SAS to Statistica: The Great Dell Migration

Part One: People

Software

Introduction

What do monarch butterflies, sockeye salmon, wildebeest, sperm whales, red crabs and Dell have in common? They have all undertaken a great migration fraught with challenge and peril, but one which ensures survival.

Within weeks of acquiring the advanced analytics product Statistica, Dell set out on one of the most ambitious migration projects since the company was founded: a window ending December 31 (approximately six months) during which all users would move off of SAS and adopt Statistica as the core analytics platform companywide.

Whether your organization is migrating 10 users or 10,000 users in a 6-week or 6-year project, we think you'll end up asking questions like the ones we asked ourselves during the project. This e-book describes the approaches to people, process and technology that contribute to the success of a migration from SAS to Statistica.

Dell set out on one of the most ambitious migration projects since the company was founded.

To give you insight into the scope and objectives of our migration project, here are a few high-level results:

- Hundreds of users migrated worldwide
- Substantial, bottom-line impact due to saved fees
- 300+ projects across multiple business units migrated from SAS to Statistica
- Migration project team consisting of 12 points of contact for users
- Overall project duration of less than nine months, with user migration in less than six months

Of course, every organization is different. We at Dell achieved these results in so little time because of our commitment and the extraordinary efforts of teams all around the company, as you'll see in Part 2. We can't guarantee similar savings from every migration to Statistica, but we have seen a bottom-line impact that similar companies making a similar transition could expect to see. In any event, we invite you to learn from our lessons.

First, why migrate to Statistica?

Without a good answer to this question, the others don't matter.

Your mileage may vary, but Dell was spending a great deal of money every year to keep SAS in place. Of course, everybody wanted to save money, but everybody also knew the migration was going to involve a lot of risk, work and change, so we had to justify the effort.

You want to enable analytics in the enterprise. We've all read about it in books, magazines and news articles — we need to do something about analytics and big data. Organizations that embed analytics within all parts of their business to make faster decisions and improve decision making, planning and forecasting have a distinct competitive advantage. Unfortunately, there is a skills shortage, so we need a software package that plays well with existing IT investments and is sufficiently easy to use. The goal is to enable all users — experts and line-of-business users alike — to make the most of their data.

You want to lower the cost of software licensing. Even before Dell acquired Statistica (formerly of StatSoft), customers and prospects had been telling us that they wanted analytics software that was less expensive. That's no surprise, but more important, they wanted software that made analytics accessible to more people in their organization. But they didn't know what a migration project would entail and like most companies, they were concerned about the downside of an unsuccessful one.

You can find better value for money. A company the size of Dell derives significant value and competitive advantage from applying analytics in areas like marketing, price optimization, forecasting, technical support, supply chain optimization, preventive maintenance and financial credit risk analysis. We believed we could get better value for less money with the full range of analytics muscle and ease of use we saw in Statistica.

You want your analysts to analyze data, not write code. SAS is powerful, but you have to hire people to write and maintain code and administer the complex system to get the most out of it. We, along with many of our customers, were having increasing trouble finding good replacements for the SAS-savvy people who were leaving or retiring, so the cost of keeping SAS was rising beyond our annual licensing fees.

It's the best way to improve the product. You know that the path to a better product leads from your front door to your loading dock. Given that Dell is making significant investments in Statistica, it made all the sense in the world to use it ourselves and see what our customers were experiencing.

It's part of why Dell acquired Statistica in the first place. Dell's acquisition strategy is driven in part by the opportunity to roll out and use acquired products internally, improve them and save money with them, then take that message out to our customers. When prospects ask whether Dell uses the product they are promoting, our sales teams need to be able to answer, "Yes, we drink our own champagne."

How many people are you going to affect? Where are they?

As with any software product, most of your users are casual, working with a subset of the product functions to accomplish their daily tasks. Then you have a group of power users who eat, sleep and breathe the product.

When we started the project, we quickly identified hundreds of users whom we needed to move to Statistica, including 170 users in Dell Global Analytics (DGA). DGA provides analytics expertise and support to different functional organizations throughout the company — for example, finance, the customer service center, dell.com, marketing and sales, supply chain, operations, pricing and product management — that do not have their own internal analytics expertise.

When we started the project, we quickly identified hundreds of users whom we needed to move to Statistica.

Our marketing operations teams all over the world rely heavily on BI and analytics tools, as do Technical Support, Supply Chain and Dell Financial Services. In short, analytics is pervasive throughout Dell and is a big part of our competitive advantage. The migration from SAS to Statistica affected our data analysts and users in all of those groups. But more important, since analytics are ubiquitous at Dell and embedded in our systems and processes, there are many more people who consume and rely on the output. Any adverse change to this output could drastically impact the business.

