a challenge as it would be today. In fact, if we had to read a 3480 tape, I'm not sure we'd know what to do, at this point.

Master: Well, some features that were, to go along with your question about what else made it so popular, one feature that was added in Gen 4 was WORM technology, which means Write Once, Read Many. And it was a different cartridge. You could tell just visually looking at it, because the bottom half of the cartridge is gray. And if data was written to that cartridge, it couldn't be changed, couldn't be modified, couldn't be erased. For some industries, like insurance, the legal profession or government regulation, the data written has to be unalterable. LTO WORM really served a great purpose.

Teale: Is that what this does? [points to a switch on the cartridge.]

Master: That's another physical way to protect a cartridge from being overwritten.

Teale: But in other words, in the WORM version, this doesn't exist as an option. It's just hard code.

Master: Yes, that one can be turned on and off. In a worm version, you can't change the data.

Gardner: Is it the medium that is not changeable, or is this some sort of electronics?

Teale: It's electronics. It's electronic recognition that the customer, by moving this switch-- and there's something that will sense that, when you load the cartridge, does not want this cartridge to be overwritten under any circumstances.

Master: Yeah, but the WORM cartridge, whether the switch is moved or not, you can't alter the data on the cartridge.

Teale: Okay, let's add one minor piece of detail, and then I'll let you tell them what a great marketing success this was. Bruce knows something uncomfortable is about to happen. One of the early requirements-- Sony hadn't gone away yet, they hadn't given up, and they announced something called Super AIT. They were one of those eight millimeter, two wheel guys like Exabyte. In fact, they were providing Exabyte with Exabyte's media. Super AIT, what was super about AIT? They took a little bitty piece of semi-conductor memory, stuck it in the cartridge, put some information in that memory, like, "Hi,



I'm a Super AIT, and my volume number is so-and-so, and my--" whatever is in there, it could tell you if you really want to know. And it turns out it was not a contactless access. When you loaded the cartridge, they actually had to reach in and make an electrical connection to read that information, and utilize that information. From an engineering perspective, we thought that was a worthless requirement, it's superfluous, it's redundant, it's unnecessary, it's bullshit. But the marketing people said, "Well, how are we going to be able to beat Super AIT? We don't have cartridge memory." So we put cartridge memory in, and in order to avoid all the Sony stuff, we put a contactless version in, so that you didn't have to establish an electrical connection. I think it's an infrared connection, I don't recall.

Master: It's RFID. , it's an interesting part of the concept-- we really didn't market it unless asked.

Teale: We could never figure out what to with it.

Master: Because most customers wanted to know what the features and benefits were, they didn't really want to know, " how the gasoline engine works inside?" Most wanted to know, the key features and benefits, and the technology overview. But sometimes we had to get into a deeper technical discussion, for instance at a business conference, making a presentation, there may be one or two people in the audience that are pretty tech-heavy and want to know those kind of details, so then we would talk about it. It was interesting, the cartridge memory proved to help a lot in implementing some of these technologies like WORM, and the next one which is encryption.

Teale: Sure, so that's how you tell.

Master: And encryption actually came in Gen 4

Teale: Encryption went into Jaguar [IBM internal name for IBM 3592 tape drive, IBM enterprise version of LTO, announced September 2003] first.

Master: And encryption is very helpful, as I mentioned a bit earlier, for securing and protecting data. If you're going to take a cartridge, fill it with data, and put it in a FedEx envelope and ship it across the country, if that gets in the wrong hands, you don't want it to be readable., You want to protect that sensitive information, whether it's employee information, customer information, IP, et cetera. It is easy to encrypt it at the drive level with LTO technology, and it doesn't impede the speed of the drive.

Teale: And it's another differentiator of tape from lots of other things, including CPUs, disk drives, because encryption is a performance hit, when you're doing real-time data processing.

Master: On other systems.

Teale: And there's been resistance to the total adoption of encryption to protect the whole chain of data, chain. But we had literal incidences of tapes falling off a truck, and some poor company's got to buy 100,000 free credit reports for their customers. And so that was a huge thing that was somewhat-- at the time, when I left the business, it was pretty much, tape was the only people encrypting data in-situ.

Master: Yes, and again, it had no impact on the performance of the tape system.

Gardner: So it sounds like capacity was more or less parity during that transition. You were a little bit ahead, they'd announced, probably like ten percent higher, but capacity wasn't really an issue, nor were many of the speeds and feeds issues, data rate-- it was fundamentally successfully driven, by the three things?

Master: No, it was still very important to have jumps in capacity, jumps in speed as the capacity increased, you still needed to have enough speed to go along with it, to get the job done in a particular period of time, a reasonable amount of time, and to compete against either other tape formats-- although they were dwindling by that time, during Gen 4 and Gen 5, but also to compete against the disk technology, and to show that tape is not the bottleneck in the scheme of moving data around an environment. The bottleneck is typically the disk, or the application, but not the tape system. So it was important to have speed in the tape system as well.

Teale: One of the things that is fundamentally true about tape, is if you're just streaming data, if it's just quarter horse, and not thoroughbred, it's just ears back, straight down the track, it's hard to beat tape performance, you just can't beat it.

Gardner: Well, it's impossible to beat tape, because tape, in that application, is really, really good.

Master: Disk evolved to get away from-- during this time, probably in Gen 4 and Gen 5, evolved in getting away from using tape-written applications to do backup to disk.

Teale: That's when they started disks writing their own.

Master: Yes, just disk backup, just using operating system technology, and then some applications so they could track where data is. Is this disk system onsite, offsite? Where did we replicate to? The advent of data deduplication came into effect, where using the same disk real estate, you can get 10 to 15 times the amount of data on the disk, by reducing the data, dedupe is a data reduction technology.

Teale: Deduping the data, yeah.

Master: So we're competing with that advent, also. And to compete against dedupe, we showed how the tape system is still less expensive than even deduped data and more secure. So the price point was still very important, and was a key part of the competitive information and the customers decision criteria.

Teale: And it kind of depends on the data. It's very data-dependent, just like compression is. So a tremendous example of redundant data would be eternally-forwarded e-mails to hundreds of people. So there's one original node in there. And so for example, an e-mail repository looks pretty good on a deduped disk.

Master: Some data is looked at once and backed up. it cannot be deduped, because dedupe needs to see multiple iterations of that data to be able to reduce it, to take out the redundant info.

Teale: But sonar data doesn't even compress. Noise doesn't even compress.

Masters: Videos, such information doesn't reduce.

Gardner: So switching a little bit from your LTO hat to now your IBM hat, sort of wrapping up, without revealing anything proprietary, how did you compete with the IBM tape deck against your partners in the consortium? What were some of the features and benefits that you publically asserted?

Master: It's an interesting question. Do we want to get into that, that discussion?

Teale: How did we differentiate from HP? I think that's another discussion another day, but let me tell you something that I think I'm comfortable telling you. IBM-- earlier I talked about a fun story about how we're going to divide royalties, way back in the beginning. Okay, we're assuming they'll come. If we sell it, we hope they'll buy it, and I'm talking about licensees now. HP, Seagate, and IBM all came into that discussion with the perspective of their go-to-market capability, their channel.

And without discussing it, IBM knew that we were largely retail, we were largely IBM server-oriented, kind of knew what our volumes would be. We figured we'd be the lowest-volume participant, was what we're thinking at the time. I'm talking pre-Gen 1. So, we came in saying, "Well, we think the royalties should be split according to how much IP was contributed by the individual company." Of course, I had a boatload, because we'd invented track following servo and they'd never even done it before. We had a boatload.

Well, this, as you can imagine, sat really well with Seagate and HP. Now, Seagate was primarily an OEM company. They had some pretty big OEM customers. Brad mentioned some earlier. One of the ones you didn't mention was Dell, very high volume. So Seagate proposed that the royalties be split according to drove the most volume of LTO. HP came into that conversation, and didn't really have any channel advantage, or IP advantage, and they said one third, one third, one third. Every single cartridge has a tax to LTO on it. And I told you I wouldn't tell you the number, but I will tell you, it's not divisible by three. What the hell is that all about? Why are you dividing it by three, and it's not-- so I was going to talk about a channel dynamic that actually, he's marketing the categories, doing a great job.

Now, what about the war between IBM, HP, Seagate? What I can tell you is a lady named Carly Fiorina, in between Gen 1 and Gen 2, bought Compaq. HP acquired Compaq. Compaq was Dell's mortal enemy. Dell had been buying all their LTO from HP. Michael Dell publically said, "No new business with HP." And we were trying to get into the OEM channel, but were just no good at it. IBM, I think Brad already described, it was a new thing for tape. They'd been doing it in disk, but we'd never done it in tape. We didn't know how to do it, we didn't have an infrastructure for supporting those customers. We had to learn it all. So we had a little traction in Gen 1, but not very much. Few little small library makers, re-integrators. This opened the door at Dell. Dell wanted LTO. By then Dell had committed to LTO and was abandoning DLT, because it was hard doing business with Quantum, because of their monopolistic, or their proprietary situation. We won Gen 2 at Dell. I don't think that's public, but they had a choice between Seagate or IBM, and they had had history with Seagate, where it wasn't as big a company to support. Dell is a tough customer. They're kind of like the Walmart of the OEM chain. There's a lot of bars you got to jump to sell product to Dell, and they beat you up on price, and they-- any rate, long story short, the channel dynamics at the end of the day weren't anything like what they thought they were going to be when we were having those earlier royalty division discussions.

