

Riverbed Technology is a developer of software, improving the performance of distributed systems through a caching, compression and protocol acceleration offering.

Startup Profile

Riverbed

Jerry Kennelly, Chief Executive Officer and Dr. Steven McCanne, Chief Technology Officer

“We’ve got fairly sophisticated analysis of the traffic that is flowing through our boxes; that will tell us directly where the problems are. We can work with the customer in terms of their priorities, what their pains are.”

Tell me a little about your company, and what you offer the marketplace.

JK: We're basically a software company, but we market our technology as a hardware appliance with a software mounted on it. You put one at either end of a WAN network link and accelerate data transfer - data transfer across a WAN; it takes the latency bandwidth constraints out of data transfer. The target markets are the Fortune 100 or 500 size corporations, who have distributed operations; they have the capability to do data transfer at LAN speed in the worldwide distributed computing environment and this allows them to optimize where they put their servers and how they do their storage and do their backup. It's a very powerful ROI.

What is the competitive landscape like? Tell me a little about that.

JK: There are no direct head-on competitors. There are people trying to do this type of functionality in one of three categories - people who have compression technologies, people who file caching technology, and people who do quality of service technology, that affect the way data comes across a WAN link. None of them actually do what we do. We sort of synthesize all of those things into a single technology and we're able to leapfrog everyone in terms of performance and delivery of the data.

When will Steelhead be launched?

JK: The product is in Beta now - it will launch in the first half of 2004.

How did you find VCs? Did they come to you or vice-versa?

JK: They were in our A round, which we did in the fall of 2002. We started the company in a classic way - the two founders had the idea and we quit our jobs and worked without pay, and did the product idea and a business plan and we visited the VC community. We had some contacts there, and we cold-called and some we had prior contacts, and we came up with Accel and Lightspeed. Utah Ventures joined in the B round.

So what is your big challenge going forward?

JK: It is fundamentally just execution. When you do these things, there's more risk. Financing risk, market risk, product risk, exe-

cution risk. We are past the first three, and those are the big ones. Now it's just execution risk. So it's actually looking pretty good.

Tell me your impressions of the marketplace, and what you can offer it with your product.

SM: Applications are chatty, and protocols that support applications, were written to run over a LAN. When you try to run them over a WAN, you run into latency problems, which means you have to go back and forth between client and server many times. When it's on the LAN, it happens very quickly. So to extend the traffic metaphor to help understand that, you have a bunch of cars all trying to go over a bridge at the same time, for example. And adding lanes to the bridge would help alleviate that problem. But the latency problem is more like if you have a truck that holds a certain amount of load, and you've got to transport that load back and forth across the bridge. It's going to take you a hundred trips if you have 100 tons of stuff, no matter how many lanes there are on the freeway. You're not going to get the job done quickly. In many ways, the TCP protocol is kind of like that truck; as you try to stuff more and more data through a TCP connection, there are problems with protocol. You can only stuff so many bits into the TCP window per round bit, like the truck. Even if you have bandwidth available, there are still problems.

So what we are all about is putting the device at either side of the network and the WAN link toward the servers and the users. It does two things; it moves the data in this sort of segmented, high-level, application-independent format, so that what you put in the truck are references to the application structures that are stored on either side of the link. And therefore you can get much larger loads in that truck for a given roundtrip. The other element is that we try to predict what the user is going to want next. Because the nature of the application is that they serially go through these steps. If every step takes the truck going back and forth across the bridge, it's going to be slow. So we try to say 'I've seen this application's behavior, I understand it's behavior, and I know the next 50 things it's going to do. So let's do them all at once; let's stuff all this information in my

Year established: n/a
 Employees: n/a
 Total investments received: n/a
 Total rounds of funding: n/a
 Select investors: n/a
 Most likely exit strategy: n/a
 Next expected round of funding: n/a
 Amount sought: n/a
 Expect date to reach profitability: n/a

(continued)

truck, and I can, because I have sufficient representation, and let's move that truck over the bridge one time and accomplish the end goal of having actually gone back and forth 50 times. So that's a metaphor that's been used a lot in thinking about bandwidth. And when you go and talk to customers, they really understand the bandwidth problem. They can buy more bandwidth if they need to; they can throw money at it. Latency? You can't pay your service provider more to overcome the speed of light. So that's sort of what we're all about.

How does caching fit into the model? Your literature lists it as a primary feature.

SM: I don't really think of what we do as caching; to me, we store data in an application-independent format. By virtue of prediction, we move results over the network much more quickly. A cache accomplishes that goal in a different way; it stores the data in a way an application would store it, and then tries to serve it up, and run a consistency protocol across the WAN to make sure the copy is up to date. Other analysts have looked at our technology, because it doesn't really have a prior, and have labeled it a fusion of caching and protocol acceleration and compression. That's a way to think about it, but on the flip side it's not really cache and it's not really compression and it's not really pure protocol acceleration. It's a way of accomplishing the blend of those three technologies.

What is the state of the product in its beta phase?

SM: We're pretty far along in our beta program that we started a while ago, and we're on the verge of shipping our generally available product. [There are 50 beta customers right now.]

What are technology challenges ahead? Tell me a little about that.

The beauty of the architecture is that it provides benefit across all TCP applications out of the gate; there are certain bottlenecks that you can't address just at the TCP layer. So we have this foundation, a core, that does TCP acceleration and segmentation, and

then we have application-specific modules, and we've focused on our version one product, delivering an e-mail and file access acceleration. Both of those elements work quite well; beta has shown they work quite well. They can always be improved. And there are other application modules that we can add to the mix over time, to do more specific things around Web application, Web services, and maybe some database applications. We're letting our customers lead us down the path, in terms of where we are going to focus our resources, in terms of lowest-hanging fruit and best opportunity.

How do the customers direct you? What is the interaction like?

SM: We've got fairly sophisticated ways to do statistical analysis of the traffic that is flowing through our boxes; that will tell us directly where the problems are. We can work with the customer in terms of their priorities, what their user pains are, and where they want to consolidate more of their infrastructure. We push those projects to the top of the list.

Is there heavy competition for you?

SM: I think that this has the potential to be a big deal in enterprises and make a big impact on the way IT architects and CIOs think about their architecture. I think it will get the attention of the big players; there is nobody out there who has the technology that we have today. There are a clump of companies that do compression, there are a clump of companies that do caching, but nobody has done what we've done. Which doesn't mean they won't see it and want to come after that opportunity. I expect that as we go forward. ■

VR

HOW THEY MAKE MONEY

Riverbed offers caching, compression and protocol acceleration in one solution.

MANAGEMENT TEAM

Jerry Kennelly, Chief Executive Officer

Dr. Steven McCanne, Chief Technology Officer

Eric Wolford, VP, Marketing & Business Development

Alan Saldich, Director, Marketing