

?Why Data-First Applications Will Come To Rule Enterprise Software

FEBRUARY 2, 2016 PETER WAGNER ESSAYS

Enterprise applications used to be mostly about business logic. Efficiency was the goal and automation was the means. Data and analytics were the province of separate systems and were a secondary priority. Looking ahead, we anticipate a new class of business applications rooted in data at their core. These will be fully enterprise-grade operational applications, but with a profound difference. Under the hood, they will be infused with sophisticated analytics that allow the business to be smarter and more agile, not just more efficient. From a design priority point of view, they are “data first”. Billions of dollars of IT spending is up for grabs here. We believe that in many of the major enterprise application categories, data-first attackers will emerge to challenge business-logic-first incumbents.

The Last Big Idea

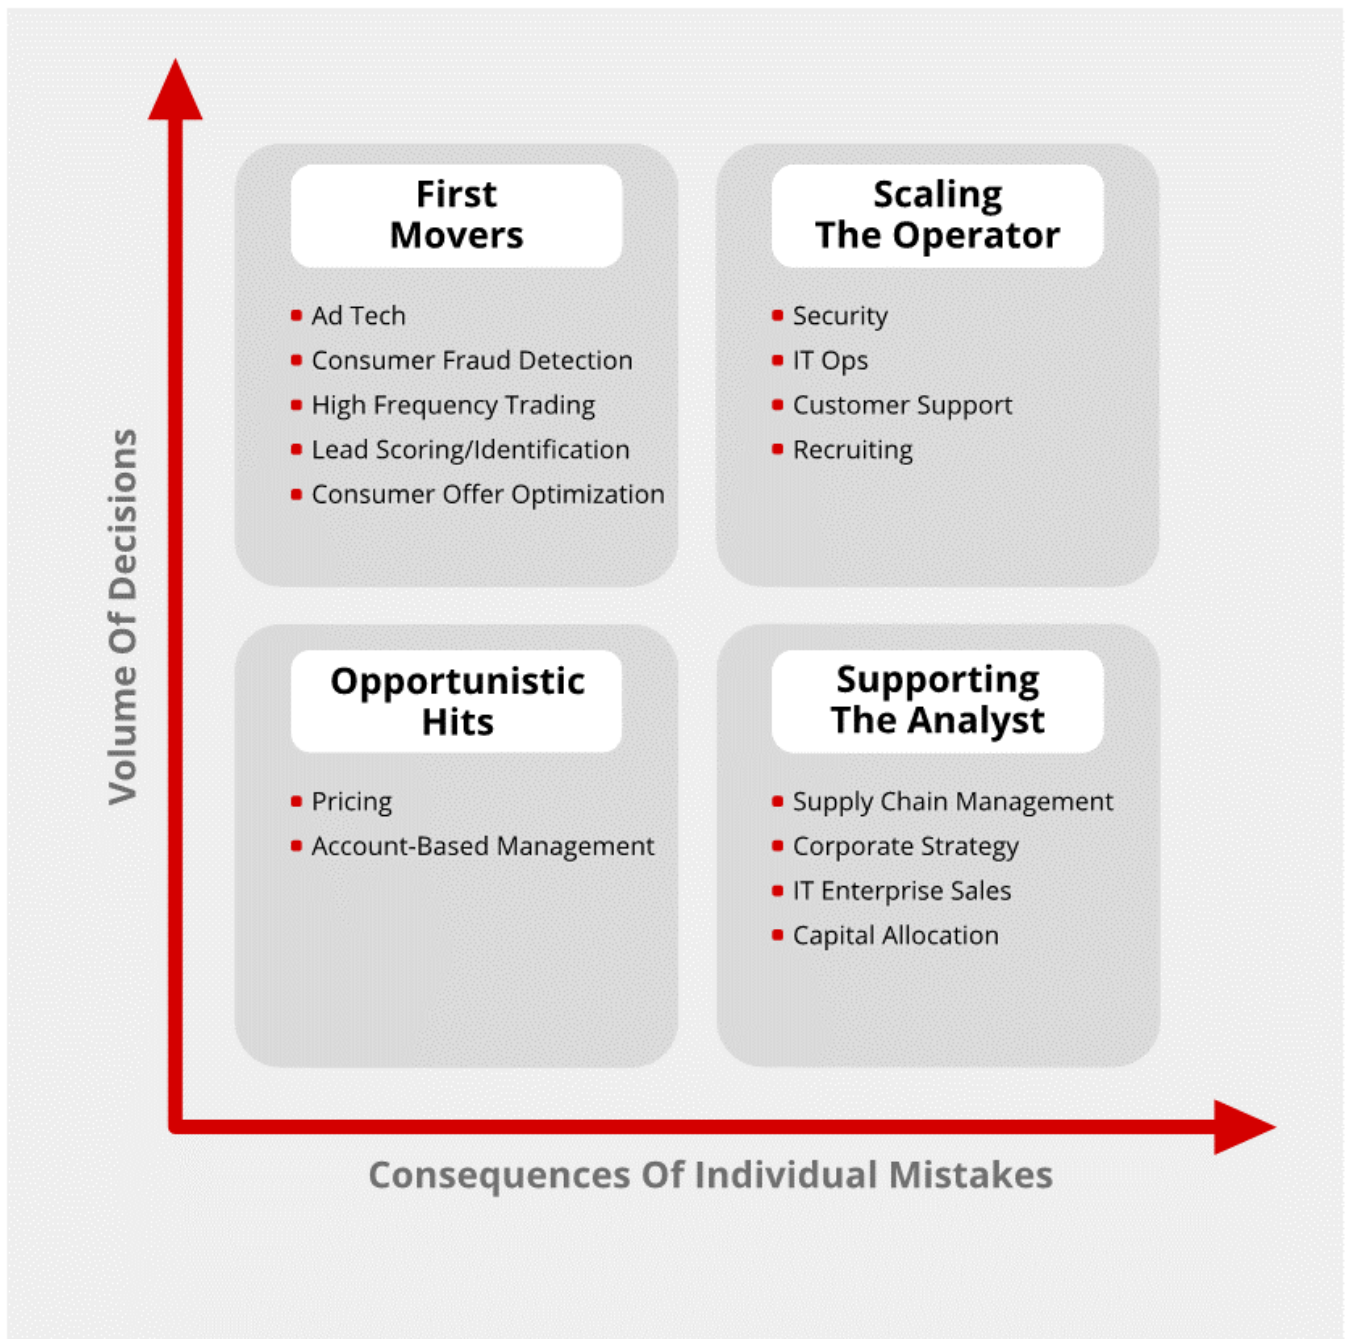
In the prior generation of applications, domain experts mapped out key business process workflows and software was written to codify them. Business logic was the value driver. This play was repeated with great success across a wide range of processes. Huge enterprise application software businesses (eg SAP and large parts of Oracle) were spawned from this and their legacy footprint remains a potent force.

