HOW A SOFTWARE COMPANY DERISKED ITS MOVE TO THE CLOUD

A Conversation with Risk Management Solutions' Hemant Shah, CEO, and Eric Yau, General Manager, Software

F OR MOST ENTERPRISE TECHNOLOGY companies, moving from on-premises software applications to a modern cloud-based platform is a business requirement. For Risk Management Solutions (RMS), the leading global provider of catastrophe-modeling software for insurers and the public sector, it was a customer-driven imperative. The insurers, reinsurers, financial institutions, and public agencies that RMS serves need technology that delivers better information and more powerful analytics for their core risk- and exposure-management processes.

Recently, Jon Brock, an associate director in BCG's London office who leads the large-project derisking and recovery business of the Technology Advantage practice, spoke with Hemant Shah, RMS's CEO and cofounder, and Eric Yau, its general manager, software, about the process the company followed to produce RMS(one) 1.0 with a high-stakes, no-leeway launch deadline: the company's annual customer conference in early 2017.

Tell us about the benefits customers get from the new platform.

Shah: Like every other industry, the global insurance industry is increasingly looking to compete on analytics. And to do so, companies need not only to master data and analytics but also to operationalize that data. Market leaders and startups are all moving analytics from the back office into the heart of how they do business. An agile analytics enterprise requires the kind of technology, platform, and systems that we are now delivering.

Yau: So many of our customers' mission-critical business processes—from risk selection to portfolio management and capacity management—are completely interconnected. But they rely on disparate sources of information and unintegrated analytics approaches and information delivery mechanisms.

In listening to our customers' challenges, we realized that we had to help them by delivering a modern analytics platform that could be used to deliver a better, faster, consistent, and auditable set of analytics.

Getting the technical solution right was an important step over the past 18 months. How did you minimize the risks to ensure that you would be ready to deliver the end product successfully and on time?

Yau: The first thing we had to do was translate our product strategy into a set of clearly defined fundamental requirements for our new software platform. We were careful to define both the functional and nonfunctional proof points. We needed to see early on that our technical solution would meet both our short-term and longer-term requirements. We undertook a very methodical and repeated set of POVs—proofs of value—to validate our decisions around our technical architecture. We could not shortchange these steps.

Once we felt comfortable with our technical choices, we moved to the first phase: developing an integrated, high-quality internal release. We built the system end to end, testing the quality and performance of the technologies and how well they integrated before we built too much functionality.

Hemant, how, in your view, did it go?

Shah: To get RMS(one) right, we had to make three strategic choices.

We had to recognize the complexity of the problem. We had to put together a design and an architecture that orchestrated the most fundamental challenges right up front instead of kicking the can down the road.

Second, we decided to embrace the open-source ecosystem of big data analytics technologies, rather than develop our own tools.

Third, we recognized that this transformation required a level of talent and experience that we just didn't have. So we had to make some difficult decisions, including changing leadership.

These people changes extended up and down the entire team, including a whole new set of capabilities to augment the deep industry and modeling-domain knowledge we already had. We brought in people with many new skills, including big data analytics, that we'd previously lacked, as well as people experienced in building mission-critical, enterprise-class SaaS, or software as a service, platforms.

One of your big decisions was embedding RiskLink, the company's original on-premises software application, directly in the platform and allowing the platform to run the old models along with the new high-definition models. Tell us about that.

Shah: That turned out to be one of the most cathartic decisions we made. And there's a cultural element to it that is quite important. To build our platform, we needed to get all our teams collaborating—our modeling teams, our data science teams, and our software-engineering teams. That was quite a challenge.

Our first attempt at implementation failed precisely because of the cultural divide. The software team's idea of this process was, "We're going to build a next-generation platform, and thank you very much, modelers, for your input." The modeling team thought the software team was just building something to implement their models.

It was almost as if these teams were working in different companies. Then when things didn't work well, the finger-pointing began.

How did you manage to overcome this divide?

Shah: One way was by implementing a loosely coupled model execution framework that let us releverage some of the existing engines. That was one of many things that resulted from having software leaders and modeling leaders actually collaborating on the design and implementation. They could learn from one another and then leverage each other's differentiated capabilities.

Can you elaborate a bit on the steps you took to create collaboration among teams?

