

March 2018

An Excerpt from

Time to Say Yes to NoSQL

Companies Need to Get Out of Their Comfort Zones and Adopt NoSQL Solutions—or Get Comfortable with a Limited Future

Stratecast Analysis by
Jeff Cotrupe



Big Data and Analytics (BDA)

Excerpt from
Volume 6, Number 1

An Excerpt from

Time to Say Yes to NoSQL

Companies Need to Get Out of Their Comfort Zones and Adopt NoSQL Solutions—or Get Comfortable with a Limited Future

Table of Contents

Executive Summary	3
Introduction	4
The Case for NoSQL	4
Barriers to Adoption of NoSQL	6
Legacy Limitations and NoSQL Improvements May Spur Adoption	7
<i>Fear and Opportunity Have Enterprises Seeking a Better Way</i>	<i>7</i>
<i>“NewSQL” Solutions Allay Old Fears</i>	<i>8</i>
Options a-Plenty; Market Leaders	8
Approaches to Reaping the Benefits of NoSQL (Couchbase)	9
<i>Couchbase Data Platform</i>	<i>9</i>
Recommendations	11
<i>For Buyers of NoSQL Solutions</i>	<i>11</i>
<i>For Providers of NoSQL Solutions</i>	<i>12</i>
Fill Out Your Portfolio	12
“Be Consistent”	12
Work Together to Educate the Market	12
The Last Word	13

List of Figures

Figure 1: Barriers to Adoption of NoSQL	6
Figure 2: Legacy DBs are Erecting Barriers to Business Success	7
Figure 3: The Most Widely Deployed NoSQL Solutions	9
Figure 4: Couchbase Data Platform	9
Figure 5: Couchbase Customers	10

An Excerpt from

Time to Say Yes to NoSQL

Companies Need to Get Out of Their Comfort Zones and Adopt NoSQL Solutions—or Get Comfortable with a Limited Future

Executive Summary

This report is the first in a two-part series on NoSQL databases.¹ It is designed to be meaningful and useful for anyone engaged in exploring, selecting, specifying, deploying, and managing solutions to equip organizations with data management strategies and solutions. The report's key themes and findings are as follows.

1. Relational databases are relics from an earlier era that are obsolete when it comes to managing the vast majority of data being created today. The world is awash in unstructured data, with exabytes more being created every day; and relational DBs can neither effectively manage the new types of data nor keep up with the sheer magnitude of it. A move to databases that can manage unstructured data—NoSQL DBs—would appear to be in order.
2. Asserting a statement along the lines of the first bullet, in a report in 2018, sounds ridiculous on its face when one considers that awareness of the changing face of data has been growing for nearly a decade. However, the vast majority of organizations (94%) continue to plod gamely forward with relational DBs at the core of their data management operations.
3. Some decision makers simply don't get it. Others do, but a combination of fear factors is holding them back. Those factors include familiarity with existing systems, and a fear of deploying new ones; fear that trained experts to run the new systems may be in short supply; and the inability of large segments of the NoSQL community to deliver systems that offer the same data consistency guarantees as relational DBs.
4. These barriers to adoption are beginning to crumble. The limitations inherent in legacy DBs are beginning to cost companies money, opportunity, and customer goodwill—across all verticals. At the same time, vendor implementations of NoSQL are beginning to allay old fears. Some NoSQL providers offer data consistency guarantees on par with relational DBs, enabling companies to avail themselves of the blazing speed and unparalleled flexibility of NoSQL without the bitter aftertaste of uneven consistency. Other providers are not quite there, but close enough to enable enterprises to cross the legacy divide.
5. NoSQL providers are visionaries who are doing an outstanding job of hitting the growing and moving target that is data, in all its forms and venues. However, as a group, they are also offering such a wide array of capable options—and competing so effectively with each other—that they are creating confusion in the market. We recommend that providers in the space give buyers more opportunities to obtain all NoSQL capabilities under one roof; and come together to educate the market.

