



Creating the Lean Startup

How Eric Ries developed a scientific method for launching profitable companies

By Eric Ries | Sep 27, 2011

*Squandered capital, wasted efforts, shattered dreams. Eric Ries, author of **The Lean Startup**, is on a mission to save entrepreneurs from such a fate. Ries, a serial entrepreneur, co-founded IMVU, an online social network that made the Inc. 500 last year. Through trial and error at IMVU, Ries developed a methodical approach to launching companies that goes beyond bootstrapping. Now he's creating a movement.*

Stop me if you've heard this one before. Brilliant college kids sitting in a dorm are inventing the future. Heedless of boundaries, possessed of new technology and youthful enthusiasm, they build a company from scratch. Their early success allows them to raise money and bring an amazing new product to market. They hire their friends, assemble a superstar team, and dare the world to stop them.

More than a decade and several start-ups ago, that was me, launching my first company. It was 1999, and we were building a way for college kids to create online profiles for the purpose of sharing...with employers. Oops. I vividly remember the moment I realized my company was going to fail. My co-founder and I were at our wits' end. By 2001, the dot-com bubble had burst, and we had spent all our money. We tried desperately to raise more capital, and we could not. It was like a breakup scene from a Hollywood movie: It was raining, and we were arguing in the street. We couldn't agree even on where to walk next, and so we parted in anger, heading in opposite directions. As a metaphor for our company's failure, this image of the two of us, lost in the rain and drifting apart, is perfect.

If you've never experienced a failure like this, it is hard to describe the feeling. It's as if the world were falling out from under you. You feel you've been duped. The stories in the magazines are lies: Hard work and perseverance don't lead to success. Even worse, the many, many promises you've made to employees, friends, and family are not going to come true. Everyone who thought you were foolish for stepping out on your own was right.

The grim reality is that most start-ups fail. Most new products are not successful. Yet the story of perseverance, creative genius, and hard work persists. Why is it so popular? I think there is something deeply appealing about this modern-day rags-to-riches story. It makes success seem inevitable if you just have the right stuff. If we build it, they will come. When we fail, as so many of us do, we have a ready-made excuse: We weren't in the right place at the right time—we didn't have the right stuff.

After more than 10 years as an entrepreneur, I have come to reject that line of thinking. Start-up success is not a consequence of good genes or being in the right place at the right time. Success can be engineered by following the right process, which means it can be learned, which means it can be taught.

Let me tell you a second start-up story. It's now 2004, and a group of founders have just started a company. They have a huge vision: to change the way people communicate online by using a new technology called avatars.

I'm in this second story, too. I'm a co-founder and chief technology officer of this company, IMVU. Although my co-founders and I were determined to do things differently, we ended up making a lot of mistakes. Despite various setbacks, the methods we developed over time at IMVU have become the basis for a movement of entrepreneurs around the world. It represents a new approach to creating continuous innovation. I call it the Lean Startup.

Our "Brilliant" Business Plan

The five of us involved in the founding of IMVU aspired to be serious strategic thinkers. Each of us had participated in previous ventures that had failed, and we were loath to repeat that experience. Our main concerns in the early days dealt with the following questions: What should we build and for whom? What market could we enter and dominate?

We decided on the instant messaging market. In 2004, that market had hundreds of millions of customers, the majority of whom were not paying for the privilege. Large companies such as AOL, Microsoft, and Yahoo ran their IM networks as a loss leader for other services while making modest amounts of money through advertising. The common wisdom was that it was more or less impossible to bring a new IM network to market without spending an extraordinary amount of money on marketing.

The grim reality is that most start-ups fail. Most new products are not successful.

At IMVU, our strategy was to build a product that would combine the mass appeal of traditional IM with the high revenue per customer of video games. Because of the near impossibility of bringing a new IM network to market, we decided to make our product compatible with existing IM networks. Customers would be able to chat online using their IMVU avatars without having to switch IM providers or learn a new user interface. They wouldn't have to persuade their friends to switch, either.

We thought the third point was essential. Every IM communication would come embedded with an invitation to join IMVU. Our product would be inherently viral, spreading throughout the existing IM networks like an epidemic. To spur rapid growth, it was important that our product be compatible with as many IM networks as possible.

With this strategy in place, my co-founders and I began a period of intense work. As the CTO, it was my responsibility to, among other things, write the software that would support the various IM networks. Because we had limited funding, we gave ourselves a hard deadline of six months to launch the product and attract our first paying customers. It was a grueling schedule, but we were determined to launch on time.

The project was so large and complex and had so many moving parts that we had to cut a lot of corners to get it done on schedule. I won't mince words: The first version was terrible. We spent endless hours arguing about which bugs to fix and which we could live with, which features to cut and which to cram in. It was a wonderful and terrifying time. We were full of hope about the possibilities for success and full of fear about the consequences of shipping a bad product.

