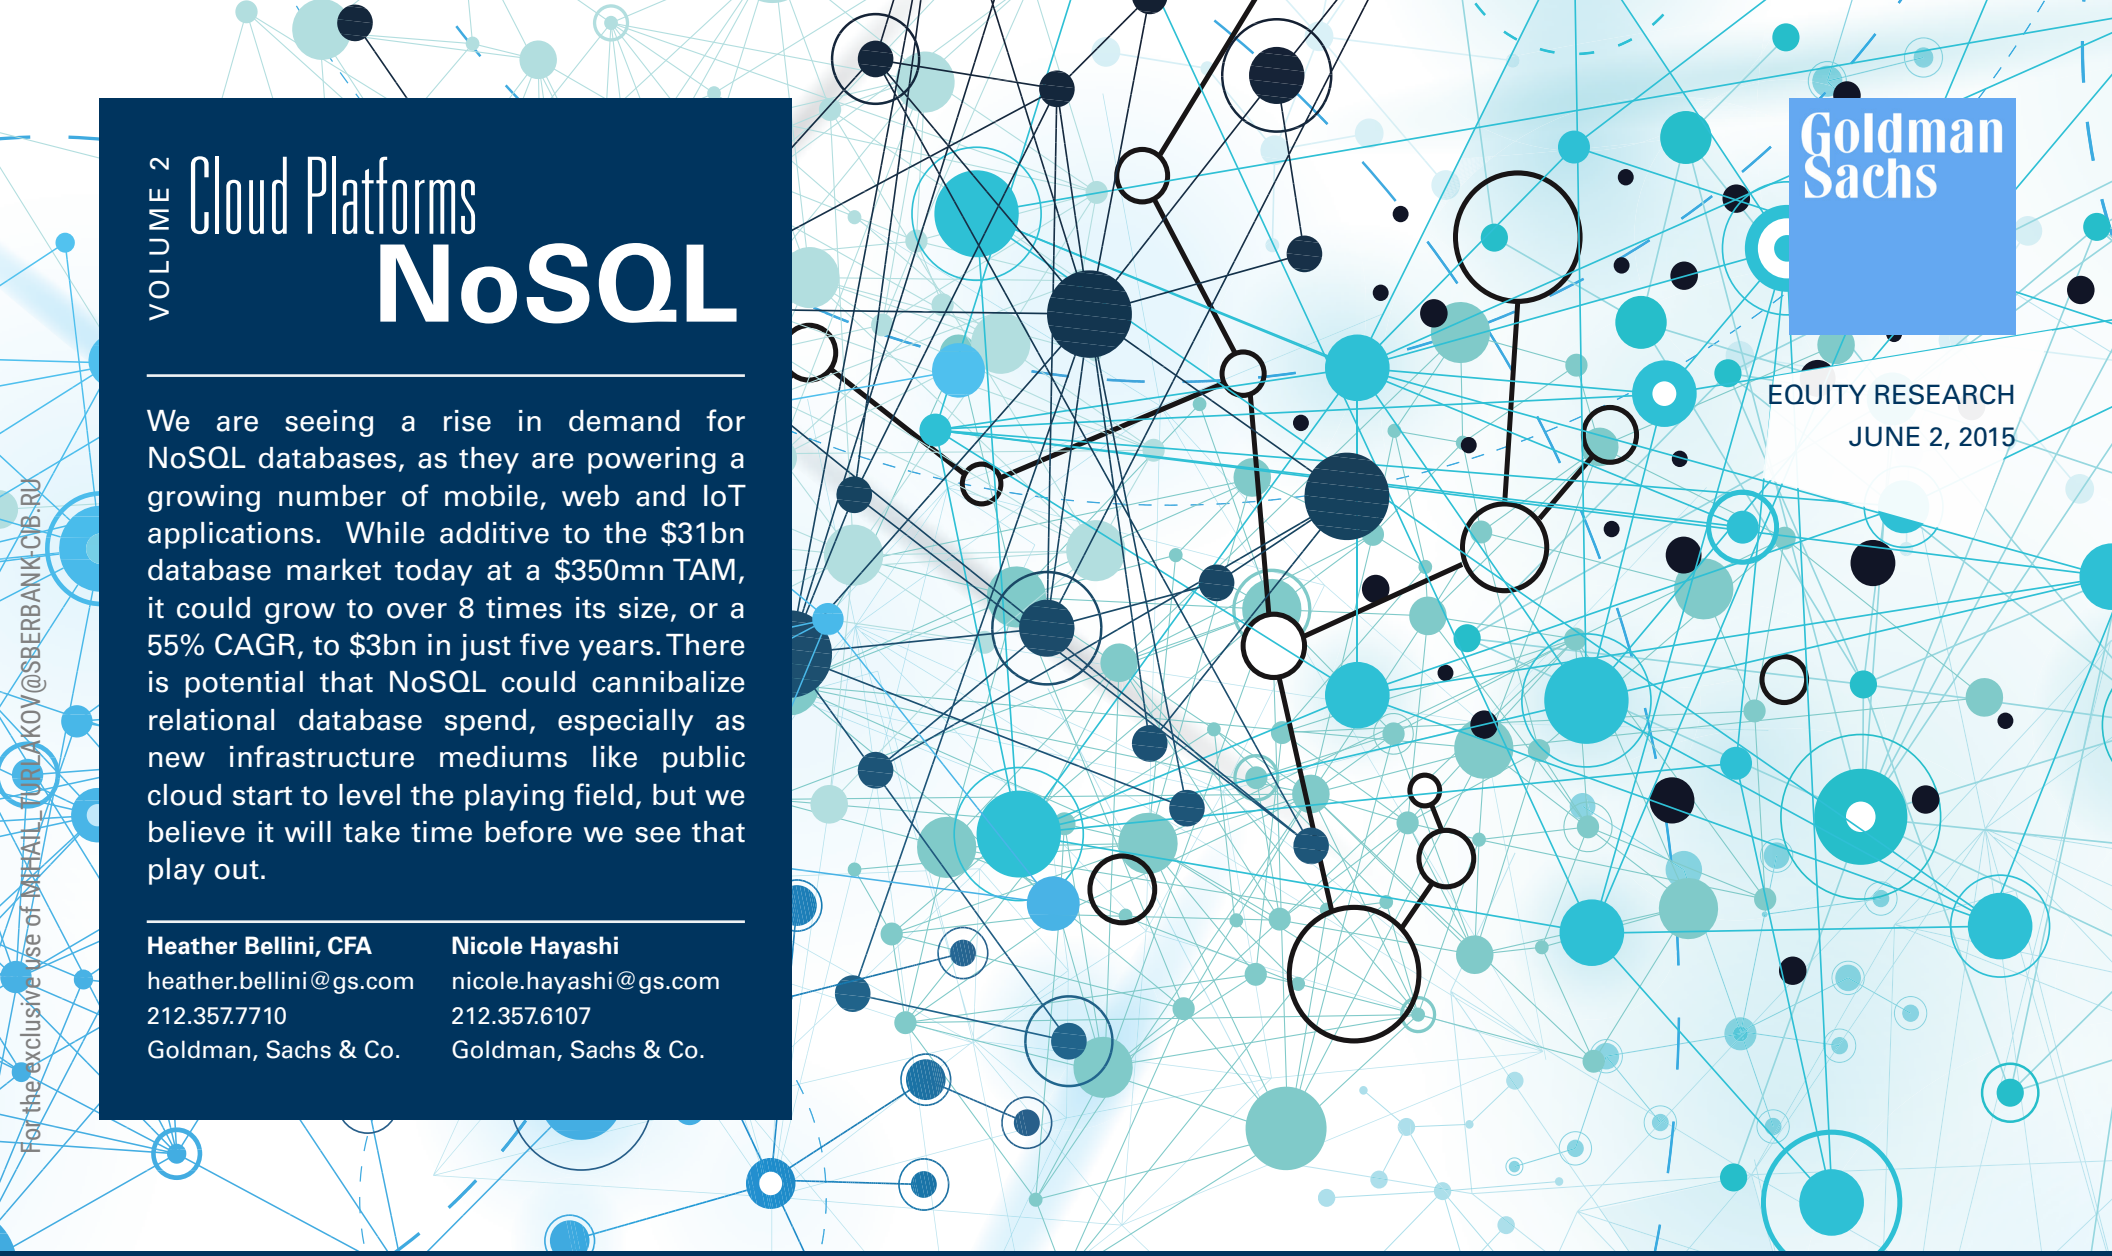

VOLUME 2

Cloud Platforms
NoSQL

We are seeing a rise in demand for NoSQL databases, as they are powering a growing number of mobile, web and IoT applications. While additive to the $31bn database market today at a $350mn TAM, it could grow to over 8 times its size, or a 55% CAGR, to $3bn in just five years. There is potential that NoSQL could cannibalize relational database spend, especially as new infrastructure mediums like public cloud start to level the playing field, but we believe it will take time before we see that play out.

Heather Bellini, CFA
heather.bellini@gs.com
212.357.7710
Goldman, Sachs & Co.

Nicole Hayashi
nicole.hayashi@gs.com
212.357.6107
Goldman, Sachs & Co.

Goldman Sachs

EQUITY RESEARCH
JUNE 2, 2015

Portfolio Manager Summary

This report is the second in a multi-part series designed to educate investors about cloud computing platforms. For more detail on the overarching trend, please see Cloud Platforms Volume One: Riding the Cloud Computing Wave, published on January 13, 2015

Background - This report examines NoSQL databases, an emerging technology that is changing the 30+ year old database industry. We examine the market, technology, key players and potential disruption.

What is it? - NoSQL stands for "not only SQL", and is a type of database that is being adopted by hundreds of large and small organizations from multiple verticals, including: Apple, Netflix, Conde Nast, McGraw Hill, and the UK National Health Service. These vendors are mostly using NoSQL databases for new applications (apps), taking advantage of their scalability at a lower cost, speed to deployment and ease of development.