¹ The companion piece is Stratecast, *Helping Enterprises Say Yes to NoSQL: NoSQL Databases, Providers, and (Open) Sources* (BDA 6-02, March 2018), available [here](#).

Introduction²

The first relational databases appeared in the early 1970s.³ They worked fine then, and still do today, for data easily organized into rows in a table: structured data. There is, however, a problem with relational DBs. The world is creating an estimated 2.5 exabytes (2.5 quintillion bytes) of data per day,⁴ and relational DBs cannot manage most of it. Fully 90% of the data that exists today has been created in the last two years. Most of it emanates from mobile sources and the Web; from corporate documents and human-readable XML files used in various systems; in sensors that reside in retail stores and on manufacturing assembly lines; from chatbots and other implementations of AI; and many other sources. The data these things are producing is not structured data, which relational DBs can effectively store and process; it is unstructured data. The rapid proliferation of the IoT is only going to accelerate this effect: more than 12 billion devices were connected to the IoT in 2016, and that will grow to more than 45 billion devices by 2020.⁵

The fact that the amount and types of data have outstripped the ability of relational DBs to handle the onslaught “is not their fault”: relational DBs were conceived and developed before the existence of the Internet. The lack of a way to handle unstructured data led to landmark developments such as Google’s MapReduce and BigTable databases in 2004-06, and Amazon’s creation of DynamoDB in 2007: the first NoSQL (“not only SQL”) databases. In the years that have followed, a number of companies have built NoSQL DBs, either directly on or modeled after this pioneering work.

Readers may be asking, understandably, “So what? Most are aware of the changing face of data.” Stratecast asserts that the “so what” is that many organizations, in 2018, still don’t get it. Or, they do, but are clinging to their legacy relational DB-driven data infrastructures as the foundation of their big data and analytics (BDA) strategies. In the process, they are leaving a lot of potential analytic insights untapped—insights that could hold the keys to their survival and prosperity.

This report discusses the primary benefits of NoSQL; barriers to adoption of NoSQL; the ready availability of NoSQL solutions, as exemplified by some of the top providers thereof; and competitive dynamics at work in the space. The report then makes specific recommendations for enterprises and vendors deploying NoSQL.

The Case for NoSQL

Organizations spend more than \$40 billion per year on databases,⁶ and incrementally more on the resources required to manage them. Infrastructure, applications, and the data itself have undergone massive changes for decades, but the underlying data management tool, the relational DB, has remained the same. In many cases, organizations are not aware of how their relational DBs are holding them back. The following technology-based factors illustrate why relational (SQL) DBs are

² In preparing this report, Stratecast conducted interviews with:

- Couchbase – Christina Knittel, Director Corporate Communications

Please note that the insights and opinions expressed in this assessment are those of Stratecast, and have been developed through the Stratecast research and analysis process. These expressed insights and opinions do not necessarily reflect the views of the company executives interviewed.

³ The relational model was first described in 1969 by Edgar F. Codd, a UK computer scientist, while working for IBM.

⁴ IBM, *10 Key Marketing Trends for 2017*, available [here](#)

⁵ Frost & Sullivan’s IoT Universe practice, *Internet of Things 2.0: Predictive Intelligence*, available [here](#)

⁶ Forbes, *MongoDB Taking Share From Oracle In \$40 Billion Market*, available [here](#)

obsolete for many business use cases, and why NoSQL is an essential ingredient in BDA deployments today:

- The majority of the data that organizations must now deal with is unstructured and semi-structured: Web content; multimedia (video and audio); email and social platform messaging; business productivity/word processing documents; books, magazines, images, and electronic objects; XML files; and more. **NoSQL DBs were built from the ground up to manage unstructured and semi-structured data.**
- Organizations today need to perform fewer data queries that are hierarchical (relational), and more queries that generate real-time insights and recommendations. Today's queries demand that users be able to find the information they need without regard to a question's syntax aligned to a set of predefined fields. Users need to be able to generate real-time analytics and business recommendations that draw subjective conclusions. **NoSQL DBs are well-suited to generating real-time insights and recommendations.**
- Relational data doesn't map well to typical programming structures that often consist of complex data types or hierarchical data. Complex objects rarely map well to a single row in a table. Key-value stores, wide column stores, graph databases, and other NoSQL DBs accommodate flexible data models with no hierarchical mapping required. **NoSQL offers the ability to concurrently obtain multiple types of analytics insights.**
- Pure algorithmic reasoning takes a long time to process large databases. For example, "SQL joins" combine columns from one or more tables in a relational DB by using values common to each; this is how companies try to map unstructured data into a relational DB. SQL joins can slow systems to a crawl, especially when a high volume of users is doing lookups against tables with millions of rows of data. **NoSQL DBs have a variety of performance-enhancing qualities, one of which is that they require no SQL joins to synergize data of various types.**
- In a relational DB, one must "measure the furniture before building the room." To use a relational DB, data scientists and IT administrators must define a schema (structure) before adding records to the DB. To add attributes to a data object in a relational DB, one must modify the schema to add columns and data types. By contrast, a NoSQL DB has metadata (records that describe the data) and the data itself. The system can use nested data to provide additional sub-categories of information about an object. To add attributes, one can simply add key-value pairs to represent the new fields. **A NoSQL DB makes it possible to add new types of data "on the fly": at a moment's notice and with little pre-planning.**
- As a result of some of the factors above, relational DBs suffer from various issues including performance constraints, loss of uptime, and inability to scale to the massive volumes of data that companies face today. By virtue of their design, **NoSQL DBs can scale to accommodate huge quantities of data, with good performance.**

Barriers to Adoption of NoSQL

With all that NoSQL would appear to have going for it, why hasn't every enterprise jumped on the NoSQL bandwagon? The reasons have to do partly with technology and partly with people factors. One area in which all relational DBs have traditionally outperformed NoSQL DBs—and still do, depending on which NoSQL DB one is assessing—is ACID compliance. ACID refers to four characteristics associated with relational DBs:

- **Atomic:** Everything in a transaction succeeds or the entire transaction is rolled back, and it is as if the transaction was never attempted.
- **Consistent:** No report can “catch data between states,” i.e., while data is being updated. The state of being consistent is sometimes termed **full** consistency.
- **Isolated:** Users cannot touch incomplete data (between states, not checked or matched).
- **Durable:** In the case of a partial DB system failure, the system, as a matter of course, recovers committed transactions (those in process at the moment of failure).

The most common point of comparison between DBs is consistency. Some NoSQL DBs offer **eventual** consistency: they guarantee consistency on individual data objects, such as a specific transaction record; as the DB catches up with all data changes system-wide, the value for a specific piece of data will, given enough time without updates, be consistent across all nodes. All relational DBs, however, offer full consistency. Barriers to NoSQL adoption appear in Figure 1.

Figure 1: Barriers to Adoption of NoSQL

Barrier	Description / Implications
The ACID Test: Consistency	<ul style="list-style-type: none"> • Forced to choose between vastly superior speed and data flexibility of NoSQL DBs versus the guaranteed consistency of relational DBs, many teams choose consistency. • In certain cases in some industries, full consistency has been mandated by regulations, such that NoSQL has been excluded by default.
Comfort / Inertia	Enterprises are comfortable with legacy relational DBs that have existed for decades.
Hiring and Training	<ul style="list-style-type: none"> • Enterprises are not comfortable taking the risk to adopt a new technology without having people in place who are hired and trained to support it. • Many IT leaders are reluctant to re-train existing staff—or to hire new staff with NoSQL expertise—so, they continue to cling to SQL and try to make it work.
Stubbornness	<ul style="list-style-type: none"> • Many companies prefer to use what they know, continuing to shoehorn SQL systems into their environment rather than truly evolving their approach to data. • SQL/RDB adherents may simply refuse to change.
Uneven Organizational Visions	<ul style="list-style-type: none"> • Some lines of business are more forward-thinking, other more set in their (legacy) ways. • Lack of consensus paralyzes organizations into “decision by indecision.”
Inflexible Development Cycles	<ul style="list-style-type: none"> • Natural development cycles of many companies force a new technology to break into the existing “flight plan.” • If outside “flight plan,” it may require re-prioritizing of schedules and rework of business apps—both of which increase organizational resistance to adoption.