I was worried that the low quality of the product would tarnish my reputation as an engineer. People would think I didn't know how to build a quality product. We envisioned the damning newspaper headlines: *Inept Entrepreneurs Build Dreadful Product*.

Six months later, teeth clenched and apologies at the ready, we released our website to the public. And then—nothing happened! It turned out that our fears were unfounded, because nobody even tried our product.

We Resort to Talking to Customers

Over the ensuing weeks and months, we labored to make the product better. We eventually learned how to change the product's positioning so that customers at least would download it. We were making improvements continuously, launching bug fixes and new changes daily. However, despite our best efforts, we were able to persuade only a pathetically small number of people to pay \$29.95 for the product.

Eventually, out of desperation, we began bringing people into our office for in-person interviews and usability tests. Imagine a 17-year-old girl sitting down with us at a computer. We say, "Try this new product; it's IMVU." She chooses her avatar and says, "Oh, this is really fun." She's customizing the avatar, deciding how it's going to look. Then we say, "All right, it's time to download the instant messaging add-on," and she responds, "What's that?"

"Well, it's this thing that interoperates with the instant messaging client," we say. She has no idea what we're talking about. But because she's in the room with us, we're able to talk her into doing it. Then we say, "OK, invite one of your friends to chat." And she says, "No way!" We say, "Why not?" And she says, "Well, I don't know if this thing is cool yet. You want me to risk inviting one of my friends? If it sucks, they're going to think I suck, right?" And we say, "No, no, it's going to be so much fun once you get the person in there; it's a social product." She looks at us, her face filled with doubt; you can see that this is a deal breaker.

Of course, the first time I had that experience, I said, "It's all right; it's just this one person. Send her away, and get me a new one." Then the second customer comes in and says the same thing. Then the third customer comes in, and it's the same thing. No matter how stubborn you are, you start to see there is something wrong.

In response to the feedback, we created ChatNow, a feature that lets you push a button and be randomly matched with somebody else anywhere in the world. The only thing you have in common is that you pushed the button at the same time. All of a sudden, people were saying, "Oh, this is fun!"

Any effort that is not absolutely necessary for learning what customers want should be eliminated.

Then, maybe they would meet somebody they thought was cool. They'd say, "Hey, that guy was neat; I want to add him to my buddy list. Where's my buddy list?" And we'd say, "Oh, no, you don't want a new buddy list; you want to use your regular AOL buddy list." You could see their eyes go wide, and they'd say, "Are you kidding me? A stranger on my buddy list?" To which we'd respond, "Yes; otherwise you'd have to download a whole new IM program with a new buddy list." And they'd say, "Do you have any idea how many IM programs I already run?"

"No," we'd say. "One or two, maybe?" That's how many each of us used. To which the teenager would say, "Duh! I run eight." It started to dawn on us that our concept was flawed.

Our early adopters didn't think that having to learn a new IM program was a barrier. Even more surprising, our assumption that customers would want to use IMVU primarily with their existing friends was also wrong. They wanted to make new friends, an activity that 3-D avatars are particularly well suited to facilitating. Bit by bit, customers tore apart our seemingly brilliant initial strategy.

Was It All a Waste?

I wish I could say that I was the one to realize our mistake and suggest the solution, but in truth, I was the last to admit the problem. I had slaved over the software that was required to make our system work with other IM networks. When it came time to abandon that original strategy, almost all of my work—thousands of lines of code—was thrown out. That was really depressing.

I wondered, in light of the fact that my work turned out to be a waste of time and energy, Would the company have been just as well off if I had spent the past six months on a beach sipping umbrella drinks?

There is always one last refuge for people aching to justify their failure. I consoled myself with the fact that if we hadn't built our first product—mistakes and all—we never would have learned these important insights about our customers. We never would have learned that our strategy was flawed. There is truth in this excuse: What we learned during those critical early months set IMVU on a path that would lead to our eventual breakout success. Today, IMVU is a profitable company with more than \$50 million in annual revenue and more than 100 employees. IMVU customers have created more than 60 million avatars.

For a time, this consolation made me feel better, but some questions still bothered me. If the goal was to learn important insights about customers, why did it take so long? How much of our effort actually contributed to that learning? Could we have learned those lessons earlier if I hadn't been so focused on making the product "better" by adding features and fixing bugs? I had created software to support more than a dozen IM networks. Was

this really necessary to test our assumptions? Could we have gotten the same feedback from our customers with half as many IM networks? With only three? With only one?

Here's the question that kept me up nights: Did we have to support any IM networks at all? Is it possible that we could have discovered how flawed our assumptions were without building anything? What if, before building anything, we simply had offered customers the opportunity to download the product solely on the basis of its proposed features? Almost no one was willing to use our original product, so we wouldn't have had to do much apologizing when we failed to deliver.

