

Profiting From Free Software. The Sleepycat & Cloudera Stories – with Mike Olson

Three messages before we get started. First, who's the lawyer that tech entrepreneurs trust? Scott Edward Walker of Walker Corporate Law. Here's what Neil Patel, founder of KISSmetrics, says about him. "Scott is a great lawyer. He is affordable, responds fast, doesn't charge you for five minute phone calls, and always gives great advice." Walker Corporate Law.

Are you still coordinating projects by e-mail and wondering why you're not productive? Would you check out TeamworkPM.net? TeamworkPM.net is trusted by Universal Studios, the University of Texas at Austin, the U.S. Army, and other large organizations. But if you're just trying to coordinate with a small company, maybe just two people, it will still help make your team more productive. How much more productive would you be if you used TeamworkPM.net?

Finally, do you know someone who wants to build an online store? Tell them that Shopify is the easiest way to build a robust store. Then you can show them who else created stores on Shopify, including Tesla Motors, Pixar, Amnesty International, Foo Fighters, and others. Shopify stores look great, are dependable, and they generate sales. Shopify.com.

Here's your program.

Andrew: Hey everyone. My name is Andrew Warner. I'm the founder of Mixergy.com, home of the ambitious upstart and the place where you come to listen to entrepreneurs talk about how they built their businesses, what they learned along the way, and mostly to pick up a lot of information and ideas that you can use in your business so you can develop it and come back here to Mixergy and do interview of your own, just as today's guest is.

Joining me today is Mike Olson. In 2006, he sold his previous company, Sleepy Cat, an open source, embedded database company which he sold to Oracle. Today he's running Cloudera, a company that contributes to and distributes Hadoop, the open source software that powers the data processing engines of the world's largest and most popular websites. Mike, welcome.

Mike: Thanks, Andrew. It's great to be here.

Andrew: Mike, I read an old article in which you were asked what you were going to do...this was right after you sold to Oracle and left the company and you were asked, "What are you going to do next?" You said, "I'm going

to take six months to myself.” I’m wondering in that time that you took to yourself, what did you do?

Mike: The first three months I spent volunteering in my kid’s school and working on my extreme skiing skills and getting the mountain bike out of storage and generally enjoying myself. That, to my surprise, and actually to my wife’s real disappointment, got pretty boring faster than I thought it was going to. We wound up, I wound up looking for other real business opportunities beginning in the summertime. I hadn’t expected to jump into that until late fall, early winter. I’m glad I did when I did. I certainly think we found a great opportunity.

Andrew: In your toughest days at Sleepy Cat, and I’d like to find out about them throughout this interview, but when it was tough and you said to yourself, I can’t wait for this to be over so I can... did you get to do what it was that you were waiting for? Whatever would’ve completed that so I can sentence?

Mike: I don’t know if this sounds self-serving, but look I never wanted to finish Sleepy Cat so I could not do it anymore. There were, of course, tough days. Any company, any interesting venture is demanding in lots of ways. I loved my job and I loved the people that I worked with and I for sure had a trajectory. I did, in fact, want to find a way to launch the company into a larger organization. I wanted to find an exit that made us a part of a very big company that was larger, and frankly more capable, than our little open source company was. I think we did pretty well on that. I enjoyed the two years that I spent at Oracle post-acquisition. I learned a great deal. I met a bunch of really smart people. I think we left the software, the employees, and the customers all in a pretty good place.

I did know that my next company was not going to be a database company. I’d been doing that since 1986, 2008 that’s plenty of time. I was glad to move on to something different. If you look at Hadoop, however, you’ll see that it’s as close as you could possibly get to a database without actually being a relational database. I don’t think I strayed too far from the field.

Andrew: Why would you want to stray at all? One of the things that I saw in a talk that you gave at Berkeley was that you’re proud of knowing as much as there is to know about databases or as much as you possibly could know. Here you were, you knew it all, why’d you want to stray even a little bit?

Mike: A couple of reasons. The first is that the classic relational database space, I think it is a very valuable technology. It's enabled businesses to do lots of really amazing things over the past 30 years, but it hasn't been that interesting a place to work for the last 20 years. That is to say much of the really interesting innovation that got done in the industry was happening outside of academic and commercial database research. Yeah, we did a good job making those platforms run 2% or 3% faster a year, to scale up 2% or 3% a year and you aggregate that over time and that's pretty fast and that's pretty capable. But nothing really dramatic, new, exciting was happening in the RDBMS space in that time. Frankly, the other problem with being a database entrepreneur is that you're active in a market with well-capitalized, capable competitors. IBM and Oracle are formidable. I spent my entire career competing with Larry Ellison until I finally sold the company to the guy, and I just didn't want to go do that again.

Andrew: What did you learn from working for him?

Mike: I was most interested in seeing how Oracle behaved in the field. How they created such a remarkably efficient and capable sales engine and managed their revenues. I had a great chance to do that. I got to fly all over the world and work with the Oracle field in different places, understand how they approach their markets, how they structured their sales team, what kind of people they hired, and what kind of compensation they paid. I learned a great deal about that stuff.