How do your people use analytics?

In almost every company, analytics are finding their way into marketing to help make sense of what customers and prospects are saying, tweeting, looking for and buying. We use analytics to personalize offers, attract prospects and keep existing customers. We study things like customer churn, cross-sell/upsell opportunities, customer sentiment and satisfaction.

In our daily operations, we apply analysis to improve the quality of our manufacturing processes. In Professional Services, our services teams embed it in customer solutions as part of service engagements. For Customer Support, we use analytics for predictive/prescriptive/preventive maintenance on hardware products we sell.

> We also use analytics to assess and control the potential cost and risk of our decisions.

Dell Financial Services relies heavily on analytics for modeling, assessing credit risk and detecting fraud. Their models are closely tied to forecasts and bank rates, so statistical analysis is part of what they do day in and day out.

Like you, we use analytics to make decisions. But since not all of our decisions are perfect, we also use analytics to assess and control the potential cost and risk of our decisions.

Analytics use case: Customer churn

Suppose that a logistic regression model tells us we're in danger of losing five customers to churn because they are likely to buy another company's products instead of Dell's. We decide to offer each of them a $50 discount on a Dell product and they take it. So it costs us $250 to keep them as customers. That sounds like a good use of $250. But suppose it wasn't the right model to use and the customers weren't really in danger of churn. Then we've wasted $250.

> With the right model and analytics, it would have cost us only $50 to save $500.

Now look at the flip side. Suppose our Random Forest model tells us that a particular customer is happy and not in danger of churning, so we don't offer her the discount. If it turns out she buys another company's product, we might lose $500 of revenue. With the right model and analytics, it would have cost us only $50 to save $500.

So, you might say that the Random Forest is a better model because it results in just one error, but it's a more costly error than the five combined from the logistic regression model. Our models need to take that into account, and these tradeoffs of model accuracy and costs associated with incorrect predictions are easy to illustrate with Statistica.

We use analytics to evaluate models based on not only the overall error rate but also the type of error that the model is making. It allows us to make more-precise, better-informed decisions about the models we apply to business factors like customer churn, product upsell, service level agreements, competitive pricing and delivery dates.

Analytics use case: Direct mail campaign

Assumptions

- $2 to mail each prospect
- 1 out of 100 will buy
- $220 profit for each response

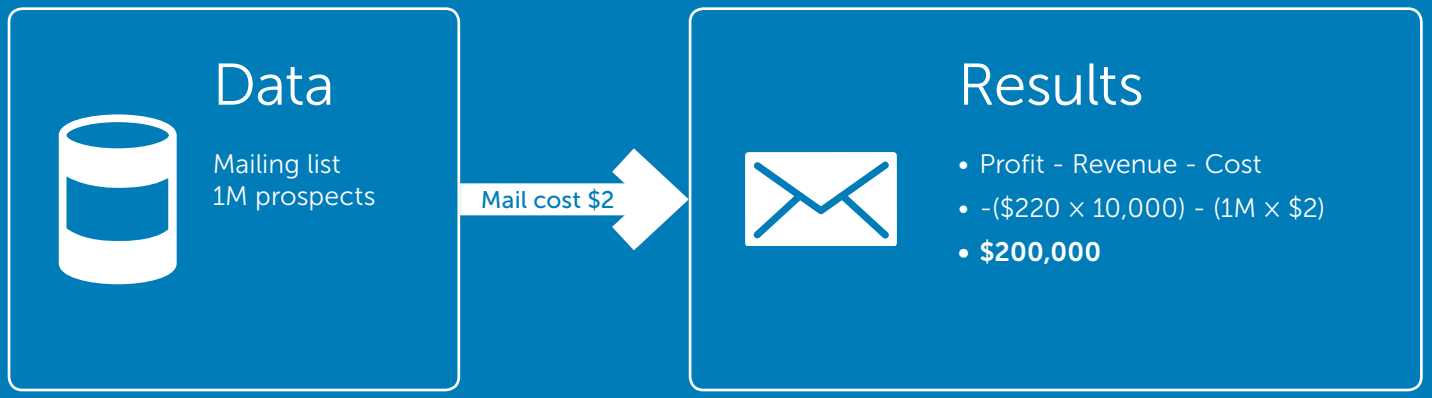

The value of a prediction

Assumptions

Analytics model output:

- 25% of the entire list are 3x more likely to respond

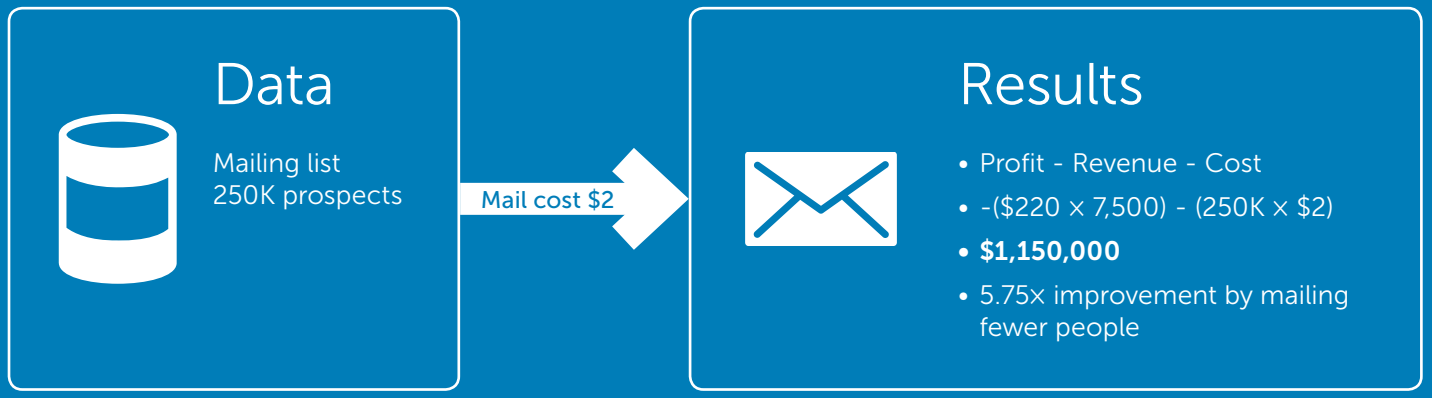

Adapted from Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie or Die by Eric Siegel

How will people in your organization react?

You'll know the answer to this question in almost no time, because word travels fast. At Dell, most people's reactions landed in one of three buckets:

1. "But we've never used Statistica."

   Since our users had been using SAS for many years, they weren't familiar with how robust an analytics platform Statistica is, so naturally they were skeptical. We pay them to be skeptical.

   They knew that their work consisted of mission-critical analytics in SAS and assumed (incorrectly, as it turned out) that Statistica wasn't up to it.

2. "We've spent years writing thousands of lines of SAS code. We don't want to just throw that away."

   Our users anticipated a great deal of work in trying to replicate in Statistica the programs they had built in SAS, so naturally they balked. Who wouldn't feel that way?

3. "We consider ourselves SAS professionals and analysts first, and employees of Dell second. For career longevity and our ability to do our job, we believe that it's really important to continue using SAS."

   That's a tough one. We found a number of heavy SAS users who had been working with the product for over 20 years. They were comfortable using it and they had grown, evolved and become pretty good with it over much of their career. Asking them to switch to something they didn't know was a huge disruption for them.

Most users had never heard of Statistica — let alone used it — so it was the devil they knew versus the devil they didn't know. Plus, many felt an emotional attachment to the tried and true product.

We addressed their reactions by having our migration leads from Statistica sit down and show them that their long years of work would not be simply discarded. The leads examined the techniques and functions our users had worked with in SAS, such as K-means clustering, polynomial regression, GLM, ARIMA and neural networks, then demonstrated how to replicate and enhance them in Statistica. Nearly all the techniques they had used in SAS were easier to implement in Statistica without the need to write thousands of lines of code. Users simply dragged and dropped icons onto the Statistica workspace in its easy-to-use graphical user interface.

You'll see that making the migration a success requires consistent communication, executive buy-in, technical support, training and encouragement, mostly from peers.

How will you get them on board?

Explaining your migration project is one thing; getting all your users on board with it is another.

First, all of our users and business units said they would need more than a year to migrate successfully to Statistica. Since the mandate was to migrate before December 31, we had to come up with several different approaches to make the project palatable to our users:

- **Making sure we had identified all current SAS users and properly communicating the change to them.**

  Nobody wants to find out about something like this at the water cooler. Our migration project team compiled a list of active SAS users and identified the affected executives and managers who would be responsible for keeping the users apprised of project status. We were mindful that many users would initially regard the change as a difficult one, so we kept the communications consistent, positive and frequent so they could digest them completely. We held several workshops early in the project to field and address users' concerns.