The SaaS revolution addressed some of the sub-optimal aspects of enterprise application software, such as problems of shelf-ware (low utilization), onboarding (long time to value) and ossification (slow pace of innovation). However, while SaaS applications have shaken up many aspects of the industry, they have still been driven primarily by business logic. Their “Big Idea” was the move to the cloud. This improved the delivery and distribution model massively—huge gains to be sure—but it didn’t fundamentally address the value proposition at a functional level. As one top SaaS CEO explained to me recently: “The cloud idea turned out to be so big that we never got to the other ideas in our plan.”

The Age of Algorithms

Many of those neglected ideas had to do with data. The heart of the new, software-based economy will be algorithms embedded in next-generation, data-driven applications. Machine learning will play a central role, but the precise type of algorithm employed is not the key point. The inflection is driven by the fact that we will be increasingly trusting algorithms to run our businesses.

High-frequency trading and consumer fraud detection were early examples of this tectonic shift. Algorithmic control also took hold early in some key online business processes such as ad targeting and consumer recommendations (see chart below). These involved an incredibly high scale and velocity of interactions, which meant that no human could feasibly be in the loop. At the same time, the consequences of getting any one decision wrong were low, so it was good territory to trust the math without fear of bankrupting the company, killing the patient, crashing the plane or any number of other, high impact outcomes. In processes where the consequences of a wrong answer are much greater, we expect to see data-driven applications deployed in support of skilled operators wherever good data sets can be found and put to use. It's possible to imagine a sort of hierarchy of data-driven application functionality, ranging from alerts (the degenerate case), through clustering / classification, scoring, recommendations and even, in some cases, fully autonomous operation!




The Virtuous Data Cycle

Using data is good, but owning it is even better. This is the essence of the “data-driving” application, which is not just a data consumer but a data generator too. Data-driving applications are designed to achieve broad deployment and they create data exhaust that is ultimately the real reason for their existence. The application is an exquisitely crafted sensor and the vendor harvests a completely proprietary, “synthetic” data set to monetize. The value of key data sets will only

increase as algorithms proliferate because data-driven applications depend on data quality and volume for their effectiveness.

The result of all this is a “virtuous data cycle”. Applications generate data that are the critical input to additional domain-specific algorithms. These algorithms, in turn, form the core of next-generation business applications. These applications generate even more refined data, which feeds additional algorithms. All this raises the hugely exciting prospect of a virtual breeder reactor of business-process optimization (see second chart). The cycle isn’t new: consumer web companies such as Google and Facebook have been running this play for years. But what is new is the fact that this phenomenon is now invading the business markets. Let’s take a look at some key enterprise software categories to highlight early examples of data-first in the wild.



CRM

Many of the early examples of data-first applications have emerged in top-line-driving segments such as Sales and Marketing. This is to be expected—provably incremental revenue makes for the simplest and most compelling of ROIs. Last year’s Dreamforce trade show was filled with data talk, including from Salesforce itself with its Wave Analytics announcements. Meanwhile out on the tradeshow floor, a legion of sales and marketing analytics startups strove to differentiate their wares. There will be a lot of value created here—and probably a lot of incumbent market cap destroyed in the process.