And when you win a Dell, you suddenly have a cost absorption infrastructure that's huge, because you

are selling tens and hundreds of thousands of things, instead of three or four thousand things. And you can leverage that in a lot of ways in your business. Not only does it naturally give you a better cost position, to differentiate within the category, but we also were able to get derivative products off of the LTO wheelbase, differentiate them by taking it to a different cartridge, putting in some frosting on the cake, like encryption, or whatever the things we did in Jaguar, before we did them in RABF {Recursive Accumulating Backhitchless Flush}. I don't think we ever did in LTO, but we did that in Jaguar.

Master: There was a version of it, with LTO.

Teale: And basically sell that same LTO cost base into the enterprise channel for 25,000 bucks a pop, instead of 1500 a pop. So the point is, it was not so much us differentiating from each other. First of all, Bruce was never really in the role to do that. He was really very committed to the LTO, and we had other people attempting to make that differentiation. I mean, maybe you were involved, I didn't remember that.

Master: Somewhat.

Teale: Somewhat. It was hard to differentiate our product. I mean, price, and availability, and maybe some reliability claims, performance-- was it dictated, or did we allow ranges of performance?

Master: Ranges.

Teale: Okay. So there could potentially be-- but I think everybody went to the top, right? So we'd all be the same. Nobody wanted to be the slowest tape drive. So it was hard to do, but the channel kind of sorted it out. So Seagate kind of found its niche, and its market share, HP got its, and IBM got this surprising Christmas present. And so no further differentiation is necessary. We've got the best price, if we want to give somebody the best price. We don't walkways want to give it to them. In fact, did you know, Quantum ended up basically exiting production of LTO, and is sharing HP's supply chain. Quantum is still participating in helping them write specs, getting royalty, but they're actually depending on HP for the actual product.

Master: I wanted to add something, too, in this chronology of moving through the generations, it occurred with Gen 5. That is, the creation of using tape in a manner like using disk which the brainchild of Ed, who's going to be the next speaker, so this may lead into Ed's presentation. Disk technology, early on, started to emulate tape. They went all disk, and the applications changed, so there was just disk to disk talking. The brainchild, or the synapse that occurred for Ed was, well why don't we have tape emulate disk? So let's write the data to tape in such a way that it presents itself to the user as disk technology. So now, the one last thing left for tape, with all its other benefits, was ease of use. When you write data to tape, 20 years from now, you would need that application to read the data. You've got to recreate that application environment to access that data on tape. Instead, not write it in such a way so that 20 years from now, you don't need any of those proprietary applications. You can access the data as though it was a disk drive. It looks like a disk drive, presents itself like a disk drive, a self-describing technology.

Teale: Yes, and we will move there, that's good segue.

Master: The technology is called Linear Tape File System, which was created by IBM, and then adopted by the LTO consortium, and now, the development is continued by SNIA. SNIA is managing the specification for LTFS now, and the development still involves IBM.

Teale: And there's a huge piece of significance to Bruce's last comment that I wanted to amplify on. Many different technologies have attempted to make marketing hay out of something called the notion of media life. And optics is good for 75 years, and there's a war out there, and it's all based on voodoo math, and it's total BS, but it used to be a huge marketing war between various technologies. It turns out there is a truth to useful media life. There is a painful truth that a tape written in 1957 has still got all the information on it, but there's no machine to read it. That's the real thing that creates obsolescence of media. It has nothing to do with mean times to failures, and magnetic things not being magnetic anymore, it has nothing to do with any of that. And so if he is suggesting that there might be a way to improve the picture of recovering in the future, because you've simplified something to a lower common denominator standard, if you will, it's important. So there is a true value proposition there that did happen after I left, and I know it's the father of that.

Gardner: Just one follow-up, sort of a summary, perhaps. Tell me if I misunderstood, but so it sounds like IBM entered this with a channel that it had established. It really was, I think, John, you said there were three companies that really were in different market channels, and even though IBM did have a fairly extensive hard disk drive OEM channel, it didn't initially see that as a tape OEM channel.

Teale: Believe it or not, as we were starting to work on Gen 1, IBM was slowly but surely starting to exit the OEM channel. And in fact, the people that started our original OEM team, Dave Wall and Glen Allen, came from the hard disk drive OEM channel to establish the tape OEM channel. And I think what was happening there was that there was a long-term positioning, ultimately, to divest the disk business to Hitachi. But I want to reinforce something I said right at the beginning of this meeting. The LTO consortium getting together and talking about the channel conflict would have been a thoroughly inappropriate, illegal discussion. We all kind of knew that there wasn't a lot of channel conflict, which helped us be comfortable about tying the knot, but we never discussed it overtly, and still don't, because that's arguably anti-trust, and potentially-- it turns out that the channel that resulted just didn't happen to be the channel we originally visualized. We figured we would be at mainly IBM blue to market, and maybe some business partners, some small OEMs, other automation manufacturers. But we never imagined that we would be a high-volume, low-cost supplier to another enterprise. That was a lucky outcome.

Gardner: Continuing our LTO discussion with John Teale. He's now joined by Ed Childers.

Teale: Okay. Without further ado, Ed, like me, was one of the founding members of LTO. Although in retrospect it was a heck of a lot of fun, there's also a lot of pain that I probably forgot about, and I'm going to let you introduce yourself, and talk about it.

Ed Childers: All right, thanks. So bio first, I guess. I started at IBM, I don't know, '78 or '79. I've been in data storage most of my time. I started off with a big floppy disk thing that was supposed to be a competitor to HDD. Prospector was the code name at the time. Then went to a floppy disk program that was essentially the founding technology for Iomega, if you remember that in the day. We called it, Bluegill

– you [points to Joel] were the program developer of that at the time, if I remember. Yeah, that's right, Bluegill. Took a stint in printer development, then went into optical. And I moved into tape when IBM exited the optical business and dropped into the LTO development part of the business. Actually, I worked for a short time buying DDS for you, John, before we started LTO. So that was kind of the timeline scenario.

Teale: Well, I think you got a couple years of [as manager of the media group].

Childers: Yes, I had media responsibility prior to LTO starting, is a better way of putting it. Now, I'm responsible for development of tape drives and LTFS at IBM, and still the LTO technical representative to the TPC consortium, as John was describing previously. So that's the background, who I am. I made a couple notes while I was listening to John and other people talk, so--

Teale: College education?

Childers: Oh, I have a bachelor's in mechanical engineering from UT-Austin, and a master's in mechanical engineering in servo controls from the University of Arizona.

Teale: Married? Children?

Childers: Single. Children, four, yes, grown, happy. Different part of life. So my notes, one thing you mentioned about the LTO TPC [Technology Providing Companies], the whole kind of initial thing there, my memory was, is that, like you said, like John had mentioned, I should put it that way. We entered this when DLT was winning the world, because there was a department server in every closet with a tape drive on it. And that had grown to be a huge business that frankly, IBM missed, and HP and Seagate had also missed. So we were all looking at it enviously. Quantum, dumb luck, fell into it, I guess is the way to put it. So we were all looking at that enviously. HP and Seagate from the bottom, DDS, upward, and IBM from the enterprise, downward.

And this is one of these things, like butterfly wings change the events of history, where we lucked into it, also internally, at IBM, because at the same time as John mentioned, Vanderslice was our group leader for the HDD business, and there was a big focus on making that a component business, if you will, inside of IBM. And we were his posterchild of whitespace opportunity for what could you do in the HDD business to actually grow it? You can grow it by entering this other big business, called a billion and a half DLT component business that kind of simulated the HDD business, but on the tape side. And I don't know if you remember, John, because we actually got pulled out of the tape GPD division, and put into the component division, the LTO team, while the rest of the tape business was still in the other division for a while because we were viewed as a component business inside of IBM for quite a while. So we survived, kind of by luck, in that sense, that there was an interest in developing a component business, that segued into this tape business, if you will, and the channel stuff that John was talking about. But it was kind of interesting that the initial desire was maybe not to develop a LTO infrastructure tape business but really just to sell components in a way. And my memory was, as we got invited to the party, Seagate and HP were already trying to do something along this line, but they didn't have timing-based servo and that was the key that got us in the door. And our desire was, if we come, we want to bring the 3570 aka Accelis as they said yes. But, of course, as we said, that didn't end up going anywhere. I don't think you mentioned it

be we were named a Brewers at the time because we didn't want to have the world know what we were doing. So there was a lot of secrecy on the meetings and such.

But the most interesting thing to me looking backwards on that is that the idea that we could create an open specification that three or more people could compete in based on standards in tape was actually never done before. All tape drives had always been developed kind of hand in hand with the media manufacturer and you get to a point where you make it work and you say, "You keep doing that. We'll keep doing this," and we write specs based on working together long enough to make sure it worked, but not really understanding the requirements. 'Cause our channel on our head has different things in it than somebody else's does. So all tape drives at the time were always proprietary thing. Sony did their thing. Quantum did their thing. Maybe you could call DDS, HP and Sony at the time semi-open, but they didn't have open standard. They just all developed kind of together and said, "Okay, we'll all do it this way."

