WiNG

What Happens When Data Meets Cloud? Why Wing Invested In Snowflake Computing

JUNE 22, 2015 PETER WAGNER CLOUD

The cloud has become one of those irresistible forces that transforms everything it touches. This profound change happens sooner in some areas of the tech landscape, later in others. But now the DNA-altering influence of the cloud is taking hold in one of the furthest reaches of enterprise technology: the data warehouse.

One of our motivations in founding Wing is what we call “The DMC Thesis”, which I outlined in “The Wing Manifesto, Part II (http://wing.vc/perspectives/essays/the-wing-manifesto-part-ii-the-dmc-shift) .” Simply put, this is the observation that all aspects of business technology are being transformed by the reinforcing power of Data, Mobile and Cloud (DMC), and that this transformation is creating an unprecedented opportunity to build major new companies. What’s going on in data warehousing today is a perfect illustration of this dramatic change.

Data warehousing is a microcosm of everything that is right and everything that is wrong with enterprise technology. Right, in that data warehousing is a powerful tool built on deep IP that allows customers to become “data-driven” across their operation, harnessing a wide range of information to improve their business results. Wrong, in that the path to this data-driven utopia is still littered with roadblocks: seven, eight and even nine-figure costs; multi-year implementation schedules; armies of consultants; and, in the end, a rigid IT estate that is hard to operate and still can’t keep up with rapidly evolving user needs.

Some would argue that the Big Data movement will fix all this. And it does address many important issues, most obviously cost, scalability and flexibility with regard to both type of data and type of analysis. This has been a superb approach for a certain class of problems, in which the distributed computing paradigm is wheeled into service for a set of jobs that benefit from it. But

many more capabilities are required before all of the needs of enterprise data warehouse users are met. Smart people are adding some of these incrementally and the list of open source projects is endless. But enterprise data warehousing is only one of many sometimes divergent targets for the Big Data community. It is certainly not the easiest, and nailing it will require relentless focus.

The Snowflake team has brought that focus, and a very fresh approach, to reinventing the data warehouse. It had a unique starting point: founders Benoit Dageville, Thierry Cruanes and Marcin Zukowski had helped lead the prior generation of data warehousing developments. I was introduced to this trio by friend and co-investor Mike Speiser, and was immediately impressed by the clarity with which they had already envisioned the future of their market. They had a deep understanding of the existing approaches, and most importantly they had a keen appreciation of where these fell short. As the public cloud's capabilities expanded, they spotted the opportunity to use it to overcome the well-known roadblocks in data warehousing without sacrificing enterprise functionality—a trade-off that is still a fact of life in the Big Data camp. We were incredibly pleased to invest in their seed financing in 2013 and have been proud to continue to invest in every financing since then. Bob Muglia has been an inspired addition as CEO, and has helped lead Snowflake from seductive vision to commercial reality. The Snowflake Elastic Data Warehouse is now generally available, and is already beginning to reshape expectations for data in the cloud.

The first things one might hope for when imagining a cloud data warehouse are the fundamental properties of scalability and elasticity, and Snowflake certainly delivers these impressively. This alone would constitute a breakthrough. Just yesterday I heard a revealing comment from a top data infrastructure executive at a leading web-scale giant: "The data tier is actually anti-cloud," he told me, lamenting the barriers he faced to realizing true agility in this part of his service. Snowflake has broken these primary technical barriers, and gone further to build an end-to-end service that removes the need for data warehouse implementation and management. The customer sees massive workflow compression, including the elimination of traditional ETL (one of the most hated aspects of a conventional enterprise data warehouse project), and expansion of the data types that can be processed directly (e.g. JSON stored in S3).

Some of these capabilities sound a lot like what the industry is hoping to achieve with Big Data technologies. However, Snowflake has delivered them in true cloud fashion without compromising on key enterprise requirements, most importantly native SQL support and performance. This last point was particularly important to our investment thesis at Wing. The world seems to be filled with

SQL layers bolted atop various distributed computing frameworks. But an elastically scalable, on-demand data warehouse with the interactivity that users of SQL-based tools demand? We believe that to be a very rare and important thing indeed.

It is the cloud that enables Snowflake's distinctiveness. But simply porting an existing data warehouse to the cloud would never have been enough. The core technology had to be broken apart and recomposed—a re-architecture catalyzed by, and beautifully adapted to, the new cloud environment. At Wing, we consider Snowflake to be an archetype for future, cloud-borne infrastructure opportunities and we look forward to supporting other similarly exciting and visionary startups that take advantage of the DMC confluence to reinvent the technologies businesses depend on.

---

http://wing.vc/blog/what-happens-when-data-meets-cloud-why-wing-invested-in-snowflake-computing
© 1990 — 2018 Wing, Menlo Park, CA