Key Drivers - We are seeing an increasing adoption of NoSQL databases for web, mobile and Internet of Things applications. Mobile and web apps are expected to grow rapidly, from slightly under 50% of total apps in a typical organization five years ago to 70% today, and ~90% ten years from now (Gartner, October 2014). The opportunity is further enhanced as data doubles every two years, creating a greater need for NoSQL databases that are able to scale well beyond traditional enterprise application needs.

Where is the debate? - Investors have been asking how large the market is and what size it can reach over time, and if this disrupts traditional database spend.

- **Total addressable market (TAM) -** Today, we see NoSQL as additive to the overall database TAM in the near term, as most companies are beyond the testing phase and are now spending incremental dollars on multiple use cases. The market is still very early in terms of size. As of CY14, the NoSQL market was $350mn or 1% of the $31bn total database market (Gartner). The remaining 99% of the total database market is primarily made up of relational databases. By the end of 2019, we believe the NoSQL market could grow as large as $3bn–$6bn, or a 7-13% share of the total database market and a +55% to +76% CAGR.
- **Is NoSQL replacing relational databases? -** We believe NoSQL databases are mostly being used for new applications today, rather than a migration of existing database workloads from relational to NoSQL. And while there are instances where users replace relational databases, we have not heard of large enterprises lowering their relational database maintenance spend with dominant database vendors such as Oracle and Microsoft.

Who are the incumbents in this space? - Oracle, Microsoft and IBM are the largest relational database vendors today (mainframe and distributed), with CY14 revenue according to Gartner of $14bn (46% share), $6bn (20% share) and $5bn (17% share) respectively. All three have NoSQL offerings, and while those products' share is extremely small today, the upstarts in the market will have to contend with the incumbents' likely attempts to leverage their vast installed base and significant number of enterprise license agreements to try and become a bigger force in the market.

Potential for future cannibalization - Our research suggests NoSQL is yet to materially dislodge the use of existing relational database deployments. However, as organizations' NoSQL deployments become more substantial in dollars spent and production use cases, there is potential for enterprises to start rationalizing their database spend, either consolidating multiple NoSQL deployments into one NoSQL vendor or potentially lowering their spending on traditional database technology.

Databases in the public cloud - While more and more workloads are being moved to public clouds, our research suggests that existing workloads will continue to leverage their current database providers versus swapping them out for an alternative. Public cloud helps level the playing field for NoSQL vendors, as customers are now able to choose a NoSQL database just as easily as they can choose a relational database.

The NoSQL market's evolution - While we believe there is potential for NoSQL to continue to garner a larger and larger share of the relational database market, the question becomes which vendors will ultimately dominate this spending. Today NoSQL databases are categorized into four different types based on the data model: document style, key value, column/table style and graph. Over time, we believe the most successful NoSQL database vendors will have offerings that span across multiple types. We outline NoSQL vendors on Exhibits 5, 11 and 16. There are very few if any public market options for investing in NoSQL databases, as the pure plays are all private. However, large incumbents, such as Oracle (ORCL) and Microsoft (MSFT) have exposure, albeit de minimis, to this trend.

How does this relate to Hadoop? - NoSQL and Hadoop address different segments of data management. NoSQL databases are often used as alternatives to transactional workloads or relational databases (as a back end data store for an application), while Hadoop is as an alternative to analytical workloads or data warehouses (where data is aggregated from multiple databases for analysis). Hadoop is negatively impacting the data warehouse industry today but NoSQL is a positive for the database market in the near term.

Authors:

Heather Bellini, CFA (212) 357-7710 heather.bellini@gs.com Goldman, Sachs & Co.

Nicole Hayashi (212) 357-6107 nicole.hayashi@gs.com Goldman, Sachs & Co.

Shateel Alam (212) 902-6785 shateel.alam@gs.com Goldman, Sachs & Co.

Jack Kilgallen (212) 357-9111 jack.kilgallen@gs.com Goldman, Sachs & Co.

What is NoSQL? And what is a database?

NoSQL is a type of database. A database (DBMS) is a piece of software that contains data. The data contained within a database could be anything, for example – first name, time, price, revenue or Big Data and Internet of Things (IoT) data such as time series data and location coordinates.

There are two types of databases, transactional and analytical. NoSQL databases and relational databases (relational database management system, RDBMS) are both transactional databases.

The purpose of a transactional database is to store and organize data, and serve as the back end data store for an application. For example, a sales person is visiting his client. To get up to speed prior to the meeting, he checks the salesforce.com app on his phone. The sales person opens up the application and selects the client's name so he can look at the client's profile. The application immediately pings the Oracle database on the back end to retrieve data about the client. Once completed, the application then serves up the client's profile on the phone which includes information from the database. Now the sales person can see the client's profile on his phone, which includes the name and address of the person he's meeting, the total dollar amount sold to them in the past year, notes from the last meeting, and more.

Exhibit 1: Example of the purpose of a database

Source: Goldman Sachs Global Investment Research. For illustrative purposes only.

Why organizations are adopting NoSQL databases

The determination of whether to use a NoSQL or relational database depends on the type of application and data it is being used for. We believe traditional enterprise applications, like a customer relationship management system or enterprise resource planning application, have and will continue to use relational databases. We outline key reasons why organizations are adopting NoSQL databases below.