Andrew: Do you have an example of one takeaway, one specific example of what you've learned?

Mike: For example, Oracle has a very valuable, big and well staffed business in Asia. You can't look at that part of the world as monolithic. The way that you engage and behave in China is very, very different from the way that you engage and behave in Japan. Hey, by the way Korea and Singapore are different from both of those geos. The kind of people you put on the ground, the pricing you're able to extract from customers, the way that you connect with them...

Andrew: How does Oracle do that well? How do they from over here know who to put over there and how to manage it?

Mike: They've done what any really successful large enterprise software company does over time. They've built out remarkably capable leadership regionally. They've got a great head of Asian operations, but then that person relies on the leadership in Japan and in Korea and in Singapore and in China to manage the businesses in those countries. They've created a very well-structured sales organization that maps nicely to how buyers behave there. Hey, by the way, what you do in Europe is different still, and the way that your North American organization is structured is different yet again. So they built a global organization with regional specialization that's appropriate to the areas where they do business.

Andrew: Okay. All right. So just before Sleepy Cat, you worked at Molecular Applications Group, and I heard you say that that was... did you say it was a mistake? Your sly to Berkeley said that it cratered, but you seemed to say that there was... well, you describe it, how was it?

Mike: I would call it my one year Mulligan.

Andrew: One year Mulligan.

Mike: Yeah, I went there to run the engineering group for an interesting bio-informatic startup. One of the key things I learned is that if you're going to work in bio-informatics, it's pretty good to have a background in bio-informatics. I didn't understand the customers. I didn't understand the applications.

Andrew: What is bio-informatics?

Mike: The use of computer technology to answer biological and medical questions. That is, if we wanted to design a new molecule to treat a disease that's particularly pesky, some variant of cancer, for example, what information systems can we use? What software, what drug design technology can we use that compliments the wet-bench science that we need to do? I knew IT very, very well. I didn't know biology at all, and I didn't know wet-bench lab science at all. I think that I didn't do the job that I needed to in that job as a result.

Andrew: Okay. At the end of that, how does it shape your sense of who you are? I feel like I, as a man, as a entrepreneur, as a businessman, the way

I see myself is based on how my work goes. At the end of that, how did you feel?

Mike: I was glad to not be responsible for the team and the product anymore, because I don't think I was doing the job that I needed to. I didn't regret the time that I spent there because that's an important thing to learn about yourself and about what you should look for in a job at which you can be successful. I will say, and this is vain, I'm actually proud of the way that I disconnected from that organization. That is I didn't just pack up and leave. I worked with the investors and with senior leadership to be sure that there was a transition plan, that I could get out cleanly and that someone else could come in and take over my responsibilities. Now at the end of the day, MAG didn't succeed. That's not that uncommon in the startup, but I don't feel that I behaved badly on my exit and that was very important to me.

Andrew: You didn't feel that maybe this wasn't for you. Maybe you needed to go to a more established companies. Maybe go knock on the door of Oracle and start working for them or IBM. You said, "I'm going to go and take a shot at a new company." Sleepy Cat was very young. There wasn't that much there. How do you get yourself comfortable with taking that leap? I feel like once you get burned you don't want to touch the stove again.

Mike: It was a personal decision. The founders of Sleepy Cat, Keith Bostic and Margo Seltzer, were both very good personal friends of mine. We had been in grad school together, and in fact I had been one of the early authors of Berkeley DB, the software that we distributed, back in the day when I was still programming for a living. I knew them and I knew the technology well. I was looking, after MAG, for something I was comfortable with. So, for sure, working with friends on database technology that I had written before, I kind of had that job wired. I knew what I was going to be doing for a living and getting together with those guys was pretty easy especially after what had been kind of a difficult and painful experience. One that was a little bit more comfortable was really attractive.

Andrew: All right. For someone like the people in my audience, and frankly for me, who doesn't understand what Berkeley DB is, can you help us understand what it is?

Mike: Yeah, it's an embeddable storage engine. There are lots of applications where saving your data reliably makes a big difference. You can think of stuff like your e-mail program or maybe the password and directory system that runs on the systems you log into. Many of those actually use Berkeley DB as a embedded store. It's not really a hardcore relational database problem. You don't ever want to type a sequel query at your e-mail database. But you need to be sure that when a message gets delivered, it doesn't get lost. That software is actually, Berkeley DB specifically, is very broadly used in the industry. It's pretty hard to touch the Internet without running into Berkeley DB in name service or in website caching or in other sort of critical services. Back in the day, it was a part of products that were distributed by vendors like Cisco and EMC and others. Amazon relied on it considerably for managing data on their website. It was very widely deployed and quite widely used even though I bet most of your users aren't aware of the number of times that they touch it in a given day. It's still absolutely ubiquitous on the Web.

Andrew: The embeddable part means that it's part of the website, part of the e-mail system?