Source: Stratecast

Beyond questions of consistency, IT and data people are like people in general in that they can be resistant to change, and prefer to use what they know rather than take a chance on something new. As a result, they continue to shoehorn SQL systems into their environments rather than commit to a

true evolution to the way they access and use data. Companies will exaggerate the capabilities at the foundation of their data management strategies, but sidestep the reality that those functionalities are bolted onto the legacy relational DB technology that is the foundation of the system. Organizations that have deployed Apache Hadoop miss the fact that it is still a relational DB-style, master-slave architecture, with multiple points of failure possible and other limitations.

Legacy Limitations and NoSQL Improvements May Spur Adoption

As a fast-moving world translates the shortcomings of relational DBs into business pain, platform improvements have NoSQL poised to fill the need.

Fear and Opportunity Have Enterprises Seeking a Better Way

A key business driver for the adoption of NoSQL solutions is fear. Retailers, for example, are trying to figure out how to transform their businesses to match the expectations in shoppers' minds that have been set by Amazon. Retailers such as Macy's are closing stores;⁷ and need to make use of all available data to get more interactive with shoppers, and make purchase recommendations.

Retail aside, in every vertical, companies are anxious about digital transformation. Decision makers are worrying about what the next disruptive force will be that sets an industry on its ear—à la Uber or AirBNB, or perhaps Amazon in healthcare—and smarter companies are choosing to fight back by harnessing the power of all data streams with NoSQL. Conversely, success stories such as Amazon, AirBNB, Netflix, Uber, Facebook, Google, and many more would be impossible without the ability to access all relevant data from all sources for real-time analytic insights.

As such, while 94% of companies are still using relational DBs from Oracle, Teradata, or Microsoft,⁸ given the negative results many are experiencing, as shown in Figure 2, that is certain to change.

Figure 2: Legacy DBs are Erecting Barriers to Business Success

Event or Characteristic	%	Description
Impacted Reliability	87	Maintenance or modifications to organization's legacy DB cause application downtime, making it harder to guarantee customer experience
Impacted Agility	86	Organization's legacy DB has impacted its ability to develop new applications and improve business operations and/or agility
Impacted Performance	74	Using legacy DB to support applications that engage with end users via increasingly complex/interconnected/varied data: performance of these applications has suffered
Lack of Scalability	61	Organization's legacy infrastructure hinders ability to scale applications up/down—so, if customer base surges, it becomes progressively harder to engage with customers
Digital Project Failure	41	Have had digital projects fail outright because legacy DB could not support them
Reduced Project Scope	29	Have had to reduce the scope of a project due to the cost of making changes to legacy technology
Significant Project Delays	14	Have incurred significant delays due to legacy technology

Sources: Couchbase and Vanson Bourne

⁷ CNBC, "Here are the stores Macy's is closing next," available [here](#)

⁸ Couchbase and Vanson Bourne, *Digital Innovation Survey*, available [here](#)

Only 41% of organizations say they can use data in real time; on average, the most recent customer data that organizations' databases can use is more than one day old—meaning that a truly engaging, real-time customer experience is still a long way away for many.

Legacy DBs are also holding companies back when it comes to harnessing the power of new technologies such as artificial intelligence (AI), virtual and augmented reality, and the IoT. Only 19% of decision makers believe their DB would be up to the task of fully supporting such technologies if their organization began using them tomorrow. And 80% of respondents were already concerned that their organization could be at risk of being left behind by competitors.