(1) **Better scalability for specific use cases –**Scalability is often the top cited reason of why users adopt or are considering adopting NoSQL databases. Gartner's survey of IT professionals found that 80% of respondents said their motivation for using NoSQL was for "scalability", higher than any other choice (February 2014). While scalability can have multiple meanings, in this context, it is often defined as (1) scaling to a large number of transactions per second, (2) scaling at a lower cost than relational databases, and (3) scaling across multiple data centers.

#1 Scaling to a large number of transactions per second

NoSQL and relational databases process and replicate data differently, impacting the speed transactions are completed and the consistency or integrity of the data. NoSQL and relational databases store multiple copies of data on different servers (also called nodes). When data is modified, added or deleted, each copy needs to be updated or stored in multiple servers. For the most part, the biggest difference between NoSQL and relational databases is that relational databases require each copy of the modified data to be updated on each server before the next transaction can be performed. In contrast, NoSQL databases can be updated on one node and then move onto the next transaction before it is done replicating on additional nodes. Over time, the data is replicated across other nodes eventually (called "eventual consistency").

For example, if a user "likes" a Facebook post, in a relational database, the "like" transaction would be written to Server 1 on the west side of the data center, replicated on Server 205 on the east side of the data center, then replicated on Server 295 elsewhere in the data center. All three transactions must be completed before the next transaction could be processed. In a NoSQL database, the "like" transaction is written to Server 1. Before the data is replicated on Servers 205 and 295, the next transaction can be processed. Eventually, the data is further replicated to Servers 205 and 295.

That means each transaction can be completed in a NoSQL database in a fraction of the time it takes a relational database to complete the process. This becomes a critical capability when the number of transactions scales to hundreds of thousands or a million per second.

The tradeoff to this method is that the data in NoSQL databases isn't 100% consistent at all times (where data is exactly the same across every instance or node containing the data). There could be times where the application could be reading stale data, or the user could lose data completely. Compare this to the relational database example, where data is not lost or inconsistent because the data is fully replicated before the next transaction can begin. We believe this difference in is one of the reasons why NoSQL applications are less likely to be used for

situations where data cannot be lost, such as an order system, financial application, or customer relationship management application.

#2 Scaling at a lower cost than relational databases

NoSQL databases can scale to support over a million transactions per second by adding dozens of inexpensive commodity servers (known as scale out, or scaling horizontally). This differs from relational databases, where users often add expensive servers, storage, and networking with more memory and CPU (known as scale up or scaling vertically).

- For example, The UK National Health Service spent three years moving their centralized database of patient health and prescription data off of Oracle and onto NoSQL database Basho Riak to be more cost-effective. With their new system, the UK National Health Service was "able to not only cut costs, but also improve the performance and reliability of the system overall" (Basho).

#3 Scaling across multiple data centers, onpremise and in the public cloud

We examine two use cases below:

- Scaling across multiple data centers - Under Armour moved off of Microsoft SQL Server and onto NoSQL database MongoDB as it was becoming difficult to scale and add new features on SQL Server as their business grew. With their NoSQL database, Under Armour was able to scale out using commodity hardware and support multi-data center replication, while maintaining high performance (MongoDB).
- Almost infinite scalability in the public cloud - Netflix moved off of Oracle and onto NoSQL database Apache Cassandra because in their specific use case they found that Oracle did not run well on virtualized hardware. Apache Cassandra allows Netflix to use multiple data centers, and because they run it on the cloud and aren't constrained by hardware, "as long as they can get nodes, [they] can scale almost infinitely" (Netflix). Netflix stress tested Apache Cassandra back in late 2011 and was able to run a workload of 1.1mn writes per second on Amazon AWS.

(2) **Speed to deployment and ease of development -** There is a fundamental difference in how the data is set up, stored and accessed between relational databases and NoSQL databases. Relational databases require schema, which is a way to logically structure and organize the data. For example, defining fields as "SKU", "cost", "units", "currency", etc. The problem with schema is that it needs to be defined before any data can be stored in the database. This often requires a separate person, a database administrator (DBA), to execute this task. This is because there can be hundreds of tables or fields within one database, plus dozens of databases outside of that original database that have to interrelate in a logical way in order to retrieve the intended data (e.g., total sales this month by region, total units sold in London by color). The time typically required to add or adjust an existing database can be very time consuming. This gives NoSQL databases a major advantage. Because NoSQL databases don't require SQL and are essentially schema-less, application developers tend to gravitate towards NoSQL databases when they need to create an application within a short deadline. Developers can write an application and start storing data immediately, instead of waiting for a DBA to get a database fully operational.