Mike: Disappears in the applications. You don't install an IBM DB2 database and then an application that runs on it. You get an app from Sun Microsystems or you connect to the Amazon web service and magically, silently without any action required on your part, you're getting good, fast, reliable, scalable storage services.

Andrew: I see. What was Sleepy Cat's role in distributing and profiting from it?

Mike: We were the company that wrote the software, that is everybody who was a Berkeley DB contributor was on payroll at Sleepy Ca. We used a technique called dual licensing. If you were an open source user, you could use the software for free and bundle it with your application and ship it anywhere you wanted to. If you didn't want to give your own source code away, you could buy a license from Sleepy Cat, the company, and we'd write you a different license for the software. We were the developers, the maintainers, the supporters, and the licensers for proprietary users of the software. It was a nice model in a very profitable business for the eight years that we ran it independently.

Andrew: If I understand this right, you were one of the innovators of the model itself, right? Today it's much more common to have a licensing system like that, isn't it?

Mike: Yeah. What I'll take credit for is being at Sleepy Cat early and making the model work. Credit where it's due, Keith Bostic, who was one of the founders of Sleepy Cat, a guy I knew when I was at grad school in Berkeley, he actually invented that open source dual licensing model, the different open source and proprietary terms. It was Keith who figured out that you could do that and created that model before I got to Sleepy Cat. The company was tiny when I joined. It had just a handful of customers and had not yet broken through a million in revenues, but that did happen soon afterward.

Andrew: As I understand it, you were the first employee there, and you were brought in to help sell it or to help introduce it to new people, right?

Mike: Yeah, that's right. Keith and Margo both had day jobs. Margo was a Harvard professor and Keith was working at a company called BSDI, commercializing the BSD operating system. They were running Sleepy Cat nights and weekends in their spare time out of a web server in their living room. I was the first full-time employee for the company when I joined in 1998.

Andrew: I see. And once you joined, what were the first things that you did.

Mike: First thing I did was convince Keith to quit his job at BSDI and start working full-time Sleepy Cat because we had enough customers and we needed the help. We worked on packaging and on the website and on marketing and basically on pursuing customers and closing deals. It was bottoms-up, grassroots. We never had any investment in the business at all. We ran on revenues, profitably, since we opened the door for the very first time. You can imagine a lot of my time was spent on sales and marketing activity.

Andrew: I see. How did you start? I don't know so much about the technology behind the business, but I do know more and I am much more fascinated about sales and marketing of an idea and a business. Can you tell me how you did it in the early days?

Mike: Yeah, you know, open source software has a huge unfair over proprietary solutions. It's everywhere in the environment. That is people can come and get it, taste it and try it out without contacting you. That means that by the time they've decided it's good for their problem, they've qualified themselves very, very well. You don't have to invest nearly as much in outbound demand gen style marketing. Now, when we got into substantial millions in sales, of course we started investing in that stuff. Early on, we took advantage of the inbound interest in the platform and worked hard to develop those opportunities into bigger ones. So, somebody would come to us. They like Berkeley DB. We help them understand how to make more use of it to deploy it in many more products than they had originally thought that they would do and build the value of those deals and those relationships over time.

Andrew: I see. And did you have to get, or was it part of your job to get more people to try the free version, to sample and then later on become customers? Or was that just working on its on?

Mike: We promoted open source Berkeley DB all that we could. The fact that it was open source and that it was, by the time we even started the business, well respected and well known, meant that we didn't have that much upfront investment to make there. We did, whenever we could, go to conferences, show up at user groups, take out magazine ads, write articles for the technical press, encouraging people to take a look at it. We engaged in that kind of demand generation, but it really was aimed more at getting people to try the open source software than hey, come and give your name because I want to call you on a sales opportunity. People will generally qualify themselves and then come to us when they're ready.

Andrew: How do they know who to come to, and how did they move themselves from one to the other, from the free to the paid?

Mike: In general, I think we did a good job of branding Sleepy Cat and Berkeley DB. If you knew what Berkeley DB was, you certainly knew where it came from. Our website was where you went for fixes and patches and the latest version, and we maintained mailing lists for people who were running it in deployment and so on. We were pretty visible there.

Andrew: I mentioned earlier that there were setbacks along the way. Can you describe one of them?

Mike: Setbacks at Sleepy Cat?

Andrew: Yes.

Mike: I would say that the challenges we had were, we decided that we didn't want to take investment capital. We were running the company profitably. That meant it was often difficult to grow as quickly as a venture backed startup might have. Hiring people in advance of revenue is impossible if you don't have a deep pile of cash that you're able to draw on. We were forced to be a lot more deliberate in growing the team. I think we might have been able to grow the business bigger and faster had we been willing to invest forward. I don't regret the model at all. I think we had a tremendous exit. I think that we built a very valuable and interesting business. But it was harder to bring people on in advance of demand. The other challenge was as the business neared the point that we wanted to sell it...

Andrew: The other challenge was what? I apologize, I didn't catch that.