In other words, which of our efforts were creating value, and which were wasteful? This question is at the heart of the lean-manufacturing revolution; it is the first question any lean-manufacturing adherent is trained to ask. Learning to see waste and systematically eliminate it has allowed lean companies such as Toyota to dominate entire industries. Lean thinking defines value as "providing benefit to the customer"; anything else is waste. But in a start-up, who the customer is and what the customer might find valuable are often unknown. I realized that for start-ups, we needed a new definition of value. The real progress we had made at IMVU was what we had learned over those first months about what creates value for customers.

Learning as Fast as We Can

If learning is the essential unit of progress for start-ups, any effort that is not absolutely necessary for learning what customers want should be eliminated. So how do we do that? By building what I call a minimum viable product—or MVP. It helps entrepreneurs start the process of learning as quickly as possible. Unlike a prototype or concept test, an MVP is designed not just to answer product design or technical questions. Its goal is to test fundamental business hypotheses.

Yes, MVPs are sometimes perceived as low quality by customers. When that happens, it's an opportunity to learn what attributes customers care about. This is infinitely better than mere speculation or whiteboard strategizing, because it provides a solid empirical foundation on which to build.

Sometimes, however, customers react quite differently. Many famous products were released in what might be called a low-quality state, and customers loved them. Imagine if Craig Newmark, in the early days of Craigslist, had refused to publish his humble e-mail newsletter because it lacked high design.

Customers don't care how long something takes to build. They care only that it serves their needs.

In the early days of IMVU, our avatars were locked in one place, unable to move around the screen. The reason? We had not yet tackled the difficult task of creating the technology that would allow avatars to walk around their virtual environments. In the video game industry, the standard is that avatars should move fluidly as they walk, avoid obstacles in their path, and take an intelligent route toward their destination. Best-selling games such as Electronic Arts's *The Sims* work on this principle. We didn't want to ship a low-quality version of this feature, so we opted instead to ship with stationary avatars.

Feedback from the customers was very consistent: They wanted the ability to move their avatars around. We took this as bad news, because it meant we would have to spend considerable amounts of time and money on a high-quality solution similar to *The Sims*. But before we committed ourselves to that path, we decided to try an experiment. We used a simple hack, which felt almost like cheating. We changed the product so that customers could click where they wanted their avatar to go, and the avatar would teleport there instantly. No walking, no obstacle avoidance. The avatar disappeared and then reappeared an instant later in the new place. We couldn't afford fancy teleportation graphics or sound effects.

Imagine our surprise when we started to get positive customer feedback. We never asked about the movement feature directly (we were too embarrassed). But when asked to name the things about IMVU they liked best, customers consistently listed avatar teleportation among the top three. It outperformed features that had taken much more time and money to make.

Customers don't care how much time something takes to build. They care only that it serves their needs. Our customers preferred the quick teleportation feature because it allowed them to get where they wanted to go as fast as possible. In retrospect, this makes sense. Wouldn't we all like to get wherever we're going in an instant? Our expensive real-world approach was beaten handily by a cool fantasy-world feature that cost much less but that our customers preferred. So which version of the product is low quality, again?

Going Lean

At its heart, a start-up is a catalyst that transforms ideas into products and services. As customers interact with those products and services, they generate feedback and data. The feedback is both qualitative (what they like and don't like) and quantitative (how many people use it and find it valuable). As we learned the hard way at IMVU, the products a start-up builds are really experiments. Learning about how to build a sustainable business is the outcome of those experiments. Each experiment essentially follows a three-step process: Build, measure, learn.

Many people have professional training that emphasizes one element of this three-step loop. For engineers like me, it's learning to build things as efficiently as possible. Plenty of entrepreneurs obsess over data and metrics. The truth is that none of these activities by itself is of paramount importance. Instead, we need to focus our energies on minimizing the total time through this loop. That way, we can avoid much of the waste that plagues start-ups today. As in lean manufacturing, learning where and when to invest energy results in saving time and money.

The Lean Startup method builds capital-efficient companies because it allows start-ups to recognize that it's time to pivot—or change direction—sooner, creating less waste of time and money. I named this loop "build, measure, learn" because the activities happen in that order. But the planning really works in the reverse order: We figure out what we need to learn, then figure out what we need to measure to get that knowledge, and then figure out what product we need to build to run that experiment and get that measurement.

So what would organizations look like if everyone were armed with Lean Startup principles? For one thing, we would all insist that assumptions about what customers want be stated explicitly and tested rigorously. We would look to eliminate waste, not build castles in the sky. We would respond to failures and setbacks with honesty and learning, not with recriminations and blame. Most of all, we would stop wasting people's time.

This article is adapted from [The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses](#), by Eric Ries, published this fall by Crown Business.

Eric Ries will discuss his new book and answer questions during a live video chat on October 5 at noon Eastern Time. To watch the chat and participate, go to www.inc.com/live.

Copyright © 2012 Mansueto Ventures LLC. All rights reserved.
Inc.com, 7 World Trade Center, New York, NY 10007-2195.