Exhibit 2: Steps to get up and running – NoSQL versus relational database

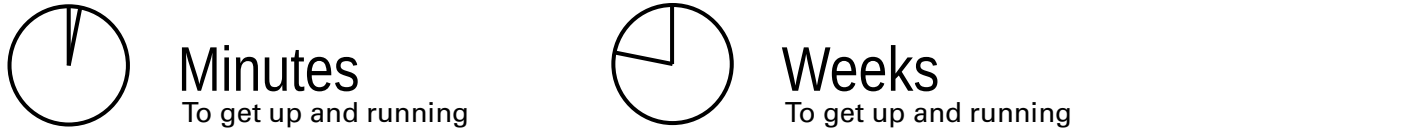

Source: Goldman Sachs Global Investment Research. In practice, each implementation varies, but this Exhibit gives an high level view of what could happen.

(3) **Not constrained to the SQL standard -** The term "NoSQL" stands for "not only SQL", inferring that that NoSQL databases offer capabilities beyond relational databases. Unlike relational databases, which use the standard programming language, Structured Query Language (SQL), to access and manipulate the data, NoSQL databases do not need to adhere to the SQL standard. This gives users more flexibility in how they write data to and retrieve data from the NoSQL database. The downside is that it could require more skilled talent and take longer to run business intelligence dashboards, reports or analytics on NoSQL databases, depending on the use case.

(4) **NoSQL databases are roughly 1/3 the cost of relational databases when there is no ELA in place –** While we estimate NoSQL databases are roughly 1/3 less expensive than enterprise grade relational databases on list price, companies who choose NoSQL databases are primarily selecting them due to scalability and ease of development, not cost. In addition, most large organizations have existing enterprise licensing agreements (ELAs) with relational database vendors. And for the most part, organizations have not cut spending on relational database maintenance contracts to use NoSQL databases. Therefore, we believe NoSQL databases are additive to the database total addressable market (TAM) in the near term.

Exhibit 3: Pricing example

User or company with no existing databases	Customers on an existing relational database ELA
$48k/processer + $10k annual maintenance **Oracle Enterprise Edition** List price is $48k/processor license plus $10k for maintenance	Multi-Million $ ELA **Existing ELAs** Many ELAs don't have incremental costs for adding an additional database
OR	PLUS
$4-18k/node per year **NoSQL** Pricing varies by vendor but we've seen list prices range from $4k - $18k/node/year including support	$4-18k/node per year **NoSQL** Pricing varies by vendor but we've seen list prices range from $4k - $18k/node/year including support

Source: Public company information and Goldman Sachs Global Investment Research. Oracle pricing from Oracle's 4/9/15 price list. A node is a server, and multiple processors can be on any given server.

The database market is expected to grow from $31bn in 2014 to $45bn in 2019 (8% CAGR)

Gartner estimates that the total database market, which includes relational and NoSQL databases, was $31bn in 2014, and will grow to $45bn in 2019, or an 8% five year CAGR. Gartner includes software license and related support and maintenance in their calculation of TAM. In the Exhibit below, we outline incremental annual database spend, which equates to $1.4-3.5bn annually.

Exhibit 4: Database TAM (includes relational and NoSQL databases)