19%

**Organizations whose
current DB can support
AI, VR/AR, and IoT**

“NewSQL” Solutions Allay Old Fears

NoSQL DBs were created to handle massive quantities of data that relational DBs could not; so, by definition, they are nearly always deployed as large-scale, distributed systems. Conventional wisdom in the data community has been that in such systems data inconsistency must be tolerated due to the CAP theorem, which states that they can have, at most, two of three characteristics: consistency (C); high availability (A), or high-speed performance; and partition tolerance (P). Partition tolerance, however, most often refers to a network failure; so, in reality, under most circumstances, no tradeoff must be made. As subsequent sections show, some providers now offer full consistency, on par with relational DBs. A few offer eventual consistency. Some offer strong consistency, which is defined as full consistency under some circumstances and eventual consistency under others. Another driver for NoSQL adoption speaks to both technical mastery and business agility: companies want all of their data available for read and write capability by authorized users, anytime, anywhere. A NoSQL-based system makes it possible to add new types of data on the fly—at a moment's notice, and with little pre-planning—to respond to changing business needs. This enhances organizational agility, and slashes time to market for new offerings.

Options a-Plenty; Market Leaders

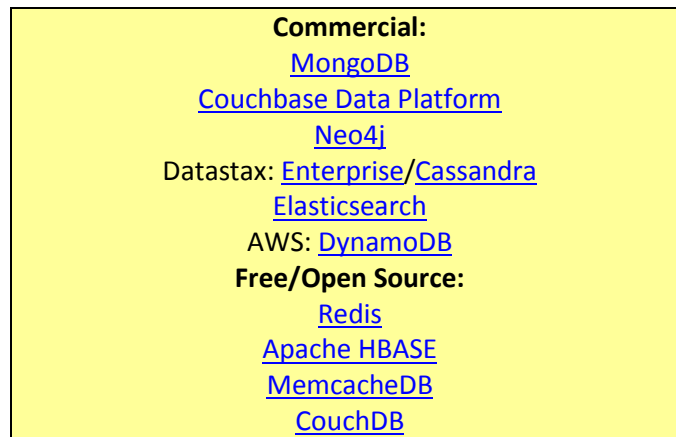
Those who want to break through barriers to NoSQL adoption have a wide array of options from which to choose. We analyze those in detail in the second part of this two-report series;⁹ but a brief review of the types of NoSQL DBs available is sufficient to illustrate. There are 12: Column/Wide Column Store DBs; Document Store DBs; Graph Store DBs; RDF Triplestore DBs; Grid Data Store DBs; Key-value DBs; Tuple Store DBs; Multi-model DBs; Multi-dimensional DBs; Multi-value DBs; Network Model DBs; and Object DBs.

As shown below in Figure 3, the 10 most widely deployed NoSQL solutions are, on the commercial side: MongoDB; Couchbase Data Platform; Neo4j; DataStax Enterprise/Cassandra; Elastic; and AWS DynamoDB. Free/open source technologies are: Redis; HBASE; Memcached; and CouchDB. This lead group accounts for 86% of total deployments, with another 125 sources accounting for the remaining 14% of deployments.¹⁰

⁹ The companion piece is: Stratecast, *Helping Enterprises Say Yes to NoSQL: NoSQL Databases, Providers, and (Open) Sources* (BDA 6-02, March 2018), available [here](#).

¹⁰ In approximate order, based on Stratecast's research and the best available confirming data from other sources

Figure 3: The Most Widely Deployed NoSQL Solutions



Source: Stratecast


Approaches to Reaping the Benefits of NoSQL (Couchbase)

The following section briefly presents Stratecast's assessment of Couchbase, one of the top commercially available solutions.

Couchbase Data Platform

Couchbase Data Platform, shown in Figure 4, is an enterprise NoSQL platform that incorporates both a JSON **document store** and a **key-value store**. Optimized for interactive applications, it guarantees high performance and strong consistency with a built-in, object-level cache, asynchronous replication, and data persistence.