So there was a lot of skepticism, especially in the IBM tape business, that this could even be accomplished, that you could write a spec that people could develop too. And, like John said, we got into a lot of issues on what are we specifying, how do we spec it so that it will work without disclosing implementation details. That was a big issue, as we said. We had the secret sauce rule sorta thing. And trying to figure out at the end of the day, once we did it, how did we tell if it actually-- somebody meets the spec, when inherently to measure the thing you have to be somebody capable enough to actually build a drive, for example. You can't test a format without a drive capable of reading the format. If you were to pay somebody to go, "Here's a format. Go tell me if it's correct," they would essentially have to do develop the whole tape drive and channel and everything independently just to go do that. So the idea that IBM, HP, and Quantum were going to compete but we were also going to set up an open table certification process to verify that we all do it when we were the only ones who were capable of doing it at the time, was quite an interesting discussion.

And we evolved into a process we called the bootstrap where we would actually have a fair means of exchange and test each other's concepts. I mean formats. And it worked out. I wouldn't say surprising but listening here, I forgot so many of the pinpoints we went through trying to get this thing to work. What else did I have that I wanted-- oh. We had a big issue on what was spec. was it just what's in your hand to be able to interchange or were there going to be some piece of quality or liability in that spec. In other words, does it have to operate once? So you could read it and it could fall apart. Or does it have to be able to run 100 passes, 2 years, 3 years? Some of our partners were really wanting to put quality into the spec because they wanted to essentially be able to market that this thing was designed for a certain performance point, was my memory. At the end of the day, talking through it, you can't really have a spec that specs quality and we ended up having to spec that spec's performance, format in hand we called it. But no measure on reliability, how many times it'd have to do it. And in fact, all we're really doing in the consortium is certifying that licensees have demonstrated the capability to produce product that conforms to spec. we're not guaranteeing any level of reliability, life and things like that. So that was a real process that we had to go through, trying to get there. Lots of arguments, lots of meetings, lots of time spent in neutral locations trying to fight this out between IBM and HP and Quantum. It was...

Teale: It was very cloak and dagger.

Childers: Yeah, we had the neutral location rule because anytime you're negotiating something it was perceived, and I think it's true, that if you went to somebody's home turf, they had a psychological comfort

advantage, could go home, whatever. And we always had to go to a neutral location to have all these discussions. I don't think we ever had a meeting in one person or other's backyard in this whole time, which meant we ended up in places like Las Vegas and San Francisco and San Diego.

Teale: Ed mentioned the Brewers because we were brewing something up and we each had nicknames. And the important thing is, when we met in a neutral location, like Ed describes, there was nothing indicating the presence of IBM, Seagate or HP.

Teale: IBM was Guinness, HP was Heineken, and Seagate was Fosters.

Childers: That's what it was.

Teale: However, I do vividly remember. Once we went to, I think, the old Hilton at San Francisco Airport that's long since been torn down. We went into a room to have our meeting and there were these notes all over the room that said things like, "LTO who? LTO what? LTO"-- it turns out the DLT people had been there the day before having a meeting. Had somehow gotten wind of our secret meeting the next day and left us all these love notes in the conference room.

Childers: That's funny.

So that, plus like Bruce said, establishing a roadmap. We put together a roadmap, arbitrarily decided four generations. We'd always have at least four generations in the roadmap at the time. I remember-- what was it? Generation four we said would be AME Media, because it was a shiny technology at the time. But it was really just a multiply everything by two, each generation, and we finally decided, "Well, we need to show something cool for the fourth generation," and we picked AME, which still hasn't happened. But also cultural between the companies was interesting. And so, like I said, IBM entered this whole thing from the enterprise looking down enviously at an open market. HP and Quantum entered it looking up and, as such, they had different things, different cultures if you will. HP seemed to be much more organized on structure and consensus and getting an agreement prior to this. I personally think it was probably due to the fact that, in a big, big volume business, you have to be darn sure what you're going to go do and you can't be fragmented. IBM entered it from a, "We're all scientists arguing about what we think is best," more of an academic sort of thing in a way almost. And we'd get into these meetings where we'd be negotiating something and we'd get to appoint where we'd get off the script with HP and they would freeze like the Borg and say, "We have to go caucus." And they'd go off and talk about it and come back where IBM would talk out loud in meetings, which was sometimes-- pissed some of us off. I remember because we would agree on something without having ever talked about it internally, and we'd have to try to back up or clean it up afterwards. So it was a really interesting-- I wouldn't say a clash of culture. It was more like a marriage of cross cultures if you will and I believe the only way anything like this would only work was if you had three people. If you had more than three, you'd form voting blocks that would essentially get stalemated. If you had any even number, you would form deadlocks. And this only worked because there was three people that couldn't always agree and, at the end of the day, we were all mutually dependent upon each of us agreeing to go forward to market and make some money. So there was a mechanism to force consensus, i.e. you wanted to get out the door. But there was also a lucky break and the three of us prevented, we'll say, blocks forming that would kind of stalemate agreement too. Where did we go from there?

Oh, I have an opinion on the DLT thing, on why we were successful. I think DLT in one sense made a classic mistake in they were suddenly 70 gigabytes at the time we announced LTO to be 100. And the expectation would be their next product would've been 140 gigabytes 'cause they had previously done 35. And, when we announced this thing, they stuck in a new thing called super DLT at 110 or 20 and tried to make it backward compatible with the old one by adding [a second head]⁷ but the net of it is it fell flat in the market and helped LTO actually, I think, get acceleration. Because it didn't succeed very well and LTO hit, and three people doing it, and a way we win. I think DLT then tried to say they were open also by recruiting Tandberg or Benchmark or somebody like that to say they're making DLT drives also. But it was too little too late and I think they could have just held the course and said, "Look, we're at 70 now," and come up with their own 2x roadmap per four generations and say, "We'll be 140 between LTO 1 and 2," and they might've turned the tide a little bit better on trying to compete in an open format sorta environment. But businesses are the way they are because it's hard for somebody who's getting revenue for a thing to be the one to take their chance to risk it and change their business model. And so, Quantum took a leap and failed and LTO took off, in my viewpoint there.

And, in fact, having this open roadmap inside of IBM at least, turned out to be one of the best things we could've done because IBM is also a company where you have a lots of people with opinions saying what the features and speeds and feeds should be. And one of the primary way development teams fail is, at the last minute, you start adding things to it to try to be more competitive. You get nervous that it's not high enough capacity. "It's not high enough data rate and we need to tweak it higher to make it more competitive at the last minute," and that's when you choke. And the advantage of the LTO consortium was, and still is, that we set a roadmap. It's the agreement with the consortium. You don't go dork with that roadmap. It is what it is and you're resistant to people coming in and saying, "I think you need to increase the capacity by 20 percent or do something by 30 percent to make it more competitive," because it's a defined open spec. And that allows you to put your head down and actually do things that you plan on doing and not get distracted by a crisis that rises or a new person coming in and feeling that we're products that needs to be modified to be more interesting somewhere.

Teale: There was also an element of job security. We've made a commitment.

Childers: Yeah, we've made a commitment. We're a part of this consortium.

The roadmap was actually and is actually a very positive thing from the development planning point of view, of you know what you're going to do once you get in agreement with the consortium. And you know it's not going to change so you could put your head down, develop the product, not have somebody come in and try to push you off.

Teale: No requirements creep.

⁷ To implement backwards compatibility to the previous-generation DLTtape products, the SDLT 220 and SDLT 320 tape drives actually have two head mechanisms—the MR heads designed specifically for Super DLTtape technology and a retractable DLT 8000-style head that enables them to load and read DLTtape IV and DLT VS1 media. Source:

"Quantum Super DLTtape Drives," Gartner, December 2003.

<http://www.bus.umich.edu/KresgePublic/Journals/Gartner/research/95500/95593/95593.html>

Childers: Yeah, thank you. Requirements creep was the word I should use.

Teale: Then went to a floppy disk program that was essentially the founding technology for Iomega, if you remember that in the day. We called it, Bluegill – you [Joel] were the program developer of that at the time, if I remember. Yeah, that's right, Bluegill. Because we could always just turn around and say, well you'll have to talk to HP and Seagate about that. I'm not going to.

Childers: Yes, because have to get a consensus to change that particular thing. So that was very powerful in the tape industry, I think, is to get that consensus. Get it established, both externally from a marketing point of view as Brad and Bruce were talking about it but unexpectedly from a development point of view internally. You can plan your technology. You know what you're going to need four years from now because it's going to be a 4x thing and you can plan your investments in heads or whatever to prepare for that roadmap which you know you're going to get to. The only question is, is it a two year, two and a half or a three-year turn. But you know what's happening. So talked a little bit about DLT, cultural change. Inside of IBM LTO is also kind of a cultural change from the established tape business. It was seeded by you, John, from people like me, optical refugees, probably about what - ten, twelve?

Teale: There were lots of optical refugees because we mentioned in passing, when Vanderslice came in, he came to Tuscan. Here we were in the process of killing tape and keeping optics. Jim came down and gave tape a big shot in the arm and got us back in business as Ed described.