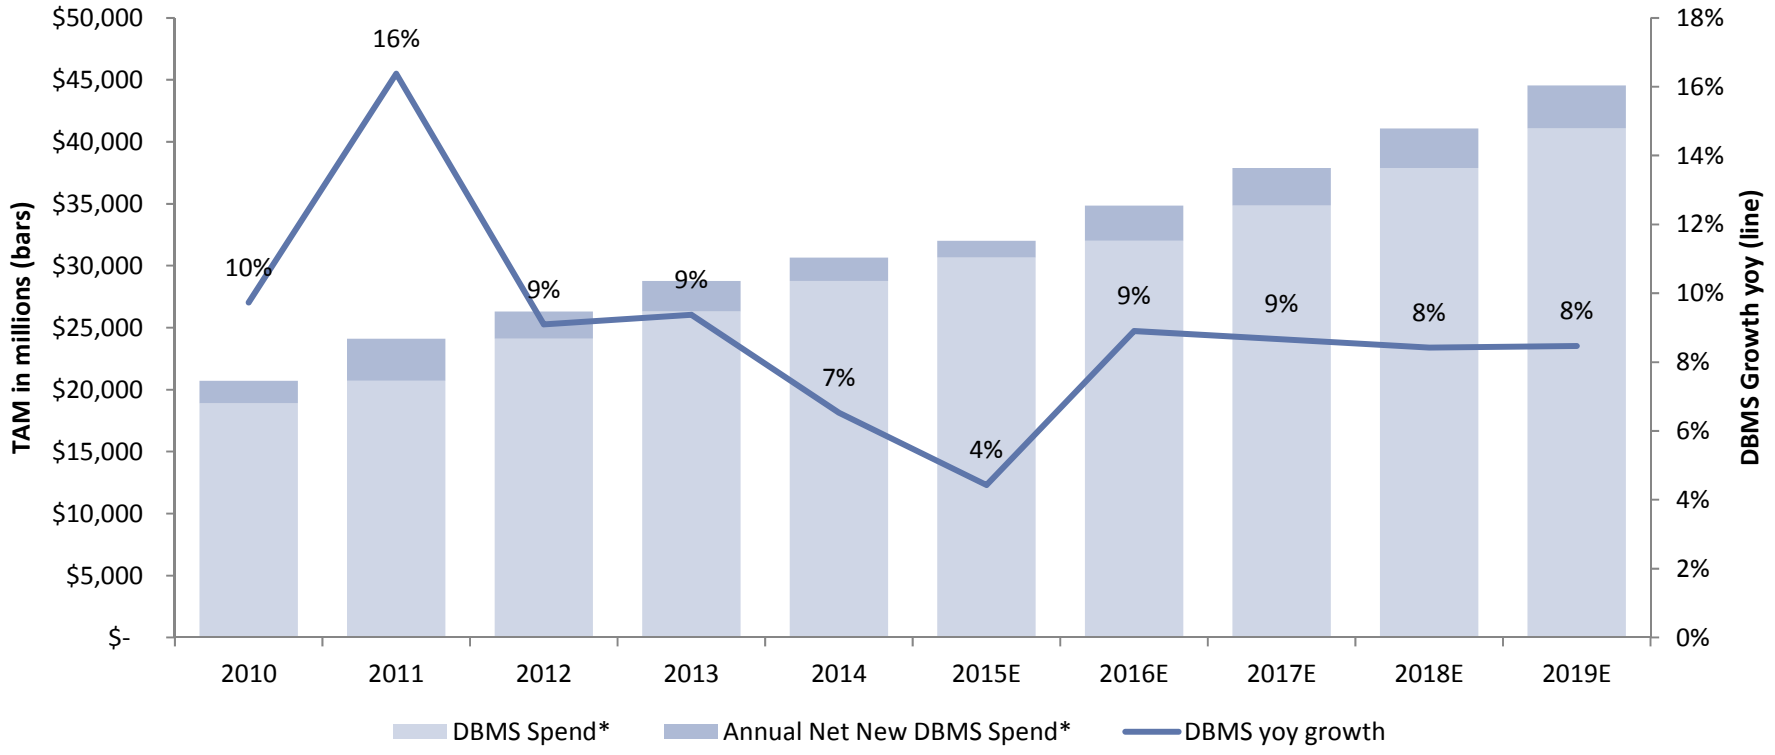

*Source: Gartner (2012-2019 estimates from 1Q15 forecast, 2009-2011 from the latest forecasts available) and Goldman Sachs Global Investment Research. For the purpose of this report, we only show post-relational (modern era) DBMS forecasts. We calculate the net new annual DBMS spend as Year 1's total database TAM minus Year 0's total database TAM. *Includes software license and related support and maintenance.*

The NoSQL market is still early, at $350mn or 1% of the market

The Exhibit below outlines the database market share for 2014. Relational databases still make up the vast majority of the market at $31bn or 99%, while the NoSQL market is still very early, at just $350mn or 1% (Gartner December 2014). We also outline the top vendors by revenue, according to Gartner estimates. Oracle is by far the largest vendor in this space, making up almost half of the market, with 46% share, followed by Microsoft at 20% and IBM at 17%. We note that traditional relational database vendors may disrupt themselves as they try to gain share of the NoSQL market. IBM's NoSQL market share grew from zero share in 2013 to 4% in 2014, as they acquired NoSQL database vendor Cloudant, while their relational database market share declined from 17.8% to 16.9% (Gartner, March 2015). We also note that there were four NoSQL acquisitions year to date, as Apple acquired FoundationDB, Amazon AWS acquired Amiato, CenturyLink acquired Orchestrate and DataStax acquired Aurelius.

Exhibit 5: Total database market share and vendors in 2014 (Gartner's vendor share forecasts as of 3/31/15)

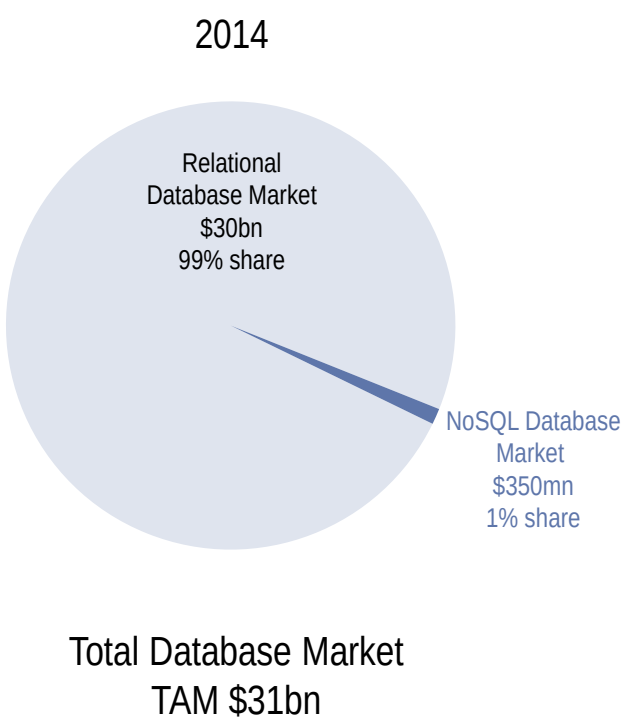

		2014 Revenue ($mns)	Y/Y Growth	Share of Total DBMS
Oracle	$	13,993	4%	45.6%
Microsoft	$	6,189	13%	20.2%
IBM	$	5,151	1%	16.8%
SAP	$	2,267	13%	7.4%
Other RDBMS Vendors	$	2,706	4%	8.8%
Total Relational*	$	30,307	6%	98.9%