Figure 4: Couchbase Data Platform



Source: Couchbase

Couchbase Data Platform is designed to independently scale out or scale up compute-, RAM-, and storage-intensive workloads to support concurrent users creating, storing, retrieving, aggregating, manipulating, and presenting data. This offers maximum capacity and flexibility as business needs grow and change. To address these demanding requirements, Couchbase Data Platform provides easy-to-scale key-value or document access with low latency and high sustained throughput, and is designed to be clustered from a single machine through very large-scale enterprise deployments.

The best way to illustrate the viability and relevance of NoSQL, or any technology, is to determine who is using it and what benefits it delivers. To illustrate the first part of this, Figure 5 shows a representative sampling of Couchbase customers.

Figure 5: Couchbase Customers



eBay, BD, Betfair, Cars.com, Amadeus, KDDI, Willis Towers Watson, Nielsen, Cox Enterprises, Centeredge, Seenit, GE, Matrix, Shop.com, Concur, United Airlines, Mirror Image, Cisco, Ryanair, Cvent, Sky, Tommy Hilfiger, Doodle, Viber, SyncThink, TransitWorks, Verizon, Nuance, Marriott, LinkedIn, SafetyCulture, Gannett, McGraw-Hill Education, DirectTV

Sources: Couchbase and the companies

Three brief snapshots illustrate the types of value Couchbase is delivering for customers:

- **Amadeus, the world's largest travel booking engine, replaced Oracle with Couchbase,** and now manages its high look-to-book ratio, including more than eight million queries per second, and more than 20 billion documents, with Couchbase Data Platform. By adopting NoSQL, Amadeus pushed the complexity of data distribution down to the database level, and now uses Couchbase across seven applications that combine to process more than 2.5 million operations per second on petabytes of data.
- **Cisco, after assessing other solutions, including Cassandra and MongoDB, chose Couchbase** primarily for four reasons: strong data consistency; reliable low latency, with consistent 500 microsecond response times at very large scale; scalability in a distributed system; and its “Swiss army knife-like” clusters, which let Cisco manage and maintain many aspects of its database with one solution.
- **Nielsen turned to Couchbase to eliminate the limitations of Oracle, and improved response times by 50%.** Nielsen's Answers on Demand (AOD) service delivers ratings data and other information for businesses in more than 100 countries. Faced with the challenge of tracking sales of fast-moving consumer goods, and reporting those insights to customers, Nielsen needed a backend solution that could store user-generated data while providing extremely fast response times and low latency. Not only has Nielsen slashed response times in half by using Couchbase, it has also attained a more detailed and granular understanding of customer buying patterns and behavior, while also reducing the time to manage and update systems. Nielsen uses Couchbase Data Platform as a document store, and Couchbase's N1QL ‘SQL for JSON’ query engine to crunch big data for analytic insights.

Recommendations

The following sections of the report provide Stratecast's specific recommendations to both prospective buyers and active providers of NoSQL-based solutions—not with regard to a choice of vendors, but regarding some big-picture concepts that are far more important to both groups.

For Buyers of NoSQL Solutions

A brief discussion is in order regarding what may appear to be a baffling array of NoSQL technologies and solutions available today. Indeed, despite the massive limitations and shortcomings that relational DBs present, it may be understandable why some organizations, faced with too many choices on the NoSQL side, make one simple choice instead: throw up their hands and keep their organizations locked into their existing Oracle, Teradata, and other legacy relational DB-based data management systems for another year.

Here is why no one should do that for another year. Or even another month.

Relational DBs have their place today, over a limited range of use cases where data is easily organized. The world, however, has changed. Data, the systems that contain it, and its impacts, are ubiquitous. The world is exploding with new markets and business models, and NoSQL is addressing the new reality. The NoSQL movement is about far more than just “new vendors looking to take market share away from Oracle and TeraData relational DBs.” It is about a new breed of providers creating new ways to apply the power of big data and analytic insights to every aspect of the organization; and beyond that, to virtually every corner of the world—as privacy advocates are quick to point out.¹¹ This is the reason for the vast array of choices in the NoSQL arena: rather than force-fit this brave new world of opportunity into the old relational model, developers of NoSQL solutions are the only ones with the vision—and the intellectual honesty—to truly take on the changing face of data, and to build DBs to handle it.