Childers: And killed optics.

Teale: And here's what he said about optics. He told a story about some couples that agreed to go out to dinner in Abilene in six months. Six months went by. Nobody really wanted to go to Abilene for dinner but they'd already agreed they were going to do it. I forget the punchline of that story but he used that for optics like- no return on investment-case closed." I remember you guys had Abilene signs in your labs. I don't remember why.

Childers: Well, since I still work at IBM, I'm not going to talk about too much confidential stuff but since you're retired you can. But the fact that we stopped optics and moved that resource to tape and seeded it with some tape people at the same time, picked up the Japan team which was also the ex-optics for code. The optics business had been a commodity business. We were making single card optical drive aimed at the commodity business. So it was kind of a different mindset of the people coming from the optical team, merging with the tape team which was a big enterprise business. Actually it was the right thing to do in LTO. It actually worked out better than I think we could've ever expected.

Teale: I would give it even more credit than that. The tape team was very enterprise-centric. We didn't know any other way to do it but to make it absolutely bulletproof and that wasn't required for LTO. Optics team not only brought the open systems perspective, lots of new blood, new ways to do things. I think it was a critical element.

Childers: Yes, it's interesting how-- I said the butterfly wings of fate. This was something that happened where the skills came together and it worked well. There were three major disruptions, if you will, from

what IBM had been doing and what we did in LTO. It was drastically different tape path and form factor from big pneumatics and air bearing controls stuff to simple rollers and small form factor. It was a single card of electronics and a controller, which we had had before. Multiple, multiple cards in a big card pack thing. And, as you mentioned also earlier, a drastically different head technology, an HDD leveraged flying head sort of thing versus a spherical head, specially lapped sorta thing. So all those were viewed as, "It can't be done," initially and, as things works out, it did-- it could be done and it evolved into LTO, which we had then leveraged into our enterprise business eventually. And those two things work well together for IBM. It's kind of LTO is the industry that is the, "If it doesn't kill you it makes you stronger," sort of thing. You learn stuff in LTO we would never learn in enterprise, just because of the magnitude of the market and the places that it gets. And drives coming back with spider webs inside of 'em and stuff like that you never see in enterprise. And enterprise gets the stuff that high-end customers want, like RABF which we did eventually roll into LTO.

Teale: Did we?

Childers: Not in the same form. It was just a sparse writing. We don't go sweep it up.

Teale: Ah ok..

Childers: Yeah, we just leave 'em sparse.

Teale: Yeah, but how many applications does LTO get behind that need it.

Childers: None.

Teale: So it was marketing.

Childers: Yeah, so it's marketing. Well, I wouldn't say none but not many. It was done because we could.

Teale: Amplifying the optics point again-- and this is a little out of school but I am retired. I think we mentioned earlier that LTO was not popular at IBM. It was an enterprise-centric place. We had business leadership, Charlie Andrews, Bob Maness, people who were very skeptical about our ability to be competitive, didn't want to hurt the enterprise business model. And I literally got no bid from certain Tucson tape teams. When I was attempting to put an organization together to actually do the development, I got no bid internally. "You're crazy. Won't work. We can't do this," by other team people. So, if it wasn't for the optical kind of good/bad thing that happened there, we would have never gone to IBM Japan and found those resources. We would've never done the things that turned out to be the exact right things to become successful. I remember I did not get a great appraisal one year and I told my boss, "I get the feeling I'm succeeding at something you don't want done." And he said, "You're right." But LTO is protected funding. It wasn't in-plan funding. IBM had this white space bucket held up at the corporate level by a guy named Dave Johnson, and as long as Dave was satisfied with our progress, he funded us outside of the process. So the local management, no matter whether they liked it or not, couldn't do a whole heck of a lot about it.

Childers: So it's just what I said. Sometimes we like to try to go back and put together the pieces to act like it was intelligent design and it was all this sort of really thought out strategy.

Teale: Act like we knew what we were doing.

Childers: And sometimes, in looking back, that's why I was saying the butterfly wings of fate is that it came together in the right way that actually worked out. That probably wouldn't've worked out if we had not had those factors come together at the right time. And you're right. It was interesting inside of IBM to watch when we finally got the first LTO drive together to be running and get tested. It was kind of eye opening that it actually was really reliable. It was quite a surprise to many that you could do something in this regime. But that's the way technology always works out, right. I mean, somebody does something somebody doesn't think can really get done and proves it works. And then everybody goes, "Oh, then we're going to go that way," and it happens inside of companies and it happens in the industry. And that was just an example of inside the company. We had a mindset and method of doing something that we were able, because of the investment, to do something a different way that turned out to be more successful and have more legs on it. That kept the business going for what? Fifteen more years from that point on, still going today.

Teale: Part of our panel is that we're going to talk about the outlook for tape because I have serious concerns. We got three more formats left and three players.

Childers: Yes, well, I think when you start getting there, it's not much different than disk or Flash. I mean, there's two and a half disk guys. There's two and a half tape guys, two and a half media guys. There's four Flash guys. So I think that's a broader thing that the-- the whole industry has built this business on top of taking commodities and putting intelligent software on it. That was our whole life. It's a runaway from the commodity market and now the real shift is no more commodities, and so we're adjusting to a no more commodities storage business going forward, both in disk and in tape and eventually in Flash. And that's the real interesting thing, I think, but we're off topic.

Gardner: No, that's actually an interesting topic and it is a little off topic but an interesting one nonetheless. Anyhow, turning back to the topic, you mentioned a couple of things like no air bearings, which is probably an internal decision. Maybe Seagate had air bearings and maybe HP did. And it's not proprietary now because they can go buy that drive and see it. And I don't want to know about LTO 7 but looking back at LTOs 1 and 2 or maybe even 0, could-- both of you guys, I think, step me through some of those driving compromises, which I take were driven towards low cost. Particularly differences on the cartridge.

Teale: Well, there's two conversations there that-- one of them is the decisions we made in our implementation. But the real driver was whoever made this decision, and this was a huge part of Ed's life. The world didn't really need another form factor with media.

Gardner: Actually, I think it did and I'd really like to hear you guys talk about that.

Teale: I'm going to encourage Ed to tell a little bit of the story 'cause he was on the TWG. He had the front row seat on this whole deal.

Childers: Yeah, like Johnson, my first job was to go convince HP and quantum to use 3590 cartridge, which immediately failed that. Because, as we said, they were both from the DDS world looking up. And it was an absolute imperative that there be a half-high form factor and that fundamentally, by math, laws of physics, unless you create anti-matter, you can't get that cartridge up and down and a half-high form factor, can't be done. So we had to come up with a thinner smaller cartridge to do that. The big fight became over the features around the thing. Does it look like a DLT? Does it look like something new? This dog bone thing in the end was a huge fight on how do you decide to thread it, because everybody's threader had to accomplish this thing. So the evolution of that was what's the minimum thing we can all agree on. It's called a pin. I mean, all the other proposals were something more around it, like that leader block or a half-moon shape, or something that would integrate into a person's tape drive easier, that particular person's implementation.

Teale: A Dogstar. Dogstar was an IBM cartridge proposal for LTO that utilized concepts from the 3570 cartridge.

Childers: Yeah, like a dog star.

Teale: No threader.

Childers: No threader. The pin became the least common denominator thing that everybody could agree on is really what it came down to.

Gardner: I thought that the judgment criteria was not least common denominator, but it was something on the order of technical excellence or best mode.

Teale: Best of breed.

Childers: That's a fuzzy line sometimes, right?

Gardner: That is true.

Childers: Sometimes, if there is only one thing that is an interception of what everybody thinks they can do, that becomes both the least common denominator and the best of breed by definition. These little notches on the side are interesting because we had a huge fight over do you take-- there's not a DLT. But do you take a 3590 notch and put it on LTO the distance from the ends so the pickers that were designed for 3590 cartridges could pick this thing? Or do you take a DLT notch and, if a picker's design for DLT it can take this thing.

Teale: I didn't remember that. That's interesting.

Childers: And, because we couldn't agree, we put both of 'em in there. So, if you held up a 3590 cartridge, you'll see this notch is exactly like that notch. If you hold up a DLT-- or actually, this is the DLT.

This is a 3590. This little funny angle in here is the DLT notch and that was to get agreement between the IBM, HP, and Quantum who were all designing pickers for different things and wanting to leverage their picker design. Now, we marketed it as a feature because it's more reliable. You say to drive holes onto one, the picker grabs the other. That's a revisionist history thing but the original was we just couldn't agree on anything. That's probably confidential but it's old enough it probably doesn't matter.

Teale: You were on the toughest TWG by far. 'Cause the cartridge, when you think about it, everything else needed to do 100 gigabytes. I mean, we had our internal execution problems changing heads, but the technology, most of it was known. There was not a lot of big technical risks other than <points to Childers>

Childers: A whole new mechanism

Teale: And that drove the whole new mechanism that drove the unbelievable pain that we had internally, had no idea what it-- and then, like I said, that's another story, another day and probably even another set of people. Ed could bring backs some engineers someday.

Childers: Well, the thing, unfortunately was that all of us in the room negotiating this cartridge were fundamentally people that had no experience on cartridges.

