RDBMS Workshop: Ingres and Sybase

Moderator:
Doug Jerger

Recorded: June 13, 2007
Mountain View, California

CHM Reference number: X4069.2007

Table of Contents

RDBMS Workshop: Ingres and Sybase

Conducted by Software Industry SIG – Oral History Project

Abstract: Early employees of Ingres and Sybase talk about how both companies were started and how they developed. Those involved with Ingres talk about its beginning as the Berkeley Ingres Project under UNIX on PDP-11s. Larry Rowe then discusses how Relational Technology (later changed to Ingres) started as a company to market Ingres as a commercial product and port it to other platforms. The founders of Sybase then talk about how they founded the company and their relationship with Microsoft for what became Microsoft SQL Server. They discuss how the companies developed and the significance of the benchmark tests in validating the products. Finally, they talk about how Ingres was sold to ASK and how Sybase grew after it became a public company.

Participants

Name	Affiliation
Doug Jerger	Moderator, Software Industry SIG
Greg Batti	Ingres
Marilyn Bohl	IBM, Ingres
Paul Butterworth	Ingres
Jerry Held	IBM, Ingres, Tandem, Oracle
Mark Hoffman	Sybase
Luanne Johnson	Software Industry SIG
Larry Rowe	UC-Berkeley, Ingres
Stu Schuster	Ingres, Sybase
Roger Sippl	Informix
Steve Weiss	Ingres
Michael Mahoney	Historian, Princeton University

Doug Jerger: This meeting is primarily intended for alumni to discuss their recollection of what happened and why. So it will be pretty open. I'd like to suggest that we not spend a lot of time talking about technical issues. I suggest that we kind of free ourselves up in this session and talk about all kinds of business-related things about each company. I'd like you to address business practices, culture of the company, why you did what you did, why you didn't do what you should have done, those kinds of things.

Jerry Held: Okay, if it's about the company, I'm going to go downstairs because I was never in Ingres, the company. I was only at Ingres the Berkeley project, then I was at Oracle the company. Unless you want to spend five minutes on Ingres the Berkeley project,

Jerger: Well, what do you think? Should we spend some time on it?

Paul Butterworth: Yes, it would be good to do that. And now that Larry Rowe is here, he knows a lot of the Ingres project history and can cover that.

Larry Rowe: Well, it was funny. Jerry left about the time I arrived.

Held: So, let's first talk about the early days of Ingres.

Jerger: We are going to, yes. But since we have two companies, rather than doing one and then the other, I'd like to suggest that we look at time periods, such as 1980 to 1985, 1985 to 1990, 1990 to 1995. And we'll start with Ingres because Sybase didn't get there until 1984.

Rowe: I'm Larry Rowe. I was a faculty member at UC-Berkeley on the Ingres Project from 1976 until I retired from Berkeley in 2003. And I was one of the co-founders with Mike Stonebraker and Gene Wong of the Ingres Corporation and was involved with Postgres at Berkeley but decided not to get involved with commercial Postgres.

Steve Weiss: My name's Steve Weiss and I was at Ingres from 1986 to 1989 in the marketing group.

Jerger: I probably ought to say who I am, too. I'm a volunteer with the Software Industry SIG and I live in Arlington, Virginia. I started in an application software business on April Fool's Day of 1970, which was probably propitious. I was in that for ten years and sold it to Dun & Bradstreet. We bought it back from Dun & Bradstreet and then we sold it to University Computing Company. We almost bought it back in 1982 but they wanted too much money and we couldn't handle that. So, then I went into the professional services business with guys like Jay Goldberg in Software Design Associates out of New York. After that I ended up at a trade association [ADAPSO which became ITAA] and then I came back to the software business and in last few years I've just done consulting for small technology companies. We didn't get big

enough to have the names that you guys had, to our regret. But I have an in-depth sense, in the pit of my stomach, of the problems of software companies and the risks that are inherent therein since I had four kids and a dog and a wife when we started our own business. I left Arthur Andersen as a CPA to go into this business, so that shows you how smart I am. I went from a safe business to a software business.

The point of this meeting is just to talk about the companies and what you feel was interesting about them. And since Ingres got there first, why don't we start with you, Jerry. You can share your thoughts and then we'll move on from there.

The Berkeley Ingres Project

Held: Okay. I can cover the first year of Ingres. That's when I was involved with it. It was in late 1973, I believe, when we met up with Mike Stonebraker and Gene Wong who were the two professors at Berkeley. Mike was only a kid who had just gotten his PhD and was an associate professor at Berkeley. His only interest in life was figuring out how to get tenure and he tells that story a lot better than I can. The name "Ingres" stood for the Interactive Graphic and Retrieval System. The Project was actually funded to do some research in trying to display demographic information graphically and you needed a database to power it. So it had a strange beginning. But we looked at Ted Codd's paper and then all of these things sort of came together. It was exactly the time that AT&T decided to release UNIX and Ken Thompson and Dennis Richie personally flew out and carried the five-megabyte disk and installed UNIX on our new PDP-11/70.

We were the first installation of UNIX outside of the Bell system. And we had a team of two or three PhD students, including , a Master's student named Peter Krepps, who's gone on to do a lot of stuff and founded a company. And there were three undergrads, Eric, Rick and Nick who did a lot of the coding. Eric Allman went on to do Sendmail; he's still doing Sendmail. It was really one of those unique situations where you just had the right chemistry for a project. Many universities were playing around with it. Everybody wanted to do sort of a toy database in memory or whatever. And my background involved Codasyl and some other things and large-scale databases. And so we said, "Why not do something that's real?" and we ended up going down a path of actually building something that could scale a little bit. And the competition with System R at IBM, I think, was also quite good for the group because it was physically close enough that Ted Codd would literally come back and forth. He'd be at our place one week and then at System R the next. And we'd have single variable queries running and they'd have them. Then we'd have joins running and the competition was good. It was five or six kids against these highly paid guys at IBM. So it was a fun, competitive time and we were able to get a product up and running.

We got it out with UNIX and it shipped out to all the different universities, basically in 12 months. By this time I'd graduated and I handed the chief programmer baton over to Bob Epstein and he took it for the next couple of years. And I guess the rest is history.

But I don't know the rest of the history. I don't know what else I can tell you about it. But it was a band of kids against the professionals in a neat little competition. And the fact that UNIX and PDP11's were good environments to develop on and we were all motivated to build something, was, I guess, reasonably good.

Jerger: With the same kind of situation, our company never got to be more than a hundred people. But we had a product that was excellent and at one point we were charging $150,000 for it. We were at some type of session and someone from a big company, said, "You know, we'd be good to be on your client list." I said, "I agree." And he said, "So you probably ought to cut the price a lot." I said, "No." He said, "Well, don't you understand? We're big guys; you're little guys." I said, "Yes. And you have to understand you cannot buy this product anywhere else. It is done and it works." He says, "Ah-hah." So the little guys were honest sometimes. But not always. So how did that comport with what you recall, Mr. Rowe?

Rowe: Well, there was a really important event at, I think, the National Computer Conference in the spring or fall of 1975. There was a big panel that had Ted Cobb, Chris Date, Mike Stonebraker and three or four other people who were standing up in front of the computer industry people and saying, "Here's this new data model. And here's a couple of people who've done implementations of the system and here's why these are important for you to understand." There were probably a thousand people in the room and, as I recall, the first hour was Chris giving an introduction to the relational model. And then afterwards Ted, I think, got up and held up Chris's book and said, "And if you didn't understand what he said, you should buy this book;" Chris's book had just been released.

There were three database implementations that I recall. One was a system in England but I forget the name of it; it kind of disappeared. One was IBM's System R. There was somebody from System R there but I don't know if it was Frank King or Mike Blasgen or Jim Gray. And then there was Mike Stonebraker saying, "And here's this Ingres Project."

I joined in probably the spring of 1976 or maybe the fall of 1975 and I joined Berkeley as a faculty member in the fall of 1976. I went to work with Mike and Gene on the Project because they were looking for somebody to do front ends. They had some programming language interfaces that they didn't like. They had a query language but they wanted other things. So I went to work with them at that point. Between 1976 and about 1979, the system got a lot more stable. When Jerry was working on it there was a lot of interesting function but it wasn't very reliable. They released a version of the code to AT&T and AT&T took the code and actually used it to store all the information about telephone lines in Manhattan. They spent quite a bit of

money developing the code, not in cooperation with us. They did a separate thing and actually later turned it into a product. But between 1975 and 1978, mostly what we did was to continue to develop the code and make it more reliable. We were releasing it to people to use and by 1978 or 1979 we had 300 people using it. We had them sign a 20-page license and charged them $75 or $150 and shipped them a big tape. And basically there were people all over the place running the damn thing, including some little town in Maine that used it for vote counting, believe it or not.