		2014 Revenue ($mns)	Y/Y Growth	Share of Total DBMS
MongoDB	$	86	66%	0.3%
MarkLogic	$	76	61%	0.2%
DataStax	$	38	34%	0.1%
Couchbase	$	24	100%	0.1%
IBM (Cloudant)	$	21	NM	0.1%
Basho	$	17	65%	0.1%
Oracle NoSQL	$	10	83%	0.0%
Other NoSQL Vendors	$	80	NM	0.3%
Total NoSQL**	$	350	NM	1.1%

Total Database	$	30,657	7%	100.0%

*Source: Gartner post-relational database market size estimate and market share forecasts as of March 2015 and compiled by Goldman Sachs Global Investment Research. Vendors forecasts are from Gartner 3/31/15. *Assumes all of Gartner's post-relational DBMS that is not NoSQL is relational. **This chart combines the Gartner vendor forecasts from 3/31/15 and NoSQL market forecast of $350mn from Gartner's December 2014 report. Revenue includes software license and related support and maintenance.*

We believe there is a time and place for each type of database (NoSQL or relational)

Disrupting the relational database market will not be a simple task. Larry Ellison, the co-founder, Chairman and CTO of Oracle, and team released the first commercial SQL relational database in 1979. Since then, relational databases have become the standard back end system for critical enterprise and consumer applications. Today relational databases are the backbone for financial management applications, customer relationship management applications, point of sale applications (in Exhibit 6 below), inventory management applications, human capital management applications and more. We believe these types of applications are best suited for relational databases and will continue to run on relational databases in the future.

On the other end of the spectrum, we are seeing NoSQL databases used for net new applications, such as mobile, web and IoT apps (for example, the fitness app in the Exhibit below). We believe these types of apps require quick deployment and scalability that could reach to over a million transactions per second. While the market today is heavily weighted towards relational databases, new mobile, web and IoT applications are growing rapidly.

Exhibit 6: Examples of different types of use cases by type of database

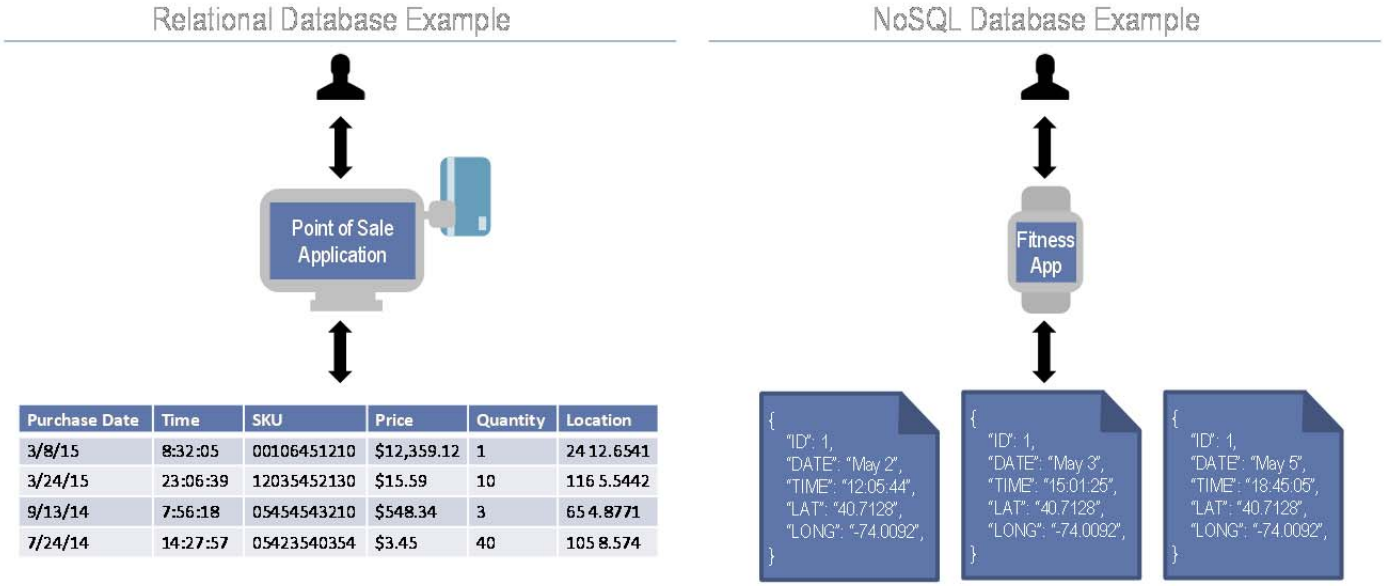

Source: Goldman Sachs Global Investment Research. NoSQL JSON file examples from Couchbase.com, May 2015.