Mike: As the business neared a point where we were ready to sell it [interference] that we thought an acquirer would be interested in, we struggled. We had been very much in control of our destiny for a long time. We were running a business that was throwing off a lot of cash. It was a great... people say lifestyle business like it's a bad thing. It was providing a wonderful life for all of us involved. We struggled a little bit with how and when we should sell the business and take the profits and let someone else grow it through the next stage.

Andrew: How did you decide to sell it then and why?

Mike: [inaudible 00:20:12] databases was clearly bigger than we were going to be able to take advantage of just by ourselves, just on our own. Establishing a field sales operation in Europe and in Asia and in the U.S. that was as various and capable as the Oracle orb that I just described was way beyond any investment that we could possibly make. We didn't think that we could live up to the opportunity in the market. Frankly, each of us was interested in making some life changing money, getting to the point where we were able to go do whatever we wanted with our next venture. We got to the point that we thought that Sleepy Cat would provide that kind of exit for us. It had been a great run. There were reasons to be concerned about our ability to grow it on our own anymore. The founders, the major shareholders and

the board sat down and hashed it out over the course of many months and came to the conclusion that the time was right and that we should seek a buyer.

Andrew: I want to dig into some of what you said so far. The first thing I noticed was the funding. Why didn't you raise more money earlier on then, if it would've helped?

Mike: When Sleepy Cat was young and pursuing a open source dual licensing model, we got laughed out of the room by the venture capitalists that we talked. It's just, you know, you're going to make a living selling free software, right. That can't possibly work.

By the time we had established the model and revenues were good, we were making a wonderful profit. We were taking that profit and paying big bonuses to all over employees and distributing the remainder to the shareholders every year. Keeping enough in the company to operate it for the beginning of the following year but it was a very cash pleasant lifestyle. If we had raised money at that point, we would've surrendered our control over the business. We no longer would've been able to pull the cash out. There was simply no compelling reason to do it at that point.

When we decided that we needed to grow, we needed to invest in sales presence overseas, for example. We actually considered raising venture capital. What we concluded was that a strategic acquirer, say Oracle, and a VC would treat us exactly the same. There would be no more opportunity to pull the profits out. We'd lose critical control of board level decisions because the investors would be demand, reasonably, that they be consulted on stuff. Given that we were going to lose that kind of control. That our big cash upside was going away, we decided it made more sense to take all the risk off the table. Find a strategic acquirer rather than sell a fraction of the business, sell the entire business.

Andrew: You also said, several times, that the business made a lot of money. What size profits were you guys generating?

Mike: They were very pleasant and we never disclosed any of that publicly.

Andrew: Can you talk about just the neighborhood that they were in without being specific enough at all to compromise the privacy that you guys have enjoyed?

Mike: We were generating millions a year in net profits.

Andrew: Okay. And you said tremendous exit, what size exit?

Mike: Also, that's never been disclosed. That's Oracle decision not to publicize the purchase price.

Andrew: Why would Oracle not want to publicize it, I could understand why you guys wouldn't because it reveals too much about your personal situation. Why does Oracle and why do acquirers not like to discuss how much they paid for a business?

Mike: You'd have to ask Oracle it's motivations. One thing that I'll suggest as a possible motivation for acquirers generally is, look if all of your deals are priced in the open, then when you go to make another acquisition, you know that you kind of set up a core for the price that you'll pay. So, keep the prices that you pay secret allows you to negotiate a little more effectively with subsequent acquisitions that you might want to make.

Andrew: Finally, I wrote down what you said earlier, you wanted life changing money. How is life going to change for you personally?

Mike: The Oracle exit, the sale of Sleepy Cat to Oracle and then one previous transaction that had happened when I was working at Illustra Information Technology, we got blocked by [inaudible 24:14].

Those had really put me and my family in a position where I've got the latitude to do what I want to do. I can choose, for example, as I did that six or eight months that I took off just to not draw salary at all, not to feel beholden to anybody. I can afford to invest in opportunities directly that I think makes sense. Just key obligations that you have as a father, my kids can go to schools that they want to go to without a lot of worry. I don't need to worry about day-to-day, month to month, even year-to-year being able to provide. We've got enough of a nest egg that we can fall back on that. That allows me to make what look like risky bets. When we started talking about Cloudera, for example, it looked like a pretty risky bet. A bunch of the VCs that I approached early thought that it was just a crazy idea and that it couldn't possibly work. I was advised a number of times to walk away from this thing because it couldn't possibly work.

Andrew: Something that strikes me every time and I haven't seen you much but I got to see some videos online in preparation for this interview, got to read about you. You seem just very comfortable and unworried. Is that true, what do you worry about today, professionally?

Mike: Am I unworried, no.

Andrew: You know what I think part of it is, you have an earring. I can't so much see it here in this Skype video but you wore an earring at Berkley. You look like you're hiking in a few of the pictures that some of the computer online magazines show of you. I felt that you feel like you don't have any worries, are there any now?

Mike: You ought to try to live a life that lets you lay the stress off. I will say that. If you want to be a successful entrepreneur, find ways to enjoy your family, eat well, get some exercise, and leave the office at the office sometimes. Am I worried, of course I am, I'm the CEO of a very fast moving, really exciting company that we think is going to transform the way that data gets managed across the industry. We have an enormous opportunity. We've had a tremendous run for a couple years and, frankly, the world and we expect a great deal out of Cloudera.