So there were a bunch of people using it, but they were all research licenses. So to kind of speed up and get to the commercialization phase, starting in 1978 we had a lot of people visit the Ingres Project. And they would all come in and look at it and say, "This is great stuff. When are you going to make it a product?" And we'd kind of go, "Well, yes, who knows. That's not what we're into." We had visits from people at DEC and I would bet that between 1978 and 1979 we probably had ten different groups of managers come from DEC and look at this technology and say, "God, this is the most important thing. We're going back to DEC. They're going to have to do something with this. We want to make this a product." And they would go back and we'd all say, "Oh, that's great they're certainly interested." And then we'd never hear from them again. So nothing happened. We didn't really have an attempt to do anything public until Larry Ellison was out selling Oracle saying that they were ten times faster than Ingres and anything else.

Held: Some things never change.

Rowe: Yes. And we didn't believe it, but we didn't know. We'd never done any tests so it kind of tweaked our noses. And we had a graduate student, Dan Reese, who was working down at Livermore. And Dan had a copy of Ingres because Livermore bought a copy. So he could do the benchmarks down there and we found out how really inconsequential Oracle was. We learned that it had no query optimizer. Stu knows this. We used to commiserate about this. And we had people coming to us saying, "Look, are you going to take this?" And so we finally said, "Well, hell, no. We can do as well as they can do." And so we decided we ought to do something if we were going to do anything with it.

Starting Relational Technology

Mike [Stonebraker], Gene [Wong] and I had been teaching and we had some money that we'd made by teaching classes in the summer. And we hired a student to start to work on porting the code. At the university when Jerry was there it ran on UNIX, on PDP11's. In 1978, 1979 we ported it to the DEC VAX and so it was running on the VAX under UNIX. But we knew as a commercial product it had to work on VMS and so we hired a student to go to work doing the port. This student was Derek Frankforth who was working in my front bedroom doing the port. And Paul was an old friend of mine from college and so I called him up and I said, "Paul, you

have to come up here and do this, too." So he came up, bought Mike's house up in the hill and started working in the front bedroom, too. And that's how we got started.

In terms of the financing, we did a dance with a couple of people who wanted to be entrepreneurs to start the company. But the deal never looked particularly good to us. And then we talked to some guys from the business incubator at Berkeley but that deal looked even worse. In fact, the best part of that deal was we met with one of the guys there and he said, "Oh, we've got this great experience for you. We'll help you commercialize this. You should meet with some of our business development people and we'll help figure out how you can fund this business. And we'll give you all the help you need to do this." So Mike and I went to lunch with two of their business development guys. And basically we were asking two questions: "Where do we get funding for the initial development of the product?" And number two, "We need a president and sales force. So we need to hire some people." We asked the business development people how to do that. And they sat there and then explained to us that the way to do it was to get a government contract and we should write a grant proposal to the government to get some money from them to support doing the productization as part of this contract. And so we looked at them; you have to understand that Mike and I were each running research projects to the tune of $300,000 to $400,000 a year on government contracts. And these guys were trying to tell us how to go write a contract for the government? That was crazy. So then we said, "How are you going to help us find a president?" And they said, "Well, what business are you in?" And we said, "Well, this is software for relational databases." And they said, "Who are some of the leading companies in your industry?" And we started to name a couple of companies. They said, "Okay. Usually what you do is talk to people who are division managers or regional managers in those companies and they'll help you do this." We turned to them and said, "Who are these people and where are you going to go to help us find what we need?" And they said, "Well, we don't know any. It's usually the case that you will know who the right people are to go talk to." At which point Mike and I looked at each other and said, "And you're going to take 65 percent of the company for doing this?" So we left.

Jerger: Just as an aside, who would you have thought of as leading software companies at that point? This is now the early 1980s.

Rowe: The database companies were pretty well known. They weren't relational database companies but there was IBM, Tymshare, NCSS, Software AG.

Weiss: Cullinet.

Rowe: Cullinet, yes. Cullinet was a big one.

Weiss: They were all there, but they weren't relational.

Rowe: None of them was relational at the time. And, in fact, we did a dance with John Cullinane because our first president was Jon Nackerud who had worked for Cullinane for years. And when we got the offer from the venture guys, Sutter Hill up here on Sand Hill Road, Cullinane called up and said he wanted to come out and make a counteroffer. We said, "Okay," and he flew out and tried to convince us that we didn't want to start a company ourselves, but that we wanted to join Cullinane because it was such a great company. We kept trying to say, "We don't see how a relational system fits with a Codasyl system," so that was kind of abortive. But we ended up getting the funding from Sutter Hill.

Butterworth: We delivered the first product, it must have been early 1981, February or March.

Rowe: Right. And our first customer was DEC which we thought was absolutely great since we could never convince them to do anything earlier. But our first substantial customer was Schlumberger. We did a deal with them and that led to a lot of business over the years.

Jerger: Did going through Sutter Hill Ventures help you get the management support, the president, the CEO, that kind of thing that you wanted? How did that work for you guys?

Rowe: Well, Sutter Hill were venture capitalists. And venture guys do certain things really well. And they certainly did okay. I wouldn't say they did great but they did okay. We actually found the first president ourselves, John Nakerud. He had been working at Cullinane as a kind of a special assistant to the president. But they brought in Gary Morgenthaler who started out as executive vice-president but eventually ended up as president and ran the company for many years. So in that sense, they contributed quite a bit.

Jerger: Well, we have a richness of people who were early in Ingres, I think. Steve, when did you get there?

Weiss: 1986.

Jerger: Greg Batti was there in 1982.

Greg Batti: Yes, I joined in early 1982.

Butterworth: And I joined in 1980.

Rowe: So you were probably employee number two or three, right?

Butterworth: Yes. Officially I had employee number 0001 because at that point Derek, who was really the first employee, was just a part timer. He was kind of a consultant and so he didn't have an employee number.

Rowe: He was still a student.

Jerger: You were the first guy on the payroll in a regular way?

Butterworth: I was the first official guy there to do work. But actually Derek was the first one there to do work. Derek and I started out in Larry's spare bedroom trying to get the code from the university to run on the VAX and VMS. And it looked like we had code that actually had been modified to run in a VAX environment. Because in the PDP11 environment they have the small address spaces, so it broke up into multiple pieces. The code had been modified to run on the VAX but it wasn't clear that it had actually ever been run because there was a lot of stuff that you would just look at and say, "That code can't possibly work." We spent about six months trying to get the code to really run in the VMS environment, first in Larry's spare bedroom and then, once the company officially formed, in Jon Nackerud's basement. That was the official home of Ingres or Relational Technology in those days. And so, yes, we got it running. It was probably about February when we started running.

Changing Company Names

Rowe: It's probably important to talk about the names because the names changed a lot over those days. Oracle started out originally as Relational Software.

Held: RSI.

Rowe: And we were Relational Technology. We were actually going to be Database Technology and we had reserved the name with the State of California but then it expired. We hadn't incorporated because we hadn't completed the deal with the VC's. You can reserve the name for something like 60 or 90 days. And we had somebody go to the place where you reserve a name the morning that it expired. But before they got it in, some kid from Arizona got the name. And so the name Database Technology wasn't available. We needed a name because we were going to incorporate but we couldn't come up with one. Mike and I drove back from Palo Alto together one day and we spent the whole time talking about names. Eventually we came up with Relational Technology, mostly because we had database technology and relational was the right phrase and so we picked that. The craziest thing was that our phone would ring and when we'd answer it the person would say, "Hi, we're interested in your Oracle database." And our people would say, "You mean Relational Software. We'll be glad to give you their telephone number but by the way, let us tell you about us, Relational

Technology." At other times, checks were sent to us that were supposed to go to Oracle and checks were sent to Oracle that were supposed to come to us. It was crazy until the names got changed. Roger, when did Informix change its name? Do you remember?

Roger Sippl: We were Relational Database Systems until 1986 when we went public. And I think a porting agreement that IBM was trying to send to Oracle was faxed to me. We were Relational Database Systems then.

Jerger: I'm glad to hear you say that you weren't sure if the software had ever been run because usually, when everybody looks backwards in the rearview mirror, they say, "Oh, the software we created was just wonderful and it was always great." But on the other hand, it seems to me in the early stages, when you would talk about software, people who were knowledgeable would ask, "Has it been written?" And then they asked, "Can you run it?" And eventually we got to the point of saying, "Does it produce what it's supposed to produce?" Have we gone beyond that now? Does that sound right?

Butterworth: Is it right? Yes.

Jerger: So we have gotten better along the way, I think. Paul, is there anything you want to add to what Larry just covered?