Childers: I didn't. I came from Optical. The guys HP that were negotiating it were Kitty Hawk guys, the ex- HDD guys. And I can't remember who from Seagate.

Teale: Those weren't tape guys, I think, but they had never been in the cartridge business.

Childers: All right, so we were all arguing about things none of us really, really knew about that much. And we ended up hiring in a retired cartridge designer from another company, the cheese head.

Teale: Leif Erikson.

Childers: To come in and be the voice of reason to say, "You really can't do that in a high production environment. This is what you have to do." And I would say that guy probably influenced the details of the design much more than any one of us who were just arguing our positions <inaudible>.

Teale: That was a good move to get that guy and he had retired from a 3M/Imation. He had probably 35 years of designing cartridges for high production, quarter inch cartridges, coyote (3570) cartridges.

Childers: He would come in and say, "High-volume molding, the tolerance you can live with on something that wall thickness is x," and we'd all go, "Okay."

Teale: He just knew and he could just tell us."

Childers: The only guy I knew that in an expensive shabu-shabu dinner shishkebabled his sushi and boiled it.

Teale: Yeah, we were- he was from Wisconsin and an avid Green Bay fan, and raw fish was not his thing. And we went to Japan and we were having one of these dinners where they had the hotpot. You stick the meet in like sukiyaki or whatever they called it. Shabu-shabu.

Childers: Shabu-shabu. It was sushi.

Teale: And they served the raw fish as an appetizer and he'd take his raw fish and take it in the hotpot and cook it, then he'd eat it. Good times. Anyway, Leif is still around somewhere in-- I think we gave him plenty of love.

Childers: Minnesota, as far as I know.

Gardner: I noticed, maybe it's my lack of knowledge but LTO tape comes out of the cartridge at a right angle as opposed to, say, the way the tape comes out of the prior cartridges.

Childers: Yeah, in previous tape paths, the tape would come out so that you could pull it forward and you'd pull it way up to a big air bearing and you have these very long tape paths. And I remember trying to draw that on the board, would always result in a picture that didn't look very good because we had this big tape path and the tube reels way back here. And the reason for the length was, in that concept, the idea was the reels were noisy and so you had to get the tape and the head a long way away from this noisy stuff and control it with guides and everything to get the stresses off of those guides so you didn't damage the edge of tape. And so it was this concept of the longer the tape path, the better for being able to manage the mechanical motion.

Teale: We did go into that in great detail yesterday so they know what you're talking about.

Childers: Okay, and so coming back to LTO, you obviously can't do that in a little form factor like this. So it had to be a short tape path, and the only way in a short path is to get it out and around the head is to bring it out the side and bring it up. And it also necessitated a new way of tracking because, in a short tape path, the tape's beam strength is very-- is high. So that's where we came up with the whole idea of surface-based guiding, don't control the tape, follow the tape, all that kind of stuff. It was an implementation thing.

Teale: Well, let me amplify on the difficulty of Ed's job. He's gotta agree on all the details of that cartridge but he's not allowed to have a tape a tape path discussion. He's not allowed to have a "How are you going to thread it?," discussion. So these guys are literally attempting to agree on something that is crucial without really even being able to share their thoughts about what they're going to do with it.

Childers: And so one of the technical issues along those lines was do you just start to take the magnetic tape at the leader pin and live with it or do you put something stiffer, a leader tape, in this thing to make it easier to thread without dinking the tape up. We had a huge fight about that.

Teale: Yeah, there used to be something called leader tapes on some technology implementations that were stiffer than the media. 'Cause you want to shove it through something.

Childers: Well, DLT had one.

Teale: DLT had one -- threading is not really treading. You're shoving something stiff through there and then, eventually, you can cinch it all up. But then, how do you have that conversation without talking about how you plan to do the threading?

Childers: Implementation, yeah. So we got in a huge fight about that. We, IBM, were interested in a leader tape because we thought it would make it easier to thread. Others were not because they felt it would add cost and that leader tape takes away from place to put capacity. And, if you're worried about technology capacity optimization, you don't want a leader tape. So we fought and fought, never got in agreement, ended up without a leader tape. And one of the other partners came back after partway through LTO 1 and said, "I think we should have a leader tape now."

Teale: Really?

Childers: Yeah, and so we added it back in.

Teale: Did you?

Childers: After we started; in Gen 1.

Teale: Really? Okay.

Childers: Yeah, after we started. Because one of the implementers decided it was a good idea after all, and we reached a consensus after the fact.

Teale: Boy, we really let that guy off the hook, didn't we?

Childers: We did.

Teale: We could have just jammed it and said tough.

Childers: You don't recall this, obviously, but we were thinking about not agreeing to extract something.

Teale: That's right. We wanted some...

Childers: But we got nice and-- 'cause we wanted it too.

Gardner: So you sort of agreed it had to be short to fit the form factor but how you actually were going to thread it was left to the implementers and you couldn't disclose it.

Childers: Each of us. Yeah, we couldn't disclose it. That was a competitive thing. Obviously, three companies who are in a business all teaming up on a, "Let's neutrally design something and implement something, go to market," there were concerns on a competitive nature. So, as John said, the line was drawn where you can talk about specification for interchange but not implementation.

Teale: But, supposing Ed is sitting there. Let's say this pin has been proposed by HP and Seagate and Ed looks at it and, in his mind, he has no clue of how he would thread this device. Let's just pretend like you didn't. That is when Ed could say, "All right, show me. Show me your secret sauce on how you're going to do it, because I don't see a way to do it, and, if I can't see a way to do it I can't agree to it. And that would've been an example. I don't think we ever got to that point on the leader pin, but it was close.

Childers: Not on the leader pin. Yeah, like you said, SNR once and something else. We only had like two secret sauces things, maybe, total.

Gardner: So that's interesting. You guys, pretty much, on the technical side were able to agree on really an incredible set of specifications when you think about it.

Teale: In a pretty short amount of time.

Gardner: With only really two big controversies, SNR and...

Teale: Well, SNR was Gen 2.

Gardner: Oh, SNR was Gen 2?

Childers: Yeah, the cartridge stuff was all Gen 1.

Teale: Gen 1, the biggest controversy on Gen 1, other than the details of the cartridge-- because that was where real heavy lifting was. Everything else we were borrowing from other pieces of-- the only other controversy was the one I already described on the argument about which compression to use.

Childers: But again, the thing that made it work that engineers are always engineers for the most part and they want to figure out how to do stuff. And so that helped. And the other thing is that it required a lot of personal negotiation relationship skills. And so, I think we evolved into roles where that became the thing that was most required to be successful and everybody had a different style. I remember one of the vendors brought in somebody specifically to counteract you. Introduced himself that way.

Teale: I don't remember who that was.

Childers: But, so there's a lot of how do we deal with each other and get what we want and not create warfare between companies. So there was a lot of relationship management that had to go into it.

Teale: I had a pretty aggressive personality, as you've already probably figured out, and when I thought I knew what I wanted, that's what I wanted. But the HP guys-- as aggressive as I was, he would become more stubborn almost. And so, I agree with Ed. I'm not sure the management team helped but you guys were able to do it at your level. But seems like all we did was fight.

Childers: I always felt like the thing that was unsung and unrecognized in the whole LTO thing was not the technical obviously but the ability to work with another company productively and get it done. We didn't have escalations but a handful of times. And I don't think we ever, inside of IBM, anybody in the

corporation, really understood or even today understands the difficulty maybe on getting agreement between three independent companies with different incentives. And, like I say, it's just because we all had a mutual need to be able to go to market that you get at the end of the day agreeing on something to go forward. And then you horse trade. Trade off what I need versus what they need and it becomes a barter system almost.

Teale: I'd like to amplify on that something Ed mentioned regarding verification. So it was marketing as much as it was technical that we have a means of certifying that whoever wanted to ship a product with LTO on it was in conformance. And the onus was on us to declare whether they were in conformance or not, because we were protecting the LTO brand. Plus, Bruce could take that independent proof of compliance and use that as a bit of a marketing apple. How do you do it? Ed mentioned how he did it but he kind of glossed over how hard it really was. The actual media's pretty easy. There were many generations of ANSI standards that taught us what you need to know about the media. It's mostly a functional spec but there are some other things. But there are things that are not in a spec that would surprise you. You don't put magnetic particle in there. You put a functional thing in there like the SNR must be and the media manufacturer says, "How do I meet it?" I had referred to what if a media manufacturer doesn't use the media formulation that we think it's going to use? And what comes out was something that is in compliance but it wears our ads out in two days. But believe it or not that's a conundrum and we couldn't solve it. They do what they do and we'll just have to let their reputation in the marketplace be the policeman of this.

Childers: Exactly.