Andrew: When you worry personally, what is it about? Is that you'll embarrass yourself in public with everyone expecting so much from you? Is the venture capitalist that looked at Cloudera and said this isn't going to work, that they might be proven right? Is it something else? Is that you'll have to spend another six months driving your wife crazy?

Mike: No, I'm not worried about what others think of what we're doing. What am I'm worried about is inside the company, we all believe deeply that we've got tremendous opportunity here. But a lot of stuff can go wrong, we've got big customers running the software and mission critical production use. We've got to live up to their expectations as service providers and as a support organization. Our product has to be wonderful for them. That's hard, hard work. We've got a lot of opportunity to sell the software to new places. Well, you know, we probably don't have as big a sales force as we need right now and we've got to solve that problem. We're only in North America, hey, what about Asia, what about Europe, got to figure out how to get there. It's what an entrepreneur calls execution risk. We've got plenty of wind at our backs. If we mess up, we could still find a way to fail. I do

worry that we'll mess up and find that way to fail. Not that I expect it but I think if you don't worry about that, then for sure it's going to happen.

Andrew: I see, okay. What was the original idea that made you see, I've got launch this business. Cloudera is my company.

Mike: Hadoop had been developed by the big web properties to solve a data management problem that they had ahead of the industry. They were collecting complex data from lots of places and combining it with sophisticated analytics, ten years ago. I read the very first map produced paper that came out of Google in 2004.

I have to tell you, I'm a database guy by training, I thought it was joke. It was this platform that didn't do transactions. There was no query language in it. Everybody that I worked with in the database industry ignored this think because it didn't look anything like a database to them. You remember earlier I said that I didn't think that we'd been doing very interesting, innovative work in the RDBMS industry for a long time. That's a great example. This piece of really transformative technology came out of the consumer internet, shame on us that we couldn't have thought of this thing ourselves. I walked away from Oracle. I knew I didn't want to do another database company so I was willing to look at stuff with open eyes and with a clearer mind than I had when I was working for Oracle and before.

When I had a chance to sit down with Jeff Hammerbacher from Facebook with Amr Awadallah from Yahoo, when those guys explained to me what this platform was doing and that the rest of this industry was going to have this problem. The web guys were just a few years in the future, they weren't different in kind. Now, my co-founders all believed that same thing. We managed to band together and instead of creating three or four different companies, just create one company. We all saw the opportunity to deliver the power of Bugle, the magic that Yahoo had created and bottled, to enterprises generally. That those kinds of hard core, large scale data problems were going to transform the industry.

Andrew: I would like to ask you for an example the way that I did before with Sleepy Cat before I get into that, I'm curious about how you ended up in a conversation with people at Yahoo, Facebook, and other companies where they showed you a new way of looking at the world.

Mike: The answer is that way back in the day, Accel Partners had been an investor in Illustra Information Technologies...

Andrew: You're talking about '93, '97 days.

Mike: Exactly right, exactly right. I got to know a bunch of the Accel guys at the time. One of the partners, one of the venture investors at Axel is a guy named Ping Li. Ping had, all by himself, gotten excited about Hadoop as transformative technology. VCs got to walk around and look at all the new hair ball new stuff that's happening in the industry. When I left Oracle and when I started thinking about what to do next, one of the places that I had stopped was at Accel. Now, I had separately decided that maybe I should take another look at Hadoop. Ping's enthusiasm for Hadoop had convinced to recruit a couple of guys out of Facebook and Yahoo to come join Accel just as entrepreneurs and residents, people who draw a modest paycheck and get to dream. [inaudible 31:00] dreams of what their next company's going to be. That was Jeff Hammerbacher from Facebook and from Yahoo, those two guys I connected with through Accel and really because of Ping's vision for the kind of company that he wanted to fund. He was looking for the entrepreneurs he would back and we basically came together as that team.

Andrew: I see, okay, all right, and an example. What's an example of how Hadoop and one of your clients are working together?

Mike: Let me give an example that has been talked about at some of our conferences. There's a major retail bank that over the course of years has acquired a bunch of small regional banks. Every time it buys one of those guys it gets a new retail banking checkbook and debit card application, it gets a new home loan app, it gets new car loan app, it gets a couple of credit cards. Over the course of ten years, you can imagine that this bank owns dozens or even hundreds of desperate siloed little applications. Each of them running in production use every single day, can't shut them down. But you know for sure, you've got the same customer multiple times across all of those databases.

Remember 2008 when bad stuff started to happen in the financial markets. Understanding your mortgage exposure was all of a sudden red hot problem for banks generally, who's going to default? This bank builds the following Hadoop based system. Hadoop is able to take any kind of data, structured or complex from any location and then run arbitrary, interesting, sophisticated, complicated analytics over it.