Butterworth: Well, that was kind of the beginning and, yes, we got the first products out so that was early in 1981. And at that point all the parts kind of worked. It wasn't the world's most robust software but it definitely got the job done. From that point forward, the big issue was trying to gear up the sales effort and gear up the technical effort to mature the product. And one of the big things we did was we started building better productivity tools on top of the core systems. We hired a bunch of people to work on that problem. And at that point, then, the competition with Oracle got to be the big thing because we were sort of neck and neck battling it out for pretty much every customer in the early days.

Jerger: When did you actually change your name and when did Oracle change their name?

Rowe: I think it happened pretty close to the same time.

Batti: We were a little later, actually. They changed their name when they went public and for a while we kept our name, thinking having the only "Relational" company name was an advantage for us. Oracle was Oracle, but we were still Relational Technology for awhile.

Held: But I think it was around 1988 when we switched names

Sippl: Oracle went public in 1985.

Batti: Yes, we were the last one. We were the last one to call ourselves Relational "something."

Weiss: But you can see the change in process and it would show up in several things like business cards. There was a succession of cards through the years where first it just had "Relational Technology." And then it added "Makers of Ingres" in small type. And then the name Ingres would get bigger and bigger while Relational Technology would get smaller. You could see which way the winds were blowing.

Sippl: Well, the customers did it to us. They started calling us all by the name of our product.

Marilyn Bohl: Who could blame them?

Rowe: People didn't think of these companies by the company name, that's true. Sybase had that right from the beginning, you know. They had the product and named the company for the product.

Jerger: Yes, that's right. They did it right, yes. Well, Fortex always wanted to keep the name of the company separate because we were going to have this huge pile of products. We didn't want to limit ourselves. So, Greg, what area of the company were you in in 1982, 1983?

Porting Ingres to UNIX and Other Platforms

Batti: I joined the engineering organization. I guess I was about the 15th person in the company at the time. It was still a pretty small company. And the first thing that they had me working on was actually getting the system to work on UNIX again. At that time we only had a VMS system; that was all we were selling. We had one VAX so that was one of the first things I got to work on. I worked at night on the VAX to get the UNIX system working. Bob Coy was also working on his optimizer at night. We worked at nights for about four or five months before we got our second VAX and then we got all the other UNIX machines. That was also the time when all the different UNIX machines were coming through from a number of small companies. And we ported all our stuff to that and worked with AT&T to get our system working on their 3B5's and things of that nature. Do you remember all that?

Rowe: Oh, yes. And this is the guy who did it and took the grief and made it work.

Jerger: Did you have your own hardware or would you rent it or lease it?

Rowe: No, we bought hardware and we started off using a spare machine up at North Berkeley Labs. They had a spare machine and they were willing to let us use it if we would give them the license for Ingres, so that was the deal we took. Eventually we ended up buying VAX 750's and at one point we had six or eight of them in a real small room. And eventually, because we were porting to so many machines, we had 50, 60 machines in the machine room when we moved down to Alameda. And it was huge.

Batti: Right. By that time we were even running on mainframes.

Butterworth: I think at the height we were supporting something like 43 different platforms.

Jerger: At this point, it's 1985 or 1986. I think we've done a pretty good job on the first five years or so. I'm going to hold back on Weiss until we get to the next five years. Is there anything else we should cover? Does anyone else have some stories like what Steve was alluding to when he talked about how the business card changed. Sometimes the fun stuff comes about when people just write stories. We found this out when we worked on the IT Corporate History Project that Luanne referred to earlier. Because each story will trigger another story and it actually gives an idea of the character of the company and its culture. And you learn a lot from those stories. So please, if you get triggered by something here, don't hesitate to pop it in because that's good stuff.

Rowe: I'll tell you one that surprised me. Early on, we were just a little company in the garage and we didn't have any money. Gary Morgenthaler wanted us to look professional so he wanted the manuals to look really good. He arranged to get special paper that had nice gray outlines and the name in red and a place in the upper right-hand corner where you'd put the page number and then the name of the product over here and it would look professional. And so we would print this stuff out on, if you could believe it, daisy-wheel printers at the time. And then we would take it down to the local copy center over on Durant and they would copy it onto the special paper that we had. At first we had a real problem because they didn't pay much attention to getting it to register correctly. And Gary would go crazy. Eventually there was one woman over there who had a night shift who understood that it was important that the stuff fit in the boxes. She was really good. And so she printed our manuals. At one point we were spending $30,000 a month in the copy store because we were shipping a lot of systems and in those days when you shipped a system it included boxes of manuals.

And then at one point Gary told us that we had just bought a new Xerox that was going to produce the manuals. The thing cost $500,000 or $600,000 and he'd hired the woman from the copy store to run the publication operation at the company. And we said, "Gary, we haven't got the money. How are we going to do this?" And he wrote down on the back of an envelop a couple of numbers that showed that the savings we got from printing on our own in-house machine made up the cost of that $500,000 printer in less than six months. I was floored. And

it was at that point I really understood what it meant when people said that when you're in the software business, you're really in the publishing business. Because we were producing books, that's what I think it amounted to, and I was floored.

Luanne Johnson: Did you get the answer to the question about when RTI went public? .

Butterworth: It was 1988.

Rowe: I'm sure that it was 1988 because the stock market meltdown happened in October of 1987. We had a board meeting the Monday after the meltdown and we were already scheduled to go public. Well, I can tell that story later.

Jerger: Yes. I want to just throw one thing in here just because it's interesting. Paul and I were talking on the side before the session started about some of the joy of having been in this business in the past and seeing what has happened to the people who work for you or with you along the way. I just happened to run across a fellow named John Newton as I was gathering some information to think about this. Do any of you know about him?

Rowe: Would you like to hear the story of the origin of John Newton?

Jerger: Okay, and then I'll read you what he says.

Rowe: During the fall of 1976 as a young assistant professor, the first thing I was told was I had to open my doors for office hours for new students. So the first day I open my doors and in walks a young freshman who says, "Hi. My name is John Newton." And so John was my advisee for four years. At the time he was convinced he was going to go into the Navy. He was in the ROTC program at Berkeley, believe it or not, in the 1970s. My dad was in the Navy for 30 years and I understood what that life was like so I took it as my task to spend the next six months to convince John that he did not want to go into the Navy for a career. Actually, a lot of the early people we hired were students either in my class or advisees or people that I met at the university. Greg was in the compiler class, Roger was in the compiler class and, frankly, if you look at the first 20 employees, probably 15 of them took classes with me. There was a bit of an underground there. So John went to England; let's see, when did he go to England?

Batti: Well, he went to England, then he came back and founded Documentum.

Rowe: Right. And he did the first of the bi forms interfaces so he was the first with the forms-based user interface.

Jerger: Well, let me read something he wrote, "Major influencers of my career" In it he said, "I started my career in 1981 at Ingres," I'll just read it quickly, "a database company started by my professors at UC Berkeley and recently resurrected out of CA as an open source database company. I spent nearly ten years watching that market grow from infancy to maturity. And from this early entry, was able to span development, sales, marketing, et cetera. Ultimately ran the database group at Ingres. Chief architect, Paul Butterworth, had an important influence on how I look at architecture and how to manage developers. Mike Stonebraker and Larry Rowe, my professors at Berkeley who started Ingres, helped me to see the larger picture of how technology affects the market and how it evolves over time." Isn't that nice to be said about you? You guys did a great job.

So let's see – in 1984 Sybase was founded by Mark Hoffman and Bob Epstein out of Mr. Epstein's home in Berkeley, California. Was that in a basement or a bedroom?

Building Sybase

Mark Hoffman: It was a back porch that was screened in and glassed in.

Jerger: Well, tell us about your experiences, please.

Hoffman: Well, we were both at Britton Lee. Britton Lee worked with the database machine and it had the concept of taking the database machine and putting it on general purpose hardware. Being able to have a very high performance relational database network is what interested us. Bob had left the company and he came and talked to me about it. I tried to do some work out of Britton Lee and decided that that was not a good idea. So then I left and Bob and I just started to put this thing together. That was in, probably, August of 1984. Then it took us nine months to raise capital at that point in time. It was a very difficult period. In 1984 the venture people weren't investing much of anything.

Rowe: Who were the venture capitalists?

Hoffman: It was Kleiner Perkins, H&Q and TRW. We raised one and a half million dollars.

Rowe: We talked to H&Q and talked to Larry Moore, as a matter of fact. And he wasn't quite sure he should do Ingres. I ran into him a couple of years later and he kind of looked at me and said, "So I probably should have done that one, shouldn't I?" And I said, "Well, I think so." And I think it wasn't too long after that that you guys got started.

Hoffman: Well, the way we got in was, again, because TRW was extremely interested in this stuff for a lot of intelligence work that they were doing. The guy at H&Q at that time was

Tom O'Rourke. He was coming out of Tymshare so he understood the high performance requirements and the hosting and the network stuff. So he put us together with a guy by the name of Bob Evans who had left IBM and was also consulting for H&Q. He said, "I'm interested but you have to pass our test." And so Bob worked us really hard to justify what we were doing and what the edge would be.