- **Giving people early access to Statistica.**

  It's easy to overlook this as a migration project starts up, but as the message went out and we talked to group managers about the timeline, they naturally replied, "We'll need access to the tool if you expect us to hit those dates." We had to show them and their users what they were in for and let them start working with Statistica. That kicked off the next set of conversations about setting up infrastructure, rolling out the product internally and paying for it.

- **Allaying the concerns of the dyed-in-the-wool SAS professionals.**

  Using a new tool does not strip those people of their SAS expertise or marketability, but they consider themselves part of a community that extends beyond Dell. We knew we couldn't address that on the emotional level, so we addressed it on the technical level by deeply discussing their analytics tasks and demonstrating how Statistica could satisfy them.

- **Putting SAS in its proper context.**

  Why use a Ferrari to deliver a load of dirt? We found that, for lack of a better (and lower-priced) tool, many people had grown accustomed to using SAS in cases where it wasn't the best fit and where it was overkill. They applied it to functions like data movement (ETL), data aggregation and preparation, which don't require advanced analytics technologies.

- **Guiding users away from inefficiencies and shadow IT.**

  By centralizing data movement tasks in officially supported environments and processes, we lowered internal risk and improved performance within Dell. In cases of ad hoc data movement, the migration leads showed users how to accomplish the same task with Statistica or with tools like Toad Data Point to get them on board with the migration.

> We knew that carrot works better than stick at Dell, so that's how we approached migration.

- **Launching a contest.**

  To inspire teams to take Statistica with both hands and apply it, we ran an internal contest. Teams of one to five people sprang up around the company, built analytic models to address real-world questions in Dell and competed against one another. The contest not only attracted a lot of attention to Statistica but also helped us solve business problems.

Every organization is a different ship, steered differently. You can try a mindset like "Between now and the end of the year we're switching to a tool you don't know and which you can't yet access, and in the meantime you have to do your day job," but we knew that carrot works better than stick at Dell, so that's how we approached migration.

Since this project was mandated by Michael Dell, it enjoyed the support of our executive leadership. Whenever the migration team encountered a substantial hurdle, its members got the support they needed. In general, we overcame most of our obstacles with additional training, support and one-on-one sessions customized to the teams that needed them.

In short: communication, communication, communication. Add executive support, cultivate the employees who are not afraid of the challenge, and persist.

Who gets to be an exception, and why? Lesson learned.

When you look back, you'll see where you had to change your original plan and how you had to change it for the people factor.

We learned plenty of lessons about how people deal with an extensive tool change in an organization. The bigger lesson, though, was about allowing for exceptions on a migration project.

By the end of the project, all but 16 of our users had moved completely off of SAS before the December 31 deadline. Most of Dell Financial Services migrated successfully, but the modeling team made the case that they needed more time to migrate their most sensitive functions and received an extension to continue using SAS for six more months. More important, they demonstrated that any interruption or fluctuation in their credit scoring and fraud models would introduce greater risk than it would be worth to meet the initial deadline.

That made sense. Some groups use analytics to predict, say, staffing levels for IT call centers. If their predictions are wrong, customers may have to wait longer, but the immediate impact is not as prominent as a hiccup in financial modeling. It was unfortunate that not every user was on Statistica by January 1; the downside, however, might have been even more unfortunate, so the project team worked with the modelers on a specific plan that extended the migration timeline for that small group.

Other than that, it speaks well to the quality of our employees (and of the product) that the transition to Statistica was smooth for more than 95 percent of them, especially considering that we had until December 31 to achieve what everyone thought would take more than a year to accomplish. We had to help some of them over their emotional obstacles to using a different tool in a different way, but what matters is arriving at the same result. Several teams performed their own comparisons and convinced themselves of that.

Learn more

Part 2 in this series covers process and the many things we put in place to make the migration project a success. Part 3 describes the technology component of the migration project, from architecture to onboarding users.

Visit software.dell.com/products/statistica/

For More Information

About Dell Software

Dell Software helps customers unlock greater potential through the power of technology—delivering scalable, affordable and simple-to-use solutions that simplify IT and mitigate risk. The Dell Software portfolio addresses five key areas of customer needs: data center and cloud management, information management, mobile workforce management, security and data protection. This software, when combined with Dell hardware and services, drives unmatched efficiency and productivity to accelerate business results. www.dellsoftware.com.

If you have any questions regarding your potential use of this material, contact:

Dell Software
5 Polaris Way
Aliso Viejo, CA 92656
www.dellsoftware.com

Refer to our Web site for regional and international office information.

Ebook-SAStoStatistica-Part1-US-GM-26380