The first thing they did was all of those dozens or hundreds of siloed databases, each of them got copied into a single big Hadoop processor. Now, you had all the data in one place for the first time in your life. Then they

did the following very smart thing. They said, “I want you to find everybody who’s defaulted on a mortgage.” Okay, fast, we can do that. Oh, I’m sorry, no, one initial step. All of those different properties I got, all of those different businesses I’ve been running, find me all of the instances of Michael’s in there. That is unify me as a Chase credit card customer and as maybe a retailing banking customer some where else into one coherent vision of me as a customer.

Find me everybody in my database that has defaulted on a mortgage. I want you to look backward in time, overall of my interactions with that person, and see if you can find any similarities. Here’s what they learned. If I stop getting direct deposits to my checking account and I switch from buying gasoline on my debit card at 7:00 in the evening to buying gasoline on my credit card at 2 in the afternoon, it’s a reasonable guess that I have lost my job. You could product with pretty high confidence that the next thing I’m going to do is start skipping credit card payments. Then I’m going to default on my car loan. Then the next thing is I’m going to default on my home loan.

Now, I laid that out to you and it all makes good sense and you ought to be able just to see that, right? That’s not how they did it. They actually observed this pattern in the data. They noticed that this was happening. That gives them a great tool because now they can look for that very first step. Hey, direct deposits have stopped and use that to predict future behavior. They can intervene earlier with their customers, to rescue them from default and, hey, by the way to eliminate a bunch of bankruptcies.

It’s a complicated analytic question over a lot of data and the answer was found observationally. That is by looking at the data and learning from it. That’s the kind of power that Hadoop delivers.

Andrew: I see, wow. I can see how you’d get excited to launch this business. I can see how you get excited to evangelize it to others. Based on the experience you had before, what were you going to do differently now with this business.

Mike: I told that Sleepy Cat was entirely funded on revenues. I knew that the opportunity for this technology was big. I believed deeply that we were, are creating a major new data management technology that the industry is going to have to take seriously. That’s not a grow it organically kind of a deal. You know you’re going to need to go raise venture capital in order to make that happen and then you want to be sure you’ve got the kind of backers that share your goals and views.

We think we're building a long-term stand alone IPO class company. I need to have A list VCs who are willing to invest forward and to swing for the fences. So, companies like Accel and Greylock, my A and B investors, Ameritech entered in [inaudible 00:17:09], those guys have had big successes in their past, they're well-connected. They can afford to wait for a billion dollar or better exit. They need to foot the company quickly. All of my co-founders have likewise had reasonable success in their past lives and can afford to put one over the center field fence this time.

We knew that we wanted to build a fast growing, well-capitalized capable company and we staffed up in order to do that. Both in terms of investors, founder background, and then in our early hiring.

Andrew: What about the culture within the company? In know the people you work with is very important to you. How do you start off with the right culture?

Mike: Oh, I hear you, culture, you broke up just a bit there. We learned this lesson, I learned this lesson at Sleepy Cat, the company is very much a reflection of the key people inside of it. My own bent is to be open and collaborative and communicative and [inaudible 00:36:44] be sure that everybody knows what I think is important and I try to understand what everybody else cares about. We need to be sure that everyone in the company knows what we are working on. The easiest way, in my view, to make it happen is to tell guys what I worry about. I sit down every two weeks and have an all hands meeting. The entire company shows up. I talk about the stuff that I'm thinking about, that I think is important. The key business leaders in each of the operational areas does the same thing. Then we throw the floor open. Employees can ask any question at all about the business.

We try very hard to be that communicative. I don't think it's any secret, you can hear it from me right now. I think we're on to a big, big thing. I think Cloudera is going to transform the industry. I try to make that as clear as I can to everybody on the team all the time. I want them to know how passionate and evangelical I am but I want them to be the same kind of people. When we meet somebody who's first question is about their compensation or their equity grant, we pretty much know that we didn't find a good Cloudera employee. When they talk about data, when they get excited about the opportunity, and when they look for killer use cases, that's somebody who really gets and is going to be passionate about the mission. Then, by the way, the company equity packages will be great and they'll come along. I want people to worry about customers and the business before

they worry about themselves.

Andrew: All right, I wrote down some of the advice that you gave Berkley students when you spoke to them so that we can talk about that here. You said, insist on torrid romance, what do you mean by that when it relates to business?

Mike: You're going to take a job in a company that is going grind you for years if you're [inaudible 38:31] in a business. If you start Cloudera for yourself, you're making a multiple year commitment and it's late nights and weekends. If you are not just crazy mad in love with the opportunity and technology and the people, you're not going to be able to do it. You've got to care deeply. It's got to be from a place other than your head that the passion comes. You've got believe in your start-up that way, if you're going to be successful.

Andrew: What about people today who see Groupon do really well and even their knock-offs make money and they say, "I'm going to jump into this, not forever, for a couple of years. Then flip the business in a couple of years or pocket the cash that I make over that period of time. Then I'll go and have the time to find what I'm really torridly passionately romantic about and pursue that."