Eventually they came around and said, "Yes, we're interested in doing this." And Kleiner went through a lot of stuff, too, to finally get there. But they all signed up in about March or April of 1985. Bob and I had been funding it ourselves and running up our credit cards, so we had a lot of debt. We had talked Sun into giving us a machine and we were just going to get funded in the next month or so and we'd pay them back.

Jerger: Did you use credit cards for Ingres, too?

Rowe: Not really because we had the money from the summer classes we taught. So between Mike and Gene and I, we each had somewhere around twenty to thirty thousand dollars total between us. And it was mostly the money we'd made in the summer.

Hoffman: Well, we ran up about $60,000 debt between the two of us. It was a lot at that point in time and we were close to being pretty broke. But anyway, we got the money. We brought two people over from Britton Lee. One was a guy by the name of Tom Hagen and the other was Jane Dowdie. During that time we were designing the database and designing what we wanted to do. And closer toward the end of it, we had actually started to hire some people and ramp up and work on this stuff. But once we got to 1985, we were off and running and got one and a half million dollars. I remember we were driving over to pick it up. I had this old Volvo and the bumper fell off right in the middle of the San Mateo Bridge. So we drove over it and, of course; you're not going to stop on the San Mateo Bridge and pick that bumper up. I thought that was kind of the end. I said, "Boy, we better get this money."

We got the money and we were off and running with building the company. And it took us until right at the end of 1986 to build the product. And we started to ship the product. Our first customers were Picatinny Arsenal and British Petroleum up in Canada. But it didn't work very well. We had some real problems in that version and it took us about six months to get it repaired.

Stu Schuster: Yes, that was an extremely well kept secret because our whole positioning was that we were the first relational database for online transaction processing in a network environment. Except, it only ran one user because our multi-user technology didn't work! And I was thinking, "We're going to run a great benchmark as long as it's single user."

Hoffman: Yes, and I was on the East Coast and just before I left to fly back west, I called Tom Hagen and he said, "It's not working." He was telling me what was going on, and my stomach was just churning. But then I got on the plane and, of course, it's a six hour flight back, and all I knew was that it wasn't working and we heading for bankruptcy or something.

Schuster: Well, the good news was that everybody was working in development. They were all building applications. No work was being done for production because we had just made this huge announcement that we were entering the transaction processing business but it actually worked with just one user. What saved us was customers had to build their applications and that gave us time to get it working.

Sybase's Relationship with Microsoft

Hoffman: Yes. So anyway that got us started. Also at that time we got a call from Microsoft and they said they were interested in coming down and talking to us. Bill Gates actually came down and looked at it and said, "You know, I'm interested. Let's follow up with that." And so we started a series of meetings to work with them. And, of course, that eventually turned into the Microsoft SQL Server in a deal we did with Rationale that they would go after the low end of the marketplace. We were definitely going after the Fortune 1000 and they wanted to go after the low end of the marketplace. And we finally decided that there was a clear division there and that this was a great play to do. So we did sign that deal with Microsoft. And, of course, once that got announced in the marketplace and Microsoft was taking off, that turned out to be a huge thing for us. So our revenues just exploded after that. I think in 1987 we did $6 million. We then went to $24 million, $52 million, $100 million. And so it was just huge growth after that.

Sippl: Were you getting any royalties of significance from Microsoft for that deal?

Hoffman: Definitely.

Schuster: You said "significance." We were getting royalties but were they significant?

Hoffman: Yes. Well, we received millions off of it, tens of millions.

Schuster: But right at the beginning, there was none.

Hoffman: Right, at the beginning there was very little.

Sippl: The burning question I've had for years is how did that deal fall apart?

Hoffman: The problem was that we were growing like crazy and we were totally focused on the enterprise marketplace and we were looking at building this other database which was an object database. We had actually started to build that with Greg Olson. Microsoft came back and said, "You're not servicing our product. We have to have these features and functionality for the low end of the marketplace." And we said, "We have no way to do it. We can help you build it. We can tell you how, but we can't do it." In hindsight we were not servicing them very well. And so that became such a point of contention between the two of us that almost every meeting was just a conflict meeting.

Rowe: One of the questions I have about this was it seemed to me when the deal was signed my expectation was that Microsoft was going to do their own product. And they were just doing this as a way to get into the marketplace early. And that they were going to use their own product eventually. Were you thinking about that when you did the deal with them originally, that they were eventually going to produce a competitive product, or was that not part of your thinking?

Hoffman: Yes, I think they were working on it and just realized it was going to take them some period of time to do that. We also thought it was also possible that they were going to do a deal with you guys or somebody else.

Sippl: Well, we tried to do that deal. We were in the running for that deal but it didn't get very far. And all of a sudden it was announced that you got it done. I was mystified as why they didn't go farther with us because we had a version running on MS DOS.

Hoffman: At no time during the period of time that we had a contract with them were they working on a product. They were just taking us and commercializing us.

Schuster: And remember in the beginning it was the Ashton-Tate Microsoft SQL Server.

Sippl: Right, it was a three-way deal. Ashton-Tate tried to acquire Informix just prior to that deal going down. When they failed to acquire Informix, we were trading at $21 and they offered us $21 and a half or something. About six months after we said no to that, this deal got announced. And we had been trying to do that OEM deal that you guys had with Microsoft. And then a couple of years afterward, it was around 1989, I decided that I was going to try to integrate office automation with databases so that you could do index searches for everything. I tried to sell Informix to Microsoft. Not very hard, but I got Gates at a trade show and I said, "This is what I want to do. Why don't we do this together. The best way to do it is if you buy Informix." And he said, "You know, we've got so much invested into this Sybase work that we'd just have to throw away too much code." So I don't think they were working on anything else. I think they had so much work to do just to put it on Windows OS2.

Rowe: Stu, do you know when Phil Bernstein went to Microsoft?

Schuster: I don't remember.

Rowe: Because we talked to Phil at one point and he led us to believe that they were actually working on a thing.

Schuster: They might have but it wasn't very serious.

Sippl: Well, they had a whole group that did do some indexing, database-ish stuff.

Rowe: It's funny because we saw that deal go down and my recollection is that we never got a shot.

Schuster: The funny thing was I got the impression during the negotiation that they were talking to you guys. So they were masters at playing both sides.

Rowe: I heard that that they were. I heard that we were renegotiating with that. But I wasn't involved with it.

Schuster: You'd have to check with Gary.

Butterworth: We did talk to him. But I don't remember it being at a really high level.

Schuster: But the main thing for us was we knew two things. We knew we would never have the resources. We had a hard time porting to the major UNIX environments and dealing with VAX in the beginning. And then we were thinking that, "We're going to have to do a mainframe version and we may have to do these fault tolerant versions." And we also had our own product plans. And the idea of doing a PC version when things were going so well for us, especially when we were struggling with the amount of things that we had to do on the enterprise, just didn't seem possible. The second thing was I actually believed early on in the negotiation that it would end some day but I just didn't know when. But during that period, whenever that period was, three years, five years, seven years, we would never get to a PC version. And then there was no such thing as NT even being discussed.

Hoffman: It was OS2 at that point.

Schuster: And that didn't look capable of running it.

Rowe: Cutler hadn't left DEC yet so NT wasn't even in the works.

Tandem and Transaction Processing

Schuster: Yes. It did two things. In our early days we had to focus, focus, focus because we were late to the market. And we had to focus on the right applications because we could not go up against Ingres with tools, Informix with VARs, Oracle with portability. We ran on two platforms. We were the Model T. You could have us as long as you liked VAX and Sun. So we had to think about the marketplace really hard, about where we were going to compete. And it was from my work with Jerry Held at Tandem where I saw this opportunity and said, "This is the first relational database that could do transaction processing." I was at Tandem for five years.

Rowe: Tandem did transaction processing.

Schuster: Yes, but they weren't relational. They had a relational query language and they weren't portable. They weren't UNIX. When I was at Ingres I saw one out of twenty customers come in and say, "We were just talking to Britton Lee and wanted to know about multiple users and up-date intensive." And when I got the call to talk to Mark Hoffman and I saw the technology I said, "It's a niche right now but it will be major." In the mainframe market it grew to be half the market because it started Decision Support.

Held: It's interesting to look back and see from an historical perspective what happens in markets and how there are missed opportunities. Before I got into that, a headhunter who was looking for a marketing guy for Sybase called me and I said, "I know a guy who's going to be the perfect guy for you." I said, "Call Stu Schuster."