A strong case can also be made that the data-first model will have the most value in industry-specific applications. Veeva is the canonical example in CRM. The company built its footprint—and its data set—with a standard CRM application for life sciences. It subsequently revealed data-first applications such as Veeva Network and OpenData. Veterans from Veeva and CRM pioneer Siebel Systems have now teamed up at Vlocity to execute a similar strategy in other verticals.

It is interesting to note that both Veeva and Vlocity are built on the Salesforce platform, in partnership with the SaaS giant. And it’s logical to ask whether this is something Salesforce should be doing by itself? The fact that it has chosen to partner in this domain speaks to the very large volume of opportunities before it, and the inability of any one company to execute well on all of them.

IT Operations

With lots of data, complex operations and highly technical users, IT is a natural place to look for signs of data-first applications. Moogsoft, a Wing portfolio company, is a great example of one of them. The company’s flagship incident-management application consumes data from various IT systems, applications and even external data sources; analyzes it using machine learning and other types of algorithms; and delivers an intelligent view of service-affecting situations as they unfold in real time. Some of the data-driven functionality includes clustering / classification of alerts and root-cause analysis at massive scale. Moogsoft also crosses over into data-generating territory with its “Situation Room” collaboration capability, which captures how each incident is resolved in order to build a historical data set of key people, symptoms and cures that is then used to derive

recommendations for actions during future incidents.

In our work as early stage investors, we are seeing an increasing number of entrepreneurs planning data-first attacks on other parts of the IT Operations arena. One interesting example is for the autonomous operation of large-scale networks. As amazing as that sounds, such an approach is already in production at the largest webscale giants. The opportunity is to productize this capability for the rest of the market.

Security

Cybersecurity is also blessed with ample data—too much in many cases. Today’s SIEM products are easily overwhelmed by the volume and diversity of data streaming towards them, and are not well equipped mathematically to detect today’s sophisticated threats. A new class of products from companies such as Securonix, Exabeam, Fortscale and Cybereason ingest existing data streams and employ big data analytics to identify anomalous behaviors and create some measure of potential risk in order to improve the effectiveness and productivity of security operations personnel.

Another group of companies distinguish themselves by bringing new data to bear. These include endpoint- and user-monitoring technologies, which then feed analytical systems for anomaly identification. Dtex Systems, another Wing portfolio company, is pioneering the field of user analytics, gathering detailed user-behavior data via a unique endpoint technology in order to help identify compromised accounts and insider threats.

Human Resources

At the other end of the spectrum lies HR, with a relatively low volume of data and far less technical users. Incumbent HRIS, Human Capital Management, Recruiting and Learning Management Systems are epitomes of business-logic-first applications. They miss opportunities to capture relevant data, make little use of the data they do have, and don’t tap much external data at all. Even relatively progressive new leaders in this arena like Workday are only now beginning to get serious about data-driven approaches. Workday deserves credit for realizing the need, but it will still be difficult for it to invert its architecture to the data-first model.

Meanwhile new entrants are championing a data-first approach. Google has been out in front and highlights key learnings from its own experience on its re:Work site. Vendors are emerging to

productize related concepts, including HiQ which uses data science and public information to identify flight risks in a company's employee base. Kanjoya has followed an unusual path to data-first, starting life as a unique social network called "The Experience Project". It has put this network to unexpected use as an incredible training data set for its algorithms. The company's applications are able to add rich qualitative analysis, including emotions and themes, to the classic "rate this 1 thru 5" employee survey, potentially reinventing the field of employee engagement.

Both companies are pursuing attractive entry points that may open up even larger swathes of the enterprise HR market. The products are particularly popular in the technology industry, where the battle for talent has risen to fever pitch.

Other enterprise application segments contain analogous opportunities. Public and internal data on products and components can be brought to bear on the bill of materials, breathing new life into supply chain and product lifecycle management. Customer support can be accelerated and improved with data as well, a phenomenon already in clear view in the form of Nimble Storage's InfoSight offering. The Internet of Things creates a path to bring this kind of intelligent, automated support to a far wider range of sectors, including consumer products of all kinds. The IoT tether enables both rich data capture as well as proactive contact with the user, and serves as a powerful accelerant of the virtuous data cycle.

First Data, Then Dominance

Are we seeing the beginning of an upheaval in enterprise application software? The data-first attackers will initially appear with very focused use cases and value propositions, and will integrate with the incumbent systems for legacy data access, as new data collectors, and as analytical co-processors. The first impression is all very complementary.

But this may prove to be just the first step in a wholesale transformation. Next-generation systems of record will be built around data at their core. The smarter incumbents are already trying to respond, but this could be an even more difficult transition for them than the leap to the cloud. Most won't make it. New data-first alternatives are already capturing beachheads that will allow them to someday topple empires.

<http://wing.vc/perspectives/essays/why-data-first-applications-will-come-to-rule-enterprise-software>

© 1990 — 2018 Wing, Menlo Park, CA