Teale: So we kind of knew when a media vendor shipped us one of these and said I'm ready to go. They would have to have-- they would have to pay a fee for the compliance testing. We had an independent third party contractor that would do all the physicals, all the measurements right down to the fiducials down to the—fractions of an inch to make sure it was compliant. But the drive was hard. So what a specification, an interchange specification, specifies is how to write the tape. It does not specify how to read it because interchange is based upon what state did the last person who touched that leave it in, and is that a valid state? So you would be surprised at how tricky it is to write an interchange specification. It's kind of more than meets the eye and less than we see at kind of the same time. So now the question was all right, there was no collaboration on schedules, no collaboration on implementation. That was arguably illegal, so none of us knew our schedules. Supposing HP declared that they're ready to go on Gen 1 and they send in their check for certification and they send in the test cartridges we had specified, you've got to supply ten cartridges that you've written and then we're going to verify that they were written up to the specification and then will either agree that you did or you didn't; well, so, what about the scenario where HP's ready? I haven't even figured out how to read the damn tape yet. That's how far behind schedule I am. Let's pretend Seagate is equally far behind schedule. What next? Well, here's what happens. Seagate and IBM could not legally filibuster HP just because our schedule sucked and theirs didn't. So we actually had a timer that said the other two companies have two months or whatever it was-- two weeks, two years-- I don't remember how long it was.

Childers: To find fault with cause.

Teale: To find fault with cause. And if no fault with cause is reported, you're free to go -- if it's discovered after the fact you made some mistakes, guess what the new revision of the standard is, and we've got to conform 'cause we were late. So in other words, the anti-trust, the pro-competitive element this was protected by all— six ways from Sunday it seemed like, I'm giving you an example of, you know, how they crafted that language. Now the reality is that nobody ever got so far ahead that the other guy couldn't read a tape. Everybody I thought behaved very honestly in this system. There was no gaming because when you say fault with cause you've got to prove okay you wrote that, that interblock gap is too large, or whatever it was we're picking on, there was a whole 'nother element of this involved timing-based servo formatting, but that was the media manufacturers job and I know how they verify format in-factory. I have no idea how MAC⁸ does it or if we ever even let them do it.

Childers: They don't do anything on the format in cartridge, just LPOS (Longitudinal POSition) verification.

Teale: But just the planning of the verification process was enormous amount of work and I know I was involved in almost all of it because..

Childers: It was all Gen 1 stuff.

Teale: All Gen 1, and also because HP and Seagate generally had never participated in standards bodies. Like you said, they kind of agreed on what they were going to do. If there was a spec it was after the fact, and..

Gardner: So you didn't specify a standard reader that the written cartridge had to work within varying parameters.

Teale: Well for things like SNR, I think we established okay, we're going to have to have a head, it's going to have to have certain characteristics, gap lengths, whatever you're measuring, but in a lot of cases all of the requirements were relative to something else. I forgot that I think Ed was more involved in it than I was is that we had reference tapes. Go ahead.

Childers: Yeah, that was just the-- so do you have a media get sent into MAC that certification entity, and it's supposed to meet a spec, but all specs are referenced like the platinum bar in Paris, the meter thing, so you have a reference tape that everybody gets and you measure your performance and your method to that reference tape and you correlate it to what the certifying entity does. So it was a series of standards, if you will..

Teale: So for example you'd say okay write a density pattern, measure the SNR; SNR shall be within plus or minus something of the same apparatus doing the same thing on a reference tape that's been

⁸ Measurement Analysis Corporation, a small metrology company in Torrance, CA. They were contracted by the LTO virtual company to verify that media licensee's products conformed to LTO specifications as part of the "compliance" aspect of LTO. [J Teale]

provided to you. So that way if there's some variation in who's supplying the test heads, we take a lot of that variation out by the relativity of the reference tape.

Childers: The channel is different.

Teale: The channels are different.

Childers: Except as you pointed out..

Teale: And there were arguments about who was going to supply heads to the verification entity and I don't know if we ever supplied any. I think we kind of dragged it down at the time.

Gardner: So John, are you happy with where we are at this point on our discussion with Ed?

Teale: Well, I did want you to talk about the cartridge; we did. To me that was one of the most important things about LTO. Why don't you-- I didn't know there was a Gen 7. I didn't know there was a Gen 10.

Childers: There is a roadmap just like we did the other roadmap. Multiply out. I'll make one comment on this whole thing from my perspective. The thing that is most disruptive, changing in the tape industry was when DLT was there it was all this plethora of proprietary formats. There was no competition. Tape was really expensive and just was going gangbusters back then. That was the 40 percent a year disk days and tape was 10 percent a year and it wasn't moving. It would have died without an LTO. What LTO really did is force all the media vendors to compete on the same thing. Prices came down, dollars per terabyte went up and became competitive on a rate with disk primarily because we created a standard that people had to go compete on, both drive and tape. I mean, there's as many people [e.g., companies] died in the tape industry on the drive side as the media side over the last 10 years.

Teale: More attrition.

Childers: Maybe more attrition maybe on the drive⁹. It was like a hard medicine thing. If tape would have been dead, dead, dead in 2002; if we kept on the same path we had all these fragmented propriety formats and HDD taking off like it did, but LTO forced this competitive nature, trial by fire. You survived it or you didn't, but it drove the price per gigabytes of tape at a rate 40 percent a year that became more competitive to disk 'cause disk slowed down after that period. And it was everybody loves to hate LTO if you're in the business because competitive is too-- can't make money. It's yada, yada, yada, yada, but I always say if you didn't have an LTO there wouldn't be a tape business at all. We wouldn't have a Jaguar; we wouldn't have any enterprise tape. It would be-- DDS would have been the last thing out the door, essentially.

Childers: So that's really, I think, the legacy of LTO is it..

⁹ From approximately 25 tape drive manufacturers circa 2003 according to Freeman Reports, the industry in 2016 has only 6 vendors actually offering tape drives but most of them do not design the drives – there maybe be as few as 2 remaining manufacturers in the sense of design ownership. [Gardner research, see Tape Drive Manufacturers 2001-2016.pdf]

Teale: I'd like to amplify on that.

Childers: Was through the squeeze point, if you will..

Teale: Amplification is in a bigger discussion of the evolution of tape technology that we've had on other days. And I didn't say that about LTO 'cause I knew you were going to say it, but I did say that the standards that became legacy, formats that became super successful, legacy things that belong in the history museum were all pretty much at multiple industry participants, practicing 3480, LTO. 3590, nobody ever followed us. It would have been nice if they would have. I think 3590 would have been a more successful product. 3480 certainly was more successful when Fujitsu made one because it consolidated large pieces of the market on a platform that-- we had a dog in the race. And so now you see the exact same thing with LTO, the sparklers, I was going to encourage Ed to do and maybe we can pull the panel members up, was this has taken over the universe with the exception of two enterprise proprietary formats, one from IBM and one from Oracle. To the best of my knowledge, this is it.

Childers: That is it.

Teale: This is it.

Childers: There's nothing else I can see.

Teale: And what worries me is it's a little scary because I do firmly believe that there is a long-term forever requirement for tape, or a solution that has tape's attributes for the ultimate emergency, the ultimate disaster, the end of the world. And it really bothers me that we've only got two companies producing this thing. I mean, the drive; and I think maybe two or three producing this. Fuji, Sony and maybe Maxell are still in the game I lived at IBM for a long number of years and I imagine Ed still has to justify his existence on a monthly basis because it was kind of that bad when I left the business. So I wouldn't rely on IBM to continue, you know. And HP's had similar problems inside of their business I know. And it kind of scares me that tape could disappear altogether and there's nothing to replace it with. Because that would drive everything very suboptimal, in my opinion, if you were subbing disks for tape everywhere.

Childers: Well, the big transition that's happening now in the tape industry is moving from a backup to an archive. We talk about backup is pretty much getting disk to disk to cloud or whatever, but the underlying scaling thing that's happened that is keeping tape relevant is that at the same time that data growth is expanding, which is fundamentally because bandwidth got better, now you can get data from Point A to Point B; data has value so everybody keeps everything.

Teale: We used to have to put cartridges on an airplane.

Childers: So bandwidth has gotten better, data growth has rocketed up. Everybody talks about the data explosion. And at the same time HDD is literally, literally slowed quite down. And so the problem is in the data center, this balancing equation of your data's growing at the same rate the HDD can support has stopped. That no longer happens. Data is growing at 300 percent, disk growing at 10 percent annually. You have an equation that doesn't close. A lot of the whole cloud thing are people saying I don't know

what to do with this; I'm going to pay somebody else to solve that problem for me. That's why people pay the cloud because I-- but at the end of there aren't Keebler elves in the cloud, it has to be something, and that's where tape is finding its niche is it's the balancing of the equation in the data center if data's growing faster than HDD can support it; tape fills the niche for part of that.

Gardner: Our final session today on LTO, John Teale will actually act as a moderator and talk to our three other panelists about LTO, John.

Teale: Okay. Thank you all for your time today. The construct for this session that I had in mind was to start with the who, what, why. A little bit of insight into what it was like having decided who, what, why. And I'd like at this point, the reason the Computer History Museum is here is because it has been recognized that at the Computer History Museum, there's some vacuum column tape drives. You won't see a 3480, and you won't see an LTO. And the storage special interest group that Tom and I are part of recognize that that was two major holes in the storage piece of the computer history museum. 3480, pretty obvious in retrospect that that was legacy, that was big. So, what I'd like to know is something about the success of LTO. In other words, removing from who, what, why to impact on the market. Maybe some sparklers, if you have any sparklers. Maybe your own theories on why you think it was successful other than part of what we've discussed. Maybe start with how many LTO drives have been shipped in the universe, I don't know.