Schuster: But the funny thing was that I was on the road launching Ingres Star for Ingres when Bob and Mark were on the road talking about Sybase. In those days, when you got back from a trip you found a bunch of pink slips on your desk for phone calls; that's what the secretaries took for messages because there was no voice mail. And I separated them into two piles. One pile was calls from people I knew and other calls from people I didn't know. I went through the ones that I did know. In that pile were ten messages from a recruiter that I didn't know was a recruiter. And I said, "Well, if somebody's calling me ten times, that must be important." And as I was talking to him I was thinking, "There's no way there can be another database company. There are already too many. I can't even count them all on my hand. I don't even want to bother." He said, "You have to come and talk." And I said, "It's got to be in the East Bay because I'm not commuting." And he said, "It's in the East Bay." And I was thinking, "Who the hell was that?"

Held: It's interesting to look at certain opportunities that came up, as well as missed opportunities, and the transition of the whole industry from a hardware-based industry to a software-based industry. I think Apple and Tandem actually were two massive missed opportunities because they were basically hardware-driven companies. And Apple has recovered through the iPod; it's gone through many iterations. But if you look at where Apple was at a similar point in time, they had the best software, by far, and if they had separated it from their hardware, they would be Microsoft today. And Tandem really was a similar kind of missed opportunity. If you look at the database industry, in the early 1980s I don't think there was anybody who could touch Tandem in terms of transaction processing. What Sybase wanted to be, Tandem already was by far. They had high performance, distributed multi-processor, non-shared memory, recovery, five-nines reliability. But it was tied to the hardware and it was a hardware company that could only see selling hardware. If that software had been removed and put on other hardware and sold as general purpose software…it was way ahead of Ingres, Oracle, Sybase at the time. It was a huge, huge missed opportunity.

Schuster: Yes, people don't realize it but Jerry Held had built the first relational query, production query language on their file system. And I came in to help manage the programmers who managed me, basically.

Held: All the big companies were running production on it in high volume; it was a classic thing. You know, if you look back to learn from history for the missed opportunity, it's hard when you don't know what business is really the business you should be in.

Jerger: We weren't in your business but when people talked about large transactional volume I always said, "Oh, talk to the Tandem guys."

Held: But it was a system. It wasn't just software.

Butterworth: Well, Britton Lee, as I said earlier, was in the same business.

Sippl: They had their own program?

Rowe: No, they had their own programming language and their own operating system.

Sippl: And their own hardware, their own processor.

The Teradata Company

Rowe: They had everything. It was kind of like Teradata.

Sippl: It was totally proprietary so porting it wasn't really an option.

Rowe: Well, it could have been. And it eventually did. Didn't you eventually port it?

Held: Well, do you know where that software is today? HP has just announced a major product called Neoview. Do you know what Neoview is? It's the non-stop SQL. It's the thing that we started 30 years ago. HP's new database appliance is based on our product from 31 years ago.

Rowe: So that's actually a pretty smart move on their part.

Held: Well, Hurd knows about that; he ran Teradata and is the CEO of HP.

Rowe: Speaking of other database companies that were around in about this time, did you see the business plan for Teradata? .

Held: Well, the Teradata guys were some of the first customers of Tandem. There were some guys in the Citibank Group who loved the architecture and said, "We could do query processing on this architecture," and they went out and got it. The Teradata architecture is a rip-off of Tandem. It's the same multi-processor bus architecture.

Rowe: Well, except they did their own network.

Held: No, it was the same. They changed the names and they changed the application, too. Their focus was query processing; ours was transaction processing.

Rowe: Not originally. The original business plan for Teradata was transaction processing and they claimed they were going to sell it to banks. They had read Jim Gray's paper that said if you're going to do transaction processing you've got to do hash access methods and that's the mistake that everybody else made in that they didn't do relational. So originally, their only access method on their system was a hash-access method. What happened was the banks were never going to put it into production because it was too risky. And some other people at the banks discovered that you could use it for data mining and decision support type queries. And, oh, by the way, what is absolutely the worst access method for doing queries that are over groups? It's hashes. But the thing still ran much better because they did that. No, they found the decision support market; that was not part of their original business plan. Their original business plan was transaction processing.

Stu Schuster's History

Jerger: Stu, how many years were you at Ingres?

Schuster: I was at Ingres for two years, left for a year and then came back for six months before I was recruited by Mark and Bob at Sybase.

Jerger: That was when? In 1986, 1987?

Schuster: 1986, yes. That's when I joined.

Rowe: He has a much better history than that. He started out as a professor at Toronto.

Schuster: The University of Toronto.

Held: But then he went to Tandem.

Schuster: And it was amazing earlier when we did a tour downstairs in the Museum. As a grad student I wrote code for the ILLIAC IV and I was wondering if they had one down there. And they did.

Held: We flew back together from the VLDB conference in Berlin and it was on that flight I recruited you to come to Tandem.

Schuster: Yes. I started out at Tandem and I wanted to work in marketing but they wouldn't let me. But I knew in my heart that that's where I ultimately wanted to be. So I worked two years in product development. Jerry wanted to work on advance systems so I started to run the tools area of the database in the languages-compiler group. And then I went out into the field because they all said if you're going to work in marketing, you have to go out to the field. So I spent three years in product support, tech support and sales. And then I got recruited by Mike Stonebraker to be one of the first product marketing guys at Ingres. And I worked my way up in the marketing group, after working my way down from technology.

Johnson: Greg gave me this photograph of Ingres employees and they carefully went through on the back and identified everyone except there's one person identified as, "The marketing guy who went to Sybase."

Early Sybase

Schuster: I was approached by someone from Oracle and he said "I know what you're building because I'm talking to the Sun guys." And he said, "So you guys are going to do an engineering database." And I told him, "Yes, we're doing an engineering database. It's obvious." I did that because I didn't want anybody to know what we were really doing. And when we went out to announce, there were two things that we had going for us. First, we had real customers since we did not announce Sybase until we actually had customers. And second, we were in production. That meant that we could demonstrate a working system and they could call actual customers. We had briefed the analysts just before this so they could call the customers. And then I called this gentleman back and I said, "Well, I was kind of fooling with you. We were really going to do transaction processing." Because I knew he would be directly talking to Larry Ellison. Everyone was saying, "Don't worry about them. They're not going to be important. They're just going to go after this little niche." And I think we really took the market totally by surprise.

Jerger: I'm going to do a Tim Russert thing on you here. Do you recall, Mr. Schuster, being interviewed by Angel Mehta back in 1987?

Schuster: Oh, yes.

Jerger: And what did you answer to this question? "You were at Sybase in the early days when it was battling Oracle. Tell me, if you can, what did you learn from Sybase about battling a bigger company?"

Schuster: Well, I think that I said a couple of things. First, one of the early decisions that we made about marketing was that we were going to only compete against one company. We just could ignore everybody else. And that, I know, pissed my friends off at RTI. But I said, "We have eight companies to compete with but we're going to make this a two-horse race." That was job one. Anything we do competitively, we're just not going to mention anybody else but Oracle. I was hoping that Oracle would get so mad at us that they would think about us. And people went to Oracle and there was an hour-long seminar during a day visit on how Sybase was no good. People had not heard about us. We figured if they're spending that much time on us, they ought to come visit us. We got a lot of leads from Larry Ellison.

And the second thing was we had to completely differentiate ourselves because in the early days, the sales force was getting slaughtered because we were just on two platforms. We just couldn't compete when people hadn't made a hardware decision; they were finally making software decisions first and making the hardware decision later. When they came to us they had to take A or B, that was all that was available. So I decided we were going to go after a totally different class of applications.

Jerger: Well, by God, you got an "A" for that one.

Schuster: Yes, I did get that right.

Silver Bullet Benchmark Wars

Weiss: Well, by late 1987 or early 1988, Sybase was a real challenge for Ingres. And I remember a couple of things that were going on. One, of course, was we were in the middle of a major rewrite for Release 6, moving to an SQL base for our technology. Paul can speak to it much better than I can. But this was a huge project; it was a massive rewrite and at the same time we had performance challenges coming up out in the marketplace with customers asking where we were. So we did two things. One, we went back to some skunkwork things we were doing as we were asking, "What advantages could we portray about our new Release 6 architecture?" Down in the innards of Release 6 was our multi-server architecture. And so that became the new Ingres Release 6 with multi-server architecture for increased fault tolerance and scalability. And this, of course, was at the time that Sequent and Pyramid and all these others were coming out with their multi-processor machines. So the idea was in tune with the times with more porting and all kinds of things like that. We had this good story on fault tolerance, so who was going to argue with that and linear scalability. So now we had a story to hang our hat on for this Release 6 which was still unproven and who knew what shape it would be in. But at the same time that we did that, you'll recall the Silver Bullet benchmark wars were happening.

Rowe: Yes, I was going to ask, when did Silver Bullet happen?