Shah: First of all, leaders have to acknowledge—and be vocal about the fact—that no individual team can solve the problem on its own. That takes a certain amount of humility, and it was fundamental to breaking down the barriers. With that in mind, we put the product, development, and CloudOps teams in the same room as much as possible. The best way to get people to work together is to bring them physically together and make them cross paths more often than may seem necessary.

In hiring—or firing—one of my key criteria was cultural fit. The Bay Area is filled with A-plus technical talent who can't work effectively in a multidisciplinary team.

And finally, we celebrated success across the teams.

Yau: In fact, I believe our greatest success so far isn't our new software platform or our models; it's the fact that we've been able to build an integrated team of diversely skilled people who work very effectively together as one team. We've seen this have a very positive impact for us and, more important, for our customers.

Early on, you decided to officially launch the first release of RMS(one) at your annual conference in March 2017. What steps did you take to minimize risks in the development plan and ensure that you would hit that date?

Yau: We knew our new platform could deliver significantly improved functionality across many dimensions. For our first release, we needed to focus our development efforts on an initial set of use cases, delivered exceedingly well. Determining this first MVP, or minimal viable product, was challenging at first, given the many directions we could take. We had to make some tough choices.

So is that what gave you contingency in your development plan? Was it working

toward an MVP that gave you enough time to complete the development, handle any overruns, and carry out adequate testing before the conference?

Shah: Yes. And remember, we had also tried to reduce the highest-risk variables. The MVP decision had to do with the functional features. We'd already reduced complexity in the architecture, scalability, performance, and quality. And during this time, we also matured our development and delivery processes, so we were continually improving our ability to estimate time and scope.

That brings us to another point: DevOps and your automation-testing work. You were passionate about getting automation testing underway early on, although not everyone understood its value. Now that the first release is done and you're moving to monthly releases, how well would you say DevOps has positioned you?

Yau: Once we started tackling automation, the team moved fast—automating the testing, the daily builds, the continuous integration, the deployment, even the environment provisioning.

Bottom line: when you're delivering a SaaS offering—one that is mission critical—you must have mature quality management and release processes that you've automated, or you will inevitably end up hitting a brick wall.

Shah: Eric really brought me along on this. In my early days as CEO, I would listen politely as he'd make the case for investing in these underlying processes. And now, I have to say, I see the benefit. Eric and his team not only hit 1.0, as expected, but they also hit 1.1, 1.2, 1.3—and now they're working on 1.4. This is monthly.

Let's talk about customer feedback. How important was it in the early stages?

Yau: It's very important—and it's important all the time. Equally critical, though, is

our ability to interpret a customer's unique requirements versus those needed by the market as a whole. As a product company, we are focused on serving the market. But we must also make sure that we provide an open and mature set of APIs and integration points that allow our customers to build and extend as they need.

Now that you've been through this transformation, is there anything you'd like to add about the experience?

Shah: Yes. It's one thing to talk about change and transformation; it's another thing to achieve them. Four years ago, I threw around the term *transformation* like it was an easy thing to do. We failed in our first round because we didn't put any of the fundamental changes in place—technology, leadership, people, or process.

I now have a much deeper appreciation for how difficult it is to execute transformation. We feel like battle-scarred veterans. As we head into the next round of battles, we are armed: we built a leadership team with the right mindset and the capacity for change. That's a precious thing.

For more on this topic, see "Taking the Risk out of Digital Projects," BCG article, February 2018.

About the Author

Jon Brock is an associate director in the London office of The Boston Consulting Group. He leads the large-project derisking and recovery business of the Technology Advantage practice. You may contact him by email at brock.jon@bcg.com.

The Boston Consulting Group (BCG) is a global management consulting firm and the world's leading advisor on business strategy. We partner with clients from the private, public, and not-for-profit sectors in all regions to identify their highest-value opportunities, address their most critical challenges, and transform their enterprises. Our customized approach combines deep insight into the dynamics of companies and markets with close collaboration at all levels of the client organization. This ensures that our clients achieve sustainable competitive advantage, build more capable organizations, and secure lasting results. Founded in 1963, BCG is a private company with offices in more than 90 cities in 50 countries. For more information, please visit bcg.com.