Mike: You know, I think people can join a company like Groupon, by the way, I love those guys. I think they are going to be hugely successful.

Andrew: I know, they're clients of yours, but if they're knocking them off, they're saying I'm only building a new business just because I think the money is there today. Then the passion in another business or even in this business will come later on. I need the first big hit.

Mike: If you're planning to join a company that's being really successful, I can't really fault you. I think if you're going to go to Groupon for example, if you're going to join LinkedIn, you'll learn a lot, you'll make a decent living, you might even be able to cash out at a respectable level. I don't begrudge anybody that education. By the way, you look at my history, I was an engineer in start-ups for a long time before I was running them.

If you're saying, I'm going to start a company and my vision of this company is that I'm going to be committed to it for three years and I'm going to flip it quickly. I'm going to get out. I worry that you will not be

successful and the reason is that if your goal is just the money, you're not going to have the commitment to the customers, the technology, to stick it through difficult times. You really do need to be signed up for some tough sledding if you're going to make a start up successful. You must be passionate about the idea if it's your company.

If you're willing to join somebody else's company, sign on to their vision, and assist them. Then I think you can have a merely pecuniary motive. But if you're really looking to build a business, you've got to care about more than cash.

Andrew: Can you give an example of a time in your business, any one of these companies that we've talked about, Illustra, Molecular Applied Group, Sleepy Cat, Cloudera, if not for the passion you would have been gone. The only reason that it made sense to stick it out was the passion and the romanced.

Mike: I'll give you two examples. The first is, let me give you an example out of Molecular Applications Group. I was there for a year. I did not have the passion. It wasn't my field. I didn't know bio-medics long enough. I think that computer science is going to contribute to the cure of major cancers in the next 20 years and life will be much better. It wasn't my thing, I didn't know enough about it. I wasn't able to stick it through what were difficult times with in that company. That's an example, I walked away.

If I look back at my time at Illustra, I love what we were doing with [post-scripts] I love the idea that you would be able to use a relational database to deal with new kinds of data in new ways. I'd done my graduate research at Berkley on that. As we were building that organization, as customers would roll in the software into production, as it was not working and we had to parachute in at the eleventh hour and revive the patient from very near death. Those things are hard. The only way you can sustain that level of pressure is if you really believe deeply in the thing that you're doing. Fortunately at that time I'd been part of the technology from it's birth at Berkley. I really didn't even care enough to make that work.

Andrew: OK and Illustra's a company that sold for hundreds of millions of dollars. Right?

Mike: Yeah, \$440 million sale to Informix in 1996. It was a tremendous exit. It was the Netscape era but I think that Dick Williams, the CEO at Illustra did a tremendous job of finding a great buyer.

Andrew: This was a company that at one point that was near death? That was in that kind of critical situation.

Mike: No, I don't think that the company was ever near death. We had a product in the market that was, frankly, early. I've got this flip thing that I say but it turns out to be true in most instances. You know how to build reliable enterprise software? You build unreliable enterprise software and you ship it to customers. Your first customers often have a lot of difficulty making new platforms work. If you're going to be successful as a business, no matter what it takes, you've got to show up and make that work. That can be very, very difficult. Certainly we had a few early customers at Illustra that required heroic measures to make successful. We were able to deliver that level of effort.

Andrew: Can you give me an example of that? Because I can understand in the consumer space, you can ship anything out and call it beta, consumers will be very understanding or very disposal. You can improve and iterate and come up with something great. In enterprise, my understanding was that you need to come up with something great right out of the gate or else you lose your credibility, you lose your customers.

Mike: I think there's no shortage of stories in the history of databases, for example.

Andrew: How about one of yours, specific, without even mentioning a client name if you don't feel comfortable.

Mike: First, I want to say that even in consumer software reliability matters. A year plus ago Twitter had this reputation for being just hugely unreliable and props to the team over there, they have made that platform work. When was the last time you saw a fail well, it use to be a joke, it just doesn't happen any more. I want to be careful not to burn any bridges with my prior colleagues, so I don't know that I want to call out any particular client. But I can say, at every single one of my employers, and hey, by the way, that includes my current employer. We've had customers who have run into problems in production with the software that we distributed that required us to work very hard to get them back up and running.

The good news for enterprises is that when they choose a reputable hard-working vendor. They know they're going to get that level of attention. Do we all wish it's work flawlessly out of the box. Yeah, we do. Nobody really

expects that to happen that's why there's such things such as bug tracking systems and software infrastructure. What matters is how you react to problems in the field. That is really the defining characteristic of a successful company. You must do what it takes to get your customers successful, the big ones and the small ones. If you can do that, I think you'll be happy. I'm very proud with the way that Cloudera responded in the instances where that's happened. Hasn't been a lot but we certainly have. If I look back at Oracle, at Sleepy Cat, at MAG, and all the way back to Illustra, there had been no shortage of examples that I learned from. The lessons that we apply here are very much driven by that experience.

Andrew: Why is shipping flawed, a flawed version, a critical step in eventually getting to that good version?