Weiss: It was late 1987 or early 1988. I think it was early 1988.

Rowe: When was it? When did you spend that month up there? Sequent was one of the early parallel processors built running UNIX systems up in Oregon, wasn't it?

Butterworth: Beaverton. And Casey Powell was the CEO up there.

Rowe: And we had a contract with AT&T for an application that all of the VAX-cloned, UNIX companies were desperate to get into. And the only way they could get it was if they ran Ingres. So they came to us to port it to their platform. Casey came down and basically made love to the secretaries in the place to get Gary Morgenthaler to call him to do a deal for this thing. And eventually we did a deal with him. But then we were in this performance war with Sybase and with Oracle. And there was one guy up at Sequent, Gary Kelly, who really wanted to do parallel processing and none of the database companies would do it. We were kind of interested because we were interested in parallel processing. And eventually Paul went up

there and spent a month and the question that we asked was, “How fast could you run on a parallel processor doing TP1 which was what the banking transactions used? What would you have to do to make it go as far as possible?” So, Paul, why don’t you tell him a little bit about what you did and what you ultimately found out?

Butterworth: There were other pieces to it because about the same time the Tandem guys with non-stop SQL had announced that they could do a hundred transactions a second. So we decided if we could do a hundred transactions a second; that wouldn’t be a problem.

Rowe: On the top of the line VAX machines we were doing maybe 15, 20 transactions a second, I think, something on that order.

Schuster: That’s about right, back then.

Butterworth: No, we felt that if they can do it, we can do it. And so we did. We started working on the system and, of course, when we started we couldn’t get anywhere close to a hundred a second. So we did a significant number of changes to the kind of internals of the system to increase concurrency and do the caching and the rest of that stuff. And eventually we ended up taking the software up to Sequent because they were the only ones that had big enough machines. They had these 32-processor machines that we needed to run this thing on. And so we got up there and we started running our benchmarks and those guys did the classic thing: they scaled up to about three processors and the whole thing just flattened out. And we’re sitting there looking at this thing and we spent hours and hours and hours between the Sequent guys and our guys. And we finally figured out by talking to the Sequent guys for awhile that the problem was actually in the machines. They were using a write-back main memory cache so every time you wrote some data, pretty much the whole bus stopped and resynchronized itself. And we were looking at this and saying, “There’s no way this is going to work.” And we started talking to one of the other engineers and he said, “You know, we just revised the processor boards. You guys are running Rev A boards.” It turned the Rev B boards had a better caching strategy. And we actually went down one night and we took the first boards off the production line. I think they got into a lot of trouble for that. But anyway, we pulled all the boards off the production lines, took them upstairs, put them in the evaluation machine and ran with them.

Schuster: So you sold them as used equipment up there.

Butterworth: Yes. And we started rerunning the benchmark. And immediately we were up to something like 95 transactions a second. And we said, “Okay, if we are getting 95, this is no problem. We’ve got to get to a hundred.”

Schuster: What year was this?

Weiss: I'm pretty sure April 1988.

Batti: It had to be earlier than 1988 because by then we were working on Terminator.

Schuster: Because the benchmark wars had started already by then, right?

Rowe: Well, the benchmark wars started back in 1982.

Schuster: No, but I mean, the real ones because we were the ones that pushed it because we had this TP processor. We started using it as marketing data early on because that's what we were about.

Rowe: Right, so you were using it and we all started to optimize that part of our systems which had not been optimized before.

Schuster: Right.

Jerger: So why did you guys start calling it the Silver Bullet benchmark? This is 1987, 1988?

Rowe: Well, whenever this experiment got done, Ellison was marketing Oracle as: "I'm the best at everything." That's been basically his marketing strategy, irrespective of whether it's true or not. And he was always saying, "We're the fastest. We're the fastest." So we thought it would be interesting to see what we could do if we really ran on a parallel processor. And ultimately this would be technology that would eventually get in a product and go out there. And so the benchmark got done and we got whatever the magic number was, 105, I think it was.

Butterworth: 104

Rowe: Okay, 104, something like that. So we then created an advertising campaign called, "The Silver Bullet Benchmark," and it was really just an advertising campaign because nobody had the processors, nobody was really running this code, although it did get into the product.

Weiss: Well, the other piece of it, though, was when I was the product spokesperson to all the press for Release 6 with multi-server architecture, I let out the news of, "Look, 104 transactions per second in the Silver Bullet benchmark wars." I was always very clear to

differentiate between product, lab prototype, experiment and how fast could we go. Somehow I knew that the headlines would always get collapsed as "New Ingres Release 6 breaks a hundred transactions per second." And that happened consistently in all the write-ups we got in the press. Somehow all of a sudden this would put to rest the performance questions about Ingres for a period of time.

DeWitt and other Benchmark Tests

Rowe: No. The performance questions are always there. If we're going to talk performance, we've got to go back and talk about the earlier performance wars with Dave DeWitt and the DeWitt benchmark.

Schuster: Well, that was more of a mixed query benchmark wasn't it?

Rowe: But it's important. And it's important for two reasons. One: when relational systems came out, the big question was could they run fast enough? And between System R and Ingres, the research projects, the view was always that System R was the commercial system; it ran fast, it was production oriented. We had a toy that had a nice function but wasn't really good at performance. So Dave DeWitt at the University of Wisconsin developed a benchmark to test query optimization and the query processing part of the system. And it became the benchmark that everybody ran. We finally got a chance to benchmark against Oracle and it was another student, Rob McQuarrie, who went to work for another one of these parallel processing companies. They had the Oracle source and he ran the DeWitt benchmark on Oracle and he ran it on Ingres because he had both systems. And we blew Oracle out of the water. So we went out and advertised that. Ellison went ballistic. He went absolutely three-star ballistic and did several things. One thing he did was to call Dave DeWitt in the middle of the night and threaten to have him fired from his faculty position at Wisconsin. And the second thing he did was he wrote a clause that is in the Oracle contract to this day. It says, "If you license this code, you are not allowed to run a benchmark and tell anybody about the results." And that's in there purely because of that one thing.

Weiss: When was that?

Rowe: That must have been 1983, 1984. So then it was roughly 1983, 1984, 1985, when IBM finally took System R and turned it into a product and it first went out as SQL/DS. DeWitt was able to take the DeWitt benchmark and run it on SQL/DS and compare the numbers against Ingres running on CMS. I guess they had moved to CMS because it was a CMS benchmark, I believe, and he ran them head to head. And what absolutely floored us was we were much faster than System R ever was. And we never thought that we were anywhere close to as fast as them because speed was always their focus. We were blown away by that; just absolutely blown away. So the benchmark wars started right at the beginning.

Schuster: Yes, they were query oriented which we took advantage of.

Rowe: Well, at the time, the issue was, were these systems fast enough to be usable? And so queries were what people were doing. They weren't running transactions because it was production. They weren't running for production. But that was kind of there from the very beginning and then once everybody optimized their systems on query processing, we were all roughly comparable. Then the next round of benchmarks was on TP1 which was transaction processing.

Schuster: And one of the Sybase sales strategies was to get into a benchmark with Oracle on transaction processing. Because we told people the amount of hardware you would save by buying Sybase would pay for the Sybase license by not buying Oracle; they would actually get a free license.

Jerger: And was that the case?

Schuster: That was more than the case.

Rowe: So the other thing on the benchmark wars that I found interesting was that in the mid-1980s we'd all gone overseas and we were selling in Europe and other places. And the thing that surprised me is if you went to Europe and you talked about your product, they didn't care about performance. Benchmarks didn't make a bit of difference to them. All they cared about was program or productivity tools, whereas in North America they couldn't care less about program or productivity tools. All they cared about was raw benchmark speed. And I didn't completely understand that until I went over there one time and realized that at least in England, they don't work very much. And the program managers were mostly concerned with, "Can we get our programs written?" And they didn't care how fast the database systems ran. You can always buy more hardware.

Schuster: Well, it was one of the reasons that Ingres did well internationally in those early days. And we struggled because we did not have a great tool set internationally. The markets we were good at were the ones where applications came out of the U.S. and migrated over to some of those areas. But it was interesting to see the switch in the middle or late 1980s where the client server message became the dominant message. Because people were getting tired of every release being the fastest and you could only play that card so many times to where the applications didn't matter. We were running faster than most people needed at that point. Then the whole architecture wars started regarding client server and by the time we went public, the client server message in 1991 was the dominant message for Sybase.

Sybase Going Public

Rowe: You went public in 1991?

Schuster: Yes. A lot of people think that we raised a ton of money. But Sybase was a $100-million company when we went public. So instead of raising money publicly, we raised it privately.

Sippl: Well, you sounded like you grew four-fold one year. That's a hell of growth.