Master: Sure, to date, so from the year 2000 to now, there's been over 5 million drives shipped and over 280 million cartridges.

Teale: Wow.

Master: I don't have a sparkler on how much data that is. There's probably some calculation like going to here to the moon and around if you strung all those tapes together, probably 100,000 times, but I don't know. I don't know if there's a sparkler like that or not, but it's a lot. It's a lot of cartridges worldwide, all industries, extremely successful.

Teale: How would you think that compared to 3480 -- do you have any feeling? My guess is, it's in order of magnitude or two larger than 3480 was, well the drives, maybe not the media, I don't know.

Master: Media too -- just compare the capacities of what each could hold. I mean, it far surpasses 3480, 3490, anything that's been out there.

Teale: Any other comments on the impact?

Johns: Well, from an IBM perspective, we benefited in a couple ways. We had this whole business that you're well acquainted with, where we had all of those industry formats when I first showed up in 1978 we had all those products because our individual server product lines, whether it was the PC division, or the AS400 division, or the RISC System. Some of them had 8 millimeters, some of them had 4 millimeters, some of had QIC, some had them under the covers. As a result of that, they wanted to have some level of automation associated with this. So we had this whole business within our own organization that really provided those kinds of solutions to those customers. So, we had a small automation with all these

different formats. So while it didn't takeoff from-- in the OEM channel right away, or the IBM offering perspective, we did a pretty effective cannibalizing our own business, and replacing these other formats that we had been selling with LTO format, right? So we convinced our AS400 brethren to adopt the LTO format. We convinced the RISC, we've got the PC company to start to remarket LTO. While we weren't in Gen 1 necessarily taking over the world. We had this internal source where we could take this format and we could put it into the IBM channel with IBM salesforces, and make it the preferred format across the IBM server product line. And we had a lot of success with that. The secondary, and I think where we timed it right, is we had a very strong fiber channel offering with Gen 2, and that was at the same time that storage area networks were starting to take off within our customer businesses. So this idea of a shared library really drove not only the disk world, but the tape library world. Where instead of selling two or three drives, we would sell a large library with 10, 12, 16 LTO tape drives in it. All attached to the fiber channel network and shareable across multiple servers, which from a customer perspective was great, because they didn't have to dedicate specific servers to specific drives. They can now share the asset. So, and we were-- we already, sort of, knew how to do that from the enterprise world. This was-- this was not a new set of skills. We had people on the disk technology, and the tape technology that new about storage area networks, and knew how to set them up. So we found ourselves in Gen 2, in a really good position strategically. And that's, I think, when you looked at the-- the bow wave effect, that's when it really kicked in for IBM with the storage area networks, with our fiber channel offerings with Gen 2, by then it'd been established in the marketplace that we really were able to bring out a second generation of our product. We were-- weren't just a one hit wonder. So all those things, kind of, came together and the 2002, 2003 time frame. Your timing was perfect Bruce, you got too.

Master: Yes, Gen 2.

Johns: Yeah.

Childers: So, I'll insert a little bit more globally perspective from the industry thing. No storage-- no technology is actually is just competing against itself. Tape competes against other storage technologies. And at the time, when this started, tape was a going out of business from an industry perspective because it was proprietary fragmented, the consumer was falling out of love with tape at that time. The consumer base was going away. We were stopping going to blockbuster and getting VHS tapes, and getting DVR's, or DVDs. And so, not only were we losing to consumer base in the industry, which is a huge cost issue. The rate of technology progression was slow because of the fragmentation of the various formats. And that was at the same time that HDD was how big of HDD do you want in your PC laptop thing? And they were really going into the consumer in a huge way. So they were going faster than they did in, in years and years. So, we were in the environment where consumer based collapsing, HDD going crazy, and the desire in the industry was to invest in growth industries. You'd be in an idiot to put a dollar into tape, when all of this HDD was growing at 40 percent a year, and there was unlimited demand, and the consumer was going away from you. So, that was the environment of the tape industry at the time we started LTO, frankly.

Teale: So bridging from that, and I have a computer history add-on right now. Not an IBM hat, so I don't have to be nice. It seems to me LTO failed in one sense. I think, in the beginning we thought that when we hung the shingle out if you will, in November of 2007, and said, "Licenses are for sale." If this industry

was so fragmented and if people were so desperate, how come all of those other people didn't buy a license and get in the game? Why did they just allow themselves to go away? And that's a drive maker question, a media maker question, it seems like LTO could have done so much more to stimulate automation and maybe even develop some standards that would enable people to switch out stuff in and out of automation. So in other words, you didn't sell very many licenses is my point, and why do you think that is?

Childers: Well, well you can answer it a couple ways. We've sold enough, is the point.

Teale: Three.

Childers: Five media vendors [Fuji Film, Imation, Maxell, Sony, TDK] then, perhaps six media vendors, if you count Tandberg. too at the time. [As of 2016 on Fuji Film and Sony remain as manufacturers].

Teale: Well, media vendors came on board, but the drive makers didn't. Exabyte didn't jump in.

Childers: Right, 'cause they all had a vested interest. I mean, LTO, whether we think it's successful from a volume licensee point of view, I don't know. But it's clearly successful from forcing the competition in the industry that allowed it to compete with hard disk. That was the point I was aiming at earlier. How many do you need to do that? Enough, three, four, something like that.

Teale: I would have thought the more the merrier.

Childers: Well, I think that you go to realize, like somebody like Quantum, there's a business unit that's making money hand over foot with the proprietary format. And there'd have to be somebody inside that company that's going to say, "I'm going to take a risk to cannibalize this moneymaker, to go after a new thing called, LTO."

Teale: Insurance policy. Business, internal business insurance policy. How could you afford not to?

Childers: Quantum did it by buying Seagate eventually, right?

Teale: Fair enough.

Childers: That was there insurance policy, Quantum rode this until it's not successful, then I'll buy the competitor. Sony, rode theirs until they weren't competitive. I think, it's inherent in any big company where if it's making a lot of money, the risk of a business decision maker to say, "I'm going to kill that and go after something else." Is an odd thing to happen. It's not typical.

Teale: I think, that's a great answer for big businesses that are diverse, and also doing other things. But a lot of these people that we're talking about were in fact one trick ponies. I would think from an Exabyte point of view, from a Tandberg Data point of view, very capable Norwegian company, you know they could have come in and competed here. They had MR heads, they had track following, and quarter inch. I can't for the life of me understand why they didn't sign up.

Childers: Tandberg did buy and LTO license and produce..

Teale: Did they produce?

Childers: They did. They were unsuccessful. Why they were unsuccessful, I don't know. That'd be a talk to Erik and ask him what the hell happened in the business, sort of, thing.

Teale: Do you know what generation, I don't remember.

Childers: Either two or three.

Teale: Oh, okay.

Childers: --Fujitsu also bought a license in the beginning.

Teale: I didn't know that either. Okay.

Childers: At the very beginning, they never produced a product to my memory, but they did also remember Fujitsu had the Diana large format tape drive with rollers that somebody in our company said, "We'll never buy a tape drive that has a roller in it." Somebody sitting in the room, I think, said that.

Childers: They bought a license. I think, Exabyte was wholly dependent upon Sony and had a whole business model that was buying the consumer based thing, and turning it into a data product.

Teale: That was going on --- Quantum practically killed them before LTO came along. [The IBM 3570 (Coyote) had also made solid inroads into the Exabyte install base. So between 3570, DLT and LTO, Exabyte finally vanished.]

Childers: Yeah, and I think they just couldn't make that transition because they didn't have a technology base-- I don't know facts with Exabyte. I'm just reading between the lines, all the technology was developed at Sony, as far as I know, and they just bought it, and put it in a different package. And so, when they had to shift to a new thing. They could have maybe gone to a IBM, HP, or Quantum or something to go buy that thing from them, and then try to turn it into an LTO drive. But why would we sell them the bricks to go do what we wanted to do our self? The only company I know that maybe could have done it was Benchmark at the time.

Johns: Oh yeah.

Childers: So, I think the companies like an Exabyte, or some of the Travan guys.

Teale: There are a bunch of them.

Childers: Anyway, they were taking consumer based tape, that were in camcorders and stuff produced by a Sony, a Panasonic, or somebody like that, taking the engine, and putting something around it to make a product. And when the technology shifted they had no juice to go do the new technology. So the companies you mentioned that had it, Tandberg, Fujitsu, they both tried. Why they failed, I don't know.

Teale: Okay. Interesting.

Master: To follow onto that, from a marketing perspective, I think it's really a matter of perception. The marketplace at the time, the key competitor, way back in the beginning, being DLT versus LTO, which is an open standard. To them, they might say, "Well I've got one provider here in this-- current market share holder, and now I'm looking at this open standard where there's three or four drive providers and six media providers, that's a lot." So the original question you had was, "Why wasn't there more." To the marketplace, that seemed like a lot already.

Teale: And I didn't mean more out of the shoot, I meant over time, you would think that as LTO was becoming successful, that it would have attracted more. And that was really the basis of the question.