Rather than force-fit this brave new world of opportunity into the old relational model, developers of NoSQL solutions are the only ones with the vision—and the intellectual honesty—to truly take on the changing face of data, and to build DBs to handle it.

The NoSQL movement is coming to grips with the importance of graph DB capabilities to break through data siloes and help humans and machines readily find, and leverage, the relationships between data points and objects. There is more competitive heat in the graph DB space than in any other area of the NoSQL landscape—including the fact that the leading graph DB provider, Neo4j, does not consider itself a NoSQL provider. In fact, the rapid rise of Neo4j and other graph technologies may signal that data connectedness is indeed a separate paradigm from the model consolidation happening across the rest of the NoSQL landscape.

For these reasons and more, at least at this juncture, there is no one-size-fits-all NoSQL solution. It is important to be aware of the comparative strengths and weaknesses of the various NoSQL options available. As such, before deciding which NoSQL DB to use in production, we recommend that buyers do two things:

¹¹ For analysis of the global impacts of data ubiquity on privacy, we recommend: Stratecast, *We Have Seen the Future of IT, and it is Big Data: Part 1 – Will IoT Privacy Issues Steal the Future?* (BDA 5-01, June 2017), available [here](#); *We Have Seen the Future of IT, and it is Big Data: Part 2: A Blueprint for Privacy, in the IoT and Everywhere* (BDA 5-02, July 2017), available [here](#); *Big Data is in Big Trouble, Starting in the EU: Our Assessment of the GDPR* (BDA 5-03, August 2017), available [here](#); and *Our Detailed Privacy Blueprint: What All Parties Should Be Doing Right Now to Protect the People and Organizations They Care About* (SPIE 2017-28, 11 August 2017), available [here](#).

- Get firm answers from providers as to the consistency guarantees of their solutions.
- Compare solutions in their own environments, using data and user interactions representative of actual production workloads.

For Providers of NoSQL Solutions

Readers are undoubtedly familiar with the phrase “divide and conquer.” We assert that in the NoSQL space, a more fitting phrase would be: remain divided and, ultimately, be conquered. Right now, the major providers in this space are pursuing separate paths, and taking highly effective competitive shots at each other. Meanwhile, the providers of legacy DBs, who have the most to lose from the NoSQL movement, seem content to sit back, “let the kids fight it out amongst themselves,” and tighten their stranglehold on the database market.

Fill Out Your Portfolio

To be clear: competition is the lifeblood of healthy markets and economies. For example, we contend that the competition and controversy around graph is healthy for the market. As our analysis of the eight categories of NoSQL DBs indicates, each brings different capabilities and characteristics to the table. Couchbase already goes to market with a two-pronged approach, combining both a key-value store and a document store in its platform. As Couchbase and other providers cover more of the NoSQL bases in their platforms, it will help buyers immensely by taking a lot of the guesswork out of choosing a solution, and eventually will prod even the most reluctant of suitors to the altar.

“Be Consistent”

Many buyers may indeed be set in their ways, and hesitant to leave their comfort zones, but there is no glossing over the data consistency questions associated with NoSQL DBs. Some providers offer full consistency or even full ACID compliance; and their very large and growing customer bases reflect, at least in part, the advantage that gives these providers when asking enterprises to entrust them with the lifeblood of their companies—their data. Every NoSQL provider should be striving to reach full consistency, and the sooner the better.

Work Together to Educate the Market

Stratecast asserts, however, that, since the main barrier to adoption of NoSQL-based solutions seems to be a fear of the unknown coupled with a bewildering array of choices, the formation of a NoSQL industry education and standards cooperative is in order. The providers in this segment are industry visionaries. They should find a way to work together to educate the market about NoSQL-based solutions, so that all innovative boats can rise together.¹²

Stratecast sees this as far more than beneficial. We predict that, without some sort of cooperative effort between NoSQL providers, we will begin seeing failures of some, while some of the more capable providers are snapped up through acquisitions by the legacy mega-vendors. The latter development would provide some of these NoSQL providers with a healthy infusion of cash to fund further development—but it could also constrain or even short-circuit their efforts to get the industry to collectively take an important step forward.