Schuster: Yes. We went from $6 million to $24 million.

Sippl: $6 million to $24 million and then $24 million to $80 million.

Hoffman: To $87 million.

Bohl: You had such a foothold with Sun and customers in Wall Street.

Schuster: And that spread internationally extremely fast that way.

Sippl: Well, you were running on two platforms. So if someone had to buy Sun in order to run Sybase, I can see why Sun people were pushing it. I couldn't get Sun to push our products at all.

Hoffman: But we were a server-based solution.

Sippl: Yes, you were server and network oriented.

Jerger: When did you get to the $100-million mark? Was that 1989 or 1990?

Hoffman: It would have been 1990.

Jerger: Marilyn, when did you come over from IBM? You had been at IBM for 19 years or so, right?

Ingres Release 6

Bohl: Yes, I was in IBM a long time and then I actually interviewed at Ingres. I decided I was going to leave IBM, and Ingres at that point was running on the mainframe or trying to run on the mainframe in1987 or thereabouts. I met with Gary Morgenthaler. But then by 1989 the Gateways strategy had come along and Release 6 was the SQL implementation. But my perception is the extent of the changes in the code versus the extent of the testing was behind. And the code hadn't been officially released yet but was out in about 50 sites before it was released. Greg, is that your perception?

Batti: We had a re-architecture of R6, we had the distributed database stuff, we had the Gateways, we had performance work and by the time you're talking about we'd started on the advanced database stuff. We had the object or relational stuff and all the other things were going on. These were all going on in that same two-year period.

Hoffman: And Release 6 wasn't just a lot of code changes; it was the whole system.

Batti: It was a whole new transaction system; it was a whole new processing system.

Bohl: But the expectation of customers was not set. So customers would take applications that had performed fine over to Release 6 and they didn't perform at all.

Rowe: One of the reasons, I think, is that people had a bunch of VAX's that didn't have very large memories and the old version 5 was a real nice little system. And now they had this new system that was big and had lots of wonderful features, but it took more memory. If you didn't upgrade the memory or you didn't upgrade your processor, your applications were going to slow down and guess what, they did. Customers weren't happy.

Sippl: You tried to keep both API's at the same time, though, didn't you?

Batti: Yes.

Sippl: You came up with a new SQL release but it would still run the old embedded code.

Rowe: We had to do that because we had a large installed base. But the harder problem and the thing that we're not talking about is Bob Coye. Ingres had a good query optimizer in the university days. System R had a good query optimizer in the System R days. And, frankly, it had two variable tactics that Ingres didn't have. We hired Bob Coye very early on, about 1982.

Batti: He started about a month before I did in early 1982, like March 1982.

Rowe: Bob Coye got his PhD from Case Western University. His PhD thesis was to redo the query optimizer in university Ingres to implement a whole bunch of additional tactics including the two variable tactics that Pat Selinger had done at IBM. And bless his little heart, he ran it on a PDP-11. On the old PDP-11, the system had five processes and he added two more because to get the address space you had to have all these damn processes. It made the thing run great. So he came and gave a talk at Berkeley and as we walked out of the talk, I grabbed Mike and I said, "We have to hire him." So we hired Bob. He came to work at Ingres and rewrote the query optimizer. And he also redid the sort package – or was it Derek who redid the sort package with him? And between those two things, our performance was absolutely fabulous. We were really a top performer.

Batti: Well, that's where we got the DeWitt benefits.

Rowe: That's right. But then what happened was we went to implement SQL. And the issue was how do you implement SQL? Do you implement it so that it works identically in terms of inputs and outputs as to what IBM does? Or do you implement it so that it's plausibly SQL when you type the stuff in? And being a technology company we couldn't resist saying, "We're going to make it work exactly the same." And the query optimizer got broken. It got broken very badly. And the complexity of getting that optimizer to work was going to be very, very difficult because Bob had optimizations built in there based on the semantics of QUEL that were then no longer true for SQL. It mostly had to do with when you did sorts and duplicate removals and stuff like that. As a consequence, the performance really got into trouble. Then Larry Ellison came knocking and hired Bob away so we didn't have the world's greatest query optimizer guy working on our code anymore. And I'm not sure the Ingres code ever recovered from Bob leaving. Bob went to Oracle; he worked on Version 9, I think it was.

Batti: 7.

Rowe: 7? And all of a sudden Oracle, who had never been competitive performance wise, suddenly got very competitive.

Butterworth: The SQL stuff was interesting because we actually did it in two phases. We did the emulation first and then did the source transform.

Selling Ingres to ASK

Jerger: So when did ASK come along?

Rowe: ASK came in 1990. There were a couple of things that happened. Sutter Hill was our lead investor. We had a secondary investor, Welsh Carson. We'd done three more rounds because we just couldn't stop spending cash. And then we were queued up to go public in 1987. But the stock market vaulted down so we had to delay. We delayed for six months until 1988. I guess it was May, 1988.

Johnson: I have the prospectus in front of me. It says March 17, 1988.

Rowe: Right. The prospectus was actually written the previous August, as a matter of fact.

Butterworth: I actually have a copy of the October 1987 prospectus. I think we were going to go out the Monday after the crash.

Batti: Yes, it was.

Rowe: Well, I know we were at that board meeting because I'd flown back from a conference back east at that time.

Johnson: I also have some handwritten notes that Gary Morgenthaler made at that board meeting.

Rowe: What happened was that Sutter Hill's investment came from a hospital group or an insurance group up in Canada. They owned, I think a million or two million dollars worth of stock. They sold that business to a Scandinavian company and the Scandinavian company took a look and said, "What's this? What's this funky little technology company that's kind of out there doing whatever?" And they wanted to get rid of that stock. So they went and whispered to Wall Street, "Here's two million dollars or a million dollars worth of stock. Please sell this if you can." Needless to say, Ingres' stock could not move because if it went up by this amount, all of a sudden this million dollars worth of stock was going to land. We were really thinly traded, maybe twenty, thirty thousand shares a day. And our price was, at the time, maybe $10, $12. We couldn't get that overhang out of market. So that was problem number one.

Problem number two was the venture capitalists that distributed the stock no longer wanted to hang on unless there was a great advantage in the future. And the problem we had from the day we got started was how are you going to beat Oracle. And you could never come up with a strategy to beat Oracle because whatever you said today, two days later Ellison was saying it with more marketing dollars and your message was blunted. So you could just never do it. So the VCs decided that they wanted out. They started looking for somebody to buy us. Sandy Kurtzig had come back into ASK and then decided that the OEM hardware business was falling

apart. Her business was to sell so-called manufacturing software. But her average sale was $200,000. And $100,000 of it was OEM hardware from HP, $50,000 was for her software and $50,000 was for her consulting. The OEM hardware business went away and her consulting and software business was being invaded by Oracle so she was in trouble. She was deciding that they had to do something to compete with Oracle. Oracle had a database plus an application. She had no database. Don Valentine, I think, up in Sequoia who was her advisor, basically said, "You should buy these guys." We were working in the back room on a new set of graphical user interface tools.

Batti: Both the new database and 4GL stuff were just out.

Rowe: And she had seen that and she went ga-ga. So she decided that she needed a database to compete with Oracle and the venture capitalists didn't want to hang on anymore so they were basically willing to sell for any price. She had to come up with the financing. The good thing the VCs did was they said, "It's got to be an all-cash deal." So she had to find the financing. One way she found the financing was to go to DEC and said, "DEC, how about investing $20 million so that I can buy this company?" And the DEC guys said, "Well, yes, okay." But then they looked at each other and said, "You know, we don't need her. We ought to buy Ingres ourselves."

So we had a board meeting at the law firm in the city on the day of the regional NCAA playoffs. I had come back over to that game with Marilyn and Paul. And it was one of these really curious things that Sandra was brought to floor number 13 and she made her pitch. Simultaneously sitting upstairs on floor 18 were the DEC guys and then they came down and said, "Well, if you want to do a deal with her, you can do a deal with her. But, in fact, we would prefer to buy you and we would like to do a deal with you. We'll make a minority investment, take 20 percent of the company." And we said, "Okay, is this serious?" And they said, "Oh, yes. We have a board meeting tomorrow and it'll be raised at the board meeting and so we'll find out." And so sure enough they came back and said they were going to do the investment.

Well, as it turned out, we had some squirrelly finances because of some bad deals in Europe. You probably know about these, Stu. And we had an employee over there who wrote bad deals and so the books got squirrelly. The technology was okay but the books scared the financial guys at DEC so they backed out of the deal. So in the end we really had no choice: either stay separate or sell out to ASK. And the VCs absolutely wanted to get out from underneath it.

Hoffman: What year is this now?