Mike: No matter how, we as technologists talk about our testing strategy and the science that we bring to our computer science. The fact of the matter is that customers do stuff with software that you didn't anticipate. The deploy on platforms and on hardware that you didn't have access to. It is still very much an art to make software work well. The way that you get it reliable is to quickly identify and fix problems when they crop up. Yeah, we test and we test hard but we're not able to test every single deployment, every single workload, every single example of a dup deployment in the world and that's true for every single database vendor, every single infrastructure vendor.

We rely, in a certain way, on our customers to help us make the product better. Our earliest customers work the hardest to make that happen and we love them most for that. Over time there are just a lot fewer of those failures and flaws in the software because more of it's been ironed out. Now Hadoop is more mature than a lot of the enterprise software that I've worked with in the past. It came out of Google and then really in production in Yahoo and Facebook. It had petabytes scale deployment before it moved into the enterprise. The number and variety of problems are different in this case than they have been at my previous employers but it is a fact for software entrepreneurs.

Andrew: Is it fair to say that you have to put out something that doesn't work just so you can figure out why it doesn't work and then build the product properly.

Mike: I'm going to say even more strongly. Every single piece of software running today doesn't work. There are bugs in all software used anywhere. There are more bugs in software early in its life than late in its life. You know when any major vendor rolls out a new feature release, they expect to find and fix bugs. That's an artifact of the software industry for the past 30 years. It's not a new thing for open-source. It's not a new thing for these platforms. I wish that computer science really had some science behind it. Whenever you have to put science on the word, like social science or Christian science, it's usually not science. I think computer science is very much in that vain.

Andrew: That's true for everywhere. One of my first jobs was working for Dale Carnegie and Associates. These guys have been teaching how to win friends and influence people for decades. Teaching as a profession has been around for millennia. Whenever they'd roll out a new class, the company's that sent their employees, this is the first time. We're going to try it out. You'll give us some feedback. We'll improve it together but we'll be here for you to go the extra mile because you're our first customer. If in that environment you can't get it right the first time and you have bugs. Then how can software that's cutting edge not have bugs.

Mike: Yeah, and I want to make that point as clearly as I can towards the entrepreneurs in your audience. You need to commit to your customers absolutely. We talked about torrid romance earlier. You need to commit to your customers absolutely. You're going to roll stuff out. They're going to depend on it, they're depending on you and if you're going to be successful as a business person you've got to absolutely do what it takes to make them successful. It will be hard but it is your job.

Andrew: All right and I can see how that would be a passion job. One last point here, you also advise those students to seek pain. What do you mean by that?

Mike: Oh, yeah. You can almost always predict when a start up company will fail when you look for pain in its planning. Let me say this a different way. A lot of start-up companies are cool technology looking for a problem to solve. You do something that will have a graduate school. It's pretty cool and you really like and you say, we're going to start a company to commercialize this. If instead, you walk around to customers and you ask them, hey, man what hurts? What doesn't work for you today? What needs to be better?

What's too slow? What is it that causes you agony, that slows you down, that costs extra money?

You use what you learn there to plan a business. You develop products aimed at that. You will much more likely be successful. Cool technology is search of a solution sometimes finds its mark sometimes doesn't. If you really understand what customers need, if you sought their pain in planning your business then you've got something you can sell. People will pay money for stuff that will ease their pain. People won't pay money for stuff that's merely cool. That's a lesson that a lot of people, especially coming out of academia have not yet learned. It's one that I repeat when I talk to those audiences.

Mike: What's a pain that you saw in 2008 before you launched Cloudera? What's one pain, one example of pain, you said I can solve that.

Mike: This was a pain in many cases that customers didn't know they had. Some had an inkling, people are throwing data away all the time. They were archiving stuff out of their big enterprise warehouse after merely a couple of months because they didn't have enough capacity to store a years worth of transactional data. Throwing data away is an irrevocable act. It says that this information will never have any more value to me no matter what. It turns out that the Hadoop platform that Yahoo and Facebook and others were using can store a lot of stuff. Those guys were discovering huge latent value in large collections of old data. The guys at Google actually have a great quote. They say it's not that our algorithms are better than your algorithms; it's just that we have more data. You can learn stuff from lots of data. You can see patterns in lots of data that are just invisible in smaller amounts of data. Our conviction in 2008 was that people were throwing valuable stuff away and if we could just give them a way to keep it and ask questions of it, they would find new top line benefit, they'd make more money and we'd be able to carve their costs, we'd save them expense by storing more cheaply and accessing it more efficiently than ever before.

Andrew: The website is cloudera.com if people go over not only are they going to get to find out what Cloudera does and what Hadoop does, prime real estate, I know you're growing, is 'help wanted'.

Mike: Yes.

Andrew: Hopefully, some one in my audience will check that out and maybe they'll be a good fit. If nothing else, I hope people will check out Cloudera just to get a sense of the company and see how much it's grown and it's going to change the future.

Mike: Andrew, thanks a lot I really appreciate the talk.

Andrew: Thanks for doing the interview. Thank you all for watching. Bye.