¹² Stratecast is willing, in fact, quite keen, to help to lead this type of industry initiative.

Stratecast The Last Word

A strong case can be made for why organizations should adopt NoSQL. About 90% of the data that exists today was created in the last two years, and relational DBs cannot effectively manage most of it. Nevertheless, nearly all organizations still cling to their legacy relational DBs, collectively paying tens of billions of dollars each year for the privilege—while failing to attain the uptime and scale they need to survive and prosper in today’s business environment.

The reasons they cling can be attributed to both humanity and technology. Reason #1, and maybe the only one that counts, is fear. Relational DBs have been entrenched in organizations for decades, and the idea of changing something so vital to the organization can be a frightening prospect. The reality that so many members of the global IT community are expert on existing systems, and fewer are trained and ready to manage NoSQL DBs, heightens the risk. Varying levels of data consistency among the various NoSQL providers have added to the concern.

Humans, however, tend to respond most readily to either pleasure or pain—and the limitations of legacy systems are now causing sufficient pain to spur organizations to action. From the retailer who lost big on Cyber Monday when its legacy DB choked on heavy shopper traffic; to the financial services company whose legacy system was not only costing it upsell opportunities but hurting customer experience; these and countless others across every vertical have had enough and are making the move to NoSQL.

The NoSQL movement itself, however, is unwittingly erecting another barrier to adoption for many: market confusion due to too many conflicting choices. This report mentions seven different types of NoSQL DBs, and some of those combined legitimate sub-categories. Each has important characteristics and benefits to offer. Yet, of the leading commercial providers in the space, most are based on one type of DB, and are quick to cast aspersions at other models.

The market is not interested in Holy Wars over database architectures; it wants to cover all of its (data) bases, preferably under one roof. To that end, we call on NoSQL providers to fill out their portfolios. The gravitational pull of Neo4j and others has caused competitors to do so in one area, graph, but this activity needs to expand across other areas of the NoSQL landscape. Providers who do not already offer full data consistency should be working toward achieving it soon—or work a great deal harder to educate the market about what they contend are acceptable tradeoffs between availability and consistency. Providers should also hammer their swords into plowshares long enough to work together to educate enterprises on the bankable benefits of NoSQL—and we are eager to help lead that effort.

Jeff Cotrupe

Industry Director – Big Data and Analytics

Stratecast | Frost & Sullivan

jeff.cotrupe@frost.com

About Stratecast

Stratecast collaborates with our clients to reach smart business decisions in the rapidly evolving and hyper-competitive Information and Communications Technology markets. Leveraging a mix of action-oriented subscription research and customized consulting engagements, Stratecast delivers knowledge and perspective that is only attainable through years of real-world experience in an industry where customers are collaborators; today's partners are tomorrow's competitors; and agility and innovation are essential elements for success. Contact your Stratecast Account Executive to engage our experience to assist you in attaining your growth objectives.

About Frost & Sullivan

Frost & Sullivan, the Growth Partnership Company, works in collaboration with clients to leverage visionary innovation that addresses the global challenges and related growth opportunities that will make or break today's market participants. For more than 50 years, we have been developing growth strategies for the Global 1000, emerging businesses, the public sector and the investment community. Is your organization prepared for the next profound wave of industry convergence, disruptive technologies, increasing competitive intensity, Mega Trends, breakthrough best practices, changing customer dynamics and emerging economies? For more information about Frost & Sullivan's Growth Partnership Services, visit <http://www.frost.com>.

CONTACT US

For more information, visit www.stratecast.com, dial 877-463-7678, or email inquiries@stratecast.com.