Rowe: This is 1990. The board meeting where we had the two groups probably was done in March of 1990. And then the final deal with ASK probably was cut the first week in September. I think the actual board meeting where the deal was agreed to was on the 12th of

September and I think the change in control happened in October or something like that. And so she bought the company.

She thought this was going to be wonderful and everything was going to go forward. But her manufacturing software business absolutely collapsed and the thing that I found interesting was the year before the deal was done, we both did about $150 million. So the expectation was the combined company would be $300 million and the next year we were expected to do $400 million, something like that. The next year, I think they did somewhere around $380 to $400 million. And 90 or 95 percent of it was Ingres business. So it turned out ASK would have been flat dead in the water if they hadn't bought the company. The deal was done at $8.25 a share and within a year the stock was trading at $15 a share and it was all based on the database business; it had nothing to do with the manufacturing software.

Jerger: So what happened to you folks, then? I know that from something I've read that Marilyn stayed on with ASK for awhile.

Bohl: Yes.

Jerger: Did you guys all just disappear then from ASK?

Rowe: Paul had left earlier.

Butterworth: It's a long story, but I actually quit about two months before this happened. But there was a bunch of things happening and we thought we'd all stay around until ASK took over. I actually left the day after the acquisition.

Weiss: I was gone before that, in 1989.

Batti: I stayed through another couple of years. I decided to leave about three or four months before the ASK thing happened and I left about two or three weeks before it got announced.

Hoffman: I left before the CA acquisition.

Bohl: I left before that. Just before the CA acquisition, too.

Jerger: When was the CA acquisition?

Batti: That was 1994.

Jerger: How about you, Larry? What did you do in the middle of that?

Rowe: Mike and I were on the board and when the takeover happened, Kurtzig's view of us was, "Oh, you're the professors from Berkeley. This is so wonderful. You're going to come teach at ASK University. You'll be so excited to teach at ASK University."

Jerger: Like Hamburger U?

Rowe: Yes. And I just kind of went, "Oh, my God. This isn't going to work." As it turned out, I had not sold any stock in the previous six months. It was an all-cash deal so I could actually sell all my stock. All my stock got sold in the deal. Mike had sold stock three or five months earlier. In the previous window he had sold some stock. So it turned out he couldn't get out right at that time.

At the change of control, Kurtzig didn't make an offer to me of any kind. So I said, "That's it, I'm out of here." Mike had to stick around until April, I think, so that he could then sell because he was locked up with the six-month sell-by restriction with the market. I think he actually stuck around a little bit longer than that as I recall.

Batti: It wasn't much longer, I don't think.

Rowe: It wasn't much but then he bailed out as well.

Jerger: We have about six minutes to wrap up. But I have a question, did you two folks have fun working with ASK or did your style change?

Batti: The only the fun part of it was we had just brought two good products to market right before the acquisition. And we brought those out and they were actually successful for awhile. So that was the good part. But the problems that we had as a company were not solved by being part of ASK.

Batti: We went through two or three presidents under that regime. And that was difficult.

Rowe: But at that time Sybase was really on the rise. Ingres was in a lot of trouble.

Sybase Growth after the IPO

Jerger: Let's let Sybase have the last five minutes here of the glory and the fireworks. So you went public in August of 1991. What was your sales volume right before you went public?

Hoffman: That was just a little over a $100 million.

Batti: You were about the same size as us at that point but then you took off from there.

Schuster: In fact, we were actually growing faster than we had for a little bit. We were actually growing faster than a hundred percent per year.

Hoffman: Yes, we wanted to go public in 1990 so instead of going from $56 million to $120 million, we went to $103 million or whatever it was. So we went back, we reorganized, we changed a lot of the management and the structure at that point. We reenergized the company and then by the end of that year we were healthy again. And so then we did the IPO.

Schuster: And I think we did $150 million and then $250 million or something like that. It was really on the upswing.

Hoffman: Yes.

Jerger: Good timing. So then you were there, too. What about 1996?

Hoffman: I was there until 1996.

Jerger: How about you, Stu?

Schuster: 1995.

Hoffman: Yes, I was actually there until 1995. I left in 1995, too. And I was just on the board.

Jerger: Why did you leave?

Schuster: I had nine years of marketing. I outlived seven V.P.s of marketing at Oracle. I decided that the ultimate goal of this industry was getting the V.P. of marketing fired from the

other company because then you knew you were doing your job. Marketing is one of those areas where you get a lot of blame and very little credit. So I figured, "Well, okay, let's see how long I could last." But I was burning out at that point. The company became very large. We had done a lot of acquisitions. I was having personal problems at home. My wife got sick. And I was really thinking that it was time to go do something else. We had just done a $250-million quarter when I said it was time to do something else. We were on a billion dollar run rate and I was just tired, just tired after nine years from zero to a billion dollar run rate. And so for six months I kind of just hung around advising in a transition, doing some business development work and then left.

And it was really strange because the people that funded Sybase were a late entrant, Brentwood Venture Capital. They came in as a mezzanine round and they asked me to look at a couple of companies. So I said, "Well, tell me about venture capital." And they said, "Well, we're doing a new fund and we want to do some software deals. Do you want to come in as a venture partner?" So that's how my venture capital career started. And I thought all the good opportunities were over in 1995. But then this thing called the Internet came along and I was actually very fortunate, because I was a seed investor in Commerce One before Mark joined. I actually tried to talk him out of joining the company saying that there were a lot of problems there.

As a little side story, after working nine years at Sybase, at this little company, putting down a seed investment before Mark ever came in, I made the same amount of money. And then he joined and turned that investment into the same profit that I worked nine years for at Sybase, surviving against Oracle, Ingres, Informix, RDB, you name it. And here was this little bet that I had written off and then Mark came into the company and completely changed it. And we ended up working together a little while in the end.

Rowe: I made a seed investment in Ikomai and I made twice as much money in Ikomai as I did at Ingres where I had 7% equity the day we wrote the VC deal.

Jerger: We're kind of at the end of our time. Is there anything else that anybody would like to mention?

PowerBuilder and Ingres

Sippl: Well, the PowerBuilder deal, how did that go?

Schuster: That's very interesting because that was very important. That was the largest software merger of its time. It was a $900-million merger, wasn't it?

Hoffman: Yes. But we were struggling. We had a front end. And part of the problem that we ran into starting in 1994 was that we didn't fit into the new applications that most of the people that were building stuff around us were building. And this all had to go back to the row-level locking stuff versus page-level locking. And as the market began to move to applications people, our business started to slow down. It dropped from 30, 40 percent growth down to 10 percent growth starting in 1994. So we felt we had to get an application development tool that was really good, that would allow us to switch over and go to the application side. At the same time, this is not generally known, but we were also talking to Hasso Plattner about him embedding us in SAP. He'd been using a lot of Oracle so he was building up Larry by reselling the database even though he was an enemy. So he was talking to us about doing this. What we found was that he had written all the SAP applications around row-level locking. We were supporting page-level locking. We couldn't get the performance required unless he rewrote some of his applications which he was not willing to do. And so that deal fell apart at the very last minute which was too bad because that would have changed the world.

Rowe: So, Roger, you want to also tell them about the two deals we didn't do. Where you and I tried twice to put Informix and Ingres together?

Hoffman: Let me finish the story about Power stuff just for a second. That, combined with the market shifting to applications and people just building less of their own stuff got us to go with Powersoft and get a tool we could use to build the application quickly. But the merger was huge at the time. It was, like, a billion dollars, the largest single deal done. But it was also a tough deal. And integrating East Coast, West Coast personalities was difficult.

Schuster: Their sales model was totally different from ours. We were direct sale and they were more of an evangelism telesales operation. It was very interesting culturally merging with them.

Hoffman: It was really hard.

Hoffman: They had a phenomenal hidden jewel, though, which was their database.

Schuster: SQL Anywhere.

Hoffman: Yes. That turned out to be a great thing actually for Sybase. Those guys were super, super good guys and did a great job. But the application stuff never really worked together so we were really struggling after that.

Sippl: So then what happened?

Hoffman: Well, that was 1994. I left in 1995. Mitchell Kurtzman came-- took over for a couple of years.

Sippl: And did revenue fall?

Hoffman: Fall, yes. It just drifted down and then a couple of years later John Chen took over and John's been running Sybase ever since.

Schuster: Yes, it peaked at a billion dollar run rate and I think it's still over $800 million and it's still, outside of Oracle, the only other independent big database company out there.

Hoffman: And it's cash flow positive.

Rowe: Do you have any idea what percentage of their sales are database oriented at this point?

Hoffman: I think it still is the biggest piece of their revenue. It's between the Sybase stuff and SQL Anywhere. SQL Anywhere, because it has a very small footprint, has really taken off in the mobile side of the business.

Sippl: And that's the database jewel that came from Powersoft?

Schuster: Yes.

Jerger: Thank you very much. It's been fun.