Childers: I think, the other answer is at the same time, is the fact that, if you had a dime to invest in technology at the time, you probably would have put it in an HDD solution or something because that was the thing that was really growing at the time.

Teale: Yes, but now I'm being a slimy reporter here on-- on purpose. I'm presupposing that now we're in Gen 4 LTO's clearly got legs. You don't have that argument anymore -- about where you're going to put that dime. Because it's looking like a sure thing. And it seems like you should have been able to start attracting more competitor instead of less. And I'm going to be even a little more critical in the media side was that an industry volume demand just didn't support six vendors or what happened there? Did your capability run away from half of them, or what happened on the media side?

Childers: That's a good point.

Johns: I think, one dimension of this is this is really hard stuff. It's not easy to transition from one of those other technology platforms that we talked about when the other tape platforms to LTO. And each of the companies we mentioned would have a slightly different challenge because they'd have different assets and liabilities, and different skill sets, but it's really hard. They'd have to fill some gaps regardless of where they were. And it would be expensive. So, it's one of those discussions as Ed, kind of eluded, well we can milk this business we're in for some period of time, or we can come up with probably, a very large investment to get into this other model, where I've already get entrenched-- entrenched competitors. What am I going to do better than they're going to do, given that we're all meeting the same specifications?

Master: Some of those might be library vendors that could have gotten to the drive business, but then it's that make versus buy decision.

Johns: Okay.

Master: Well, we could buy it, and OEM the drives then, and put them in our libraries where we already have an investment, instead of gearing up like you're saying to get into the drive business. Or even media business.

Teale: I'm going to go back to Ed here, I'm sorry, but real quickly, I just wanted to clarify my comment on automation. It seems like LTO could have been more aggressive, in terms of, standardizing the universal

picker of the universe, rather than accommodating two things that happen to exist. And that would arguably increase participation by auto makers of those types of products and then you-- starting with just standardizing around the cartridge, you could have expanded into other ways of getting automation vendors to conform to LTO. So we got two things-- sorry, Ed.

Childers: On the first, media vendors surviving.

Teale: Yeah.

Childers: I don't know facts again on this. My personal feeling is that, that was a case where some media vendors had business models that assumed a high margin product. They had infrastructure assuming a high margin return on what they were doing. They didn't have a large consumer base on anything.

Teale: Nothing hardly at all.

Childers: And those guys died first. And the people who had a consumer, or that were commodity based businesses, tended to survive. So I think there was a business model. A shift from the high margin dinosaur, sort of, businesses to the commodity, more competitive businesses.

Teale: And that, sort of, implies that there was a world capacity in excess of demand.

Childers: There was.

Teale: You're right, once upon a time, there was an infrastructure to pop out gazillions of VCR tapes. And it's gone.

Childers: So another way of putting that is that as tape left the consumer industry, there was an oversupply of capacity. And the people who survived, were the ones that had a business infrastructure that could compete in the cost structure at the time. The ones that didn't, died. And it died on two ends of that. All the really cheap, cheap like Turkey and Korea, and all these places that were making VHS tapes by the gazillions, those died, of course. The guys that focused their business on high margin enterprise, sort of, tape products only, they died. And only the guys that had a consumer, an enterprise, kind of, both business model internally to their media business are the ones that tended to survive. So I think it's just economics in that sense. Now, to your other question, the other thread.

Teale: Yes.

Childers: I think it is true that LTO would have done better, or it arguably would make sense to have standardized interfaces that mattered to automation. So that the SCSI connector is in the same place, so you didn't have to dork with cabling that the throat height and the way you, insert is in the same way. So that theoretically, a vendor could say, "I can buy any one of these tape drives and plug it into my thing just like that.

Just like a disk drive vendor can plug in any HDD into their rack and make it work, 'cause they're all exactly the same. We, of course, couldn't agree on those details in the consortium, and that's where we

might have arguably fell on short, from a crew open, open, open. But, there was some talk about that, but we never could get a consensus. So I think that is a place the consortium could have done better in retrospect.

Teale: But maybe it doesn't matter now that you're headed to this cloud archive thing. It's all going to be big boxes, right?

Childers: Well it is an interesting thing that tape has gone from the aggregation at the mainframe, big tape, reel to reel, 3480, 3490, to a distribution and department servers all over the world that's DLT, sort of, stuff. LTO entered that, and now it's going back again. But now, the mainframe and that aggregation is called Cloud, or hybrid Cloud. So, it's actually going back to a place where tape is more competitive. 'Cause tape is more competitive at scale of course. So, big libraries make more sense than small libraries, 'cause your amortizing cost of mechanism across more cartridges. And that's actually a positive sign for the tape industry because that aggregation is making tape more competitive than less.

Teale: All right, I am Mr. Big Cloud and you're asking me to adopt, clearly a very good technology that I can only get from two people who have demonstrated intermittent lapses and commitment. Is there some reason I shouldn't be scrambling for an alternative to tape in that application.

Childers: Yeah, I'm sure that people do. Dollars count. So when you're Big Cloud, these little difference in cost add up to be big differences in cost. So they're more focused on dollars count. And that concern of how many vendors doing it is also off set with the fact that now, disk drive vendors are only a couple too. I mean, arguably, there's the same number of tape as disk. So from a vendor stability, point of view, there's not much difference. And, disk drives are starting to airline price-- airline seat price their disk. In other words, you buy the same thing at different prices depending upon what you're using it. And so, I think they're in a rock and hard space. In other words, this is the commitment for 20 years of love, it's a, "This is the best thing for the next five years. And we'll see if Flash kills disk, and maybe I'll move everything to Flash."

Teale: Yeah, and that's very consistent with the overall financial trend toward a shorter amortization schedules in general. Did you have any questions Joel?

Joel: No.

Teale: Okay.

Childers: Elaborate a little further on that point. The thing that serves tape really no value in it a cloud environment is longevity of media.

Teale: That's right.

Childers: Because it doesn't matter. All that matters is that you don't have to replace it as often as disk. And you make the upgrade of technology when it's the most cost effective timeframe, which is like every eight years, instead of every three years, if it's HDD, so.

Teale: Okay. A little speculation. LTO has taken over the universe, everybody else is gone except for two things. How are those two things surviving and why are they surviving, and when are they leaving, and are they leaving?

Childers: Two things? Or two companies?

Teale: Well, there's two formats that aren't LTO. That are still shipping in the universe. One is a proprietary from IBM, and one is a proprietary from Oracle. Excuse my hand, I'm wondering how those are continuing to exist, and are they serving a different market from LTO? Or is it all headed to Mr. Big Cloud.

Childers: I think from my perspective, I'll speak from an IBM perspective, this is what we say publically, so it's not big secret, confidentiality is that these two things both have a place. And they make sense in concert, in other words, I-- we, kind of, say internally it's a Toyota Lexus model. You invented that, so you know it. But, the idea of that-- I think, it'd be very hard to do only LTO, or only enterprise. But I think it makes a lot of economic sense to do an enterprise and an LTO where you leverage the investment between the two to serve different parts of the market, and it still makes economic sense.

Teale: Let me make a note of that. So you're saying that what Oracle's doing doesn't make any sense.

Childers: I'm not saying that at all -- Oracle has LTO drives.

Teale: But they don't produce them, they don't get any cost benefit.

Childers: I don't know that.

Teale: True-- you don't.

Childers: That is confidential information that you're assuming. But I don't know how one would know that, unless you work for Oracle.

Master: Let me take that--

Childers: Or you're disclosing confidential information you should not disclose.

Teale: You'll have a chance to review this transcript, Ed. And you can go <makes noise>.

Master: I think, I'm going to take that from a one step above that, and some pundits predict that the way the storage market place will look will be Flash drives, and tape, and that the middle medium sized disks, are squeezed out. In fact, the brains over at Wikibon coined the term FLAPE. Which is Flash and tape combined, will be the storage environment.

Teale: I love that vision.

Childers: That's Keniston's.

Master: Pardon me?

Childers: Steve Keniston did that.

Master: Anyway, I saw it at Wikibon. But you said let's speculate and that's what a lot of the pundits would agree that's going to take place in the future.

END OF INTERVIEW

Editor's process notes:

1. Initial edit completed on March 15, 2016 and sent to all participants as:
102738023-05-01_LTO_Virtual_Company_TEG.docx
2. Brad Johns' proposed changes received on May 8, 2016 as:
102738023-05-01_LTO_Virtual_Company_TEG_Johns.docx.
3. Bruce Master's proposed changes received on August 31, 2016 as:
102738023-05-01_LTO_Virtual_Company_TEG_Master.docx
4. Master's and Johns' proposed changes reviewed and incorporated as appropriate and forwarded to John Teale. Teale's proposed changes received on December 12, 2016 as:
102738023-05-01_LTO_Virtual_Company_TEG_Johns_Master_Teale.DOC.
5. No feedback received from Ed Childers.
6. Memorabilia offered by interviewees and accepted into the museum's permanent collection have CHM Lot Number's X7617.2016, X7620.2016, X7677.2016, X7678.2016 and X8091.2017.
7. All unknown or uncertain names, dates, places, products and other facts were verified to the extent possible. A collection of more than 100 pdf documents collected in conjunction with editing this transcript was provided to CHM as Incoming Receipt A2017.5820.