The NoSQL market's evolution

We believe we are in the early days of Phase II of the NoSQL market. Overall, NoSQL has been additive to the relational database market, as we are seeing much of the growth in NoSQL from net new use cases or applications, rather than a migration of existing database workloads from relational to NoSQL. And while we have seen some replacements of traditional relational databases, we have not heard of many large enterprises lowering their relational database maintenance spend with dominant database vendors such as Oracle and Microsoft yet. However, as we start to see customers progress towards Phase III, we may start to see disruption as companies will likely start rationalizing their spend with NoSQL vendors (for example, consolidating spend from multiple NoSQL vendors down to one or two) or their total database spend (lowering spend with relational database vendors).

Page 13 discusses our thoughts on market sizing if the NoSQL market were still in Phases I or II in five years (Scenarios 1-3). We believe Phase III will take over five years to achieve at the market's current trajectory. In the long run, as the NoSQL market reaches maturity, we believe its size could reach to over 20% of total database spend. This equates to 20% of cumulative net new database spend, plus 1-2% cannibalization of the relational database market.

Exhibit 7: Evolution of NoSQL

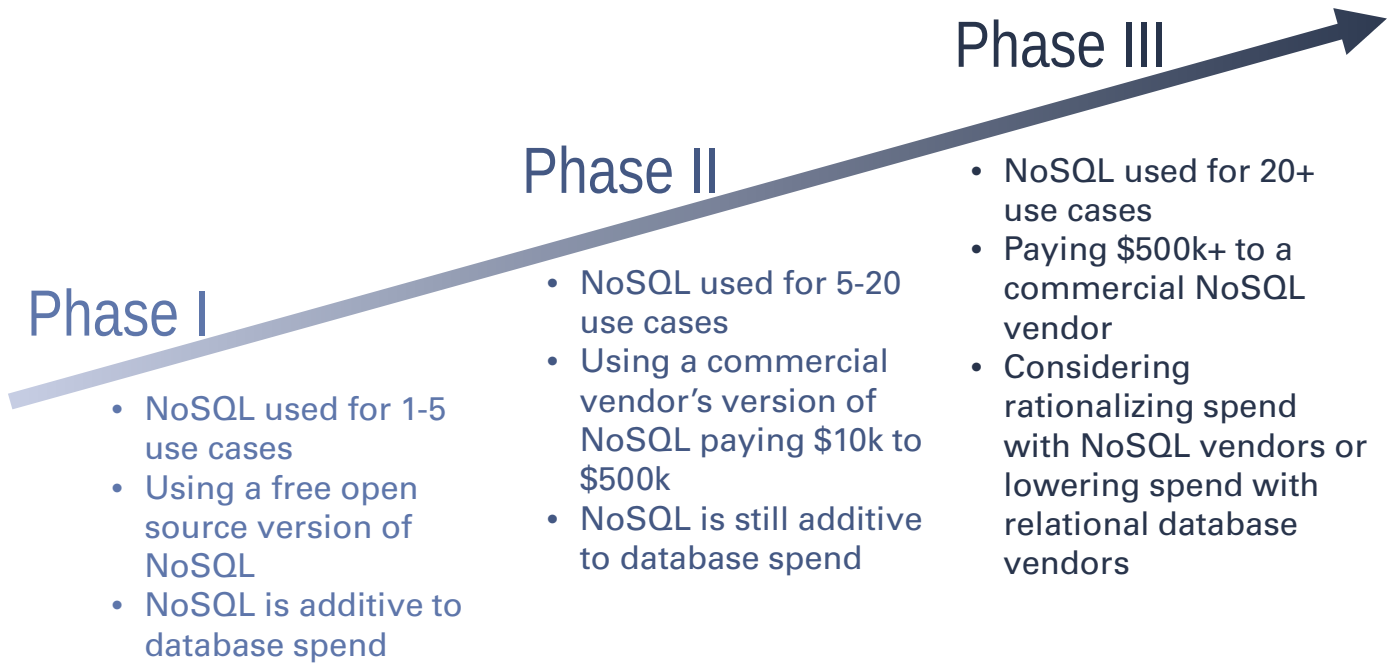

Source: Goldman Sachs Global Investment Research.

We believe the NoSQL market could reach $3-6bn in 5 years (+55% to +76% CAGR or 7-13% of the total database market)

Because the NoSQL market is still in its infancy, research firms Gartner and IDC do not forecast future market share expectations. We attempt to quantify this market using three potential scenarios, ranging from 20-40% of net new database spend over the next five years: Scenario 1 assumes 20% of net new dollar database spend between 2015 and 2019 is NoSQL (similar to CY14), equating to a $3.1bn market or 7% market share ($41.4bn RDBMS), Scenario 2 assumes 30% of net new dollar database spend between 2015 and 2019 is NoSQL, equating to a $4.5bn market or 10% market share ($40.0bn RDBMS), and Scenario 3 assumes 40% of net new dollar database spend between 2015 and 2019 is NoSQL, equating to a $5.9bn market or 13% market share ($38.6bn RDBMS).

Exhibit 8: NoSQL TAM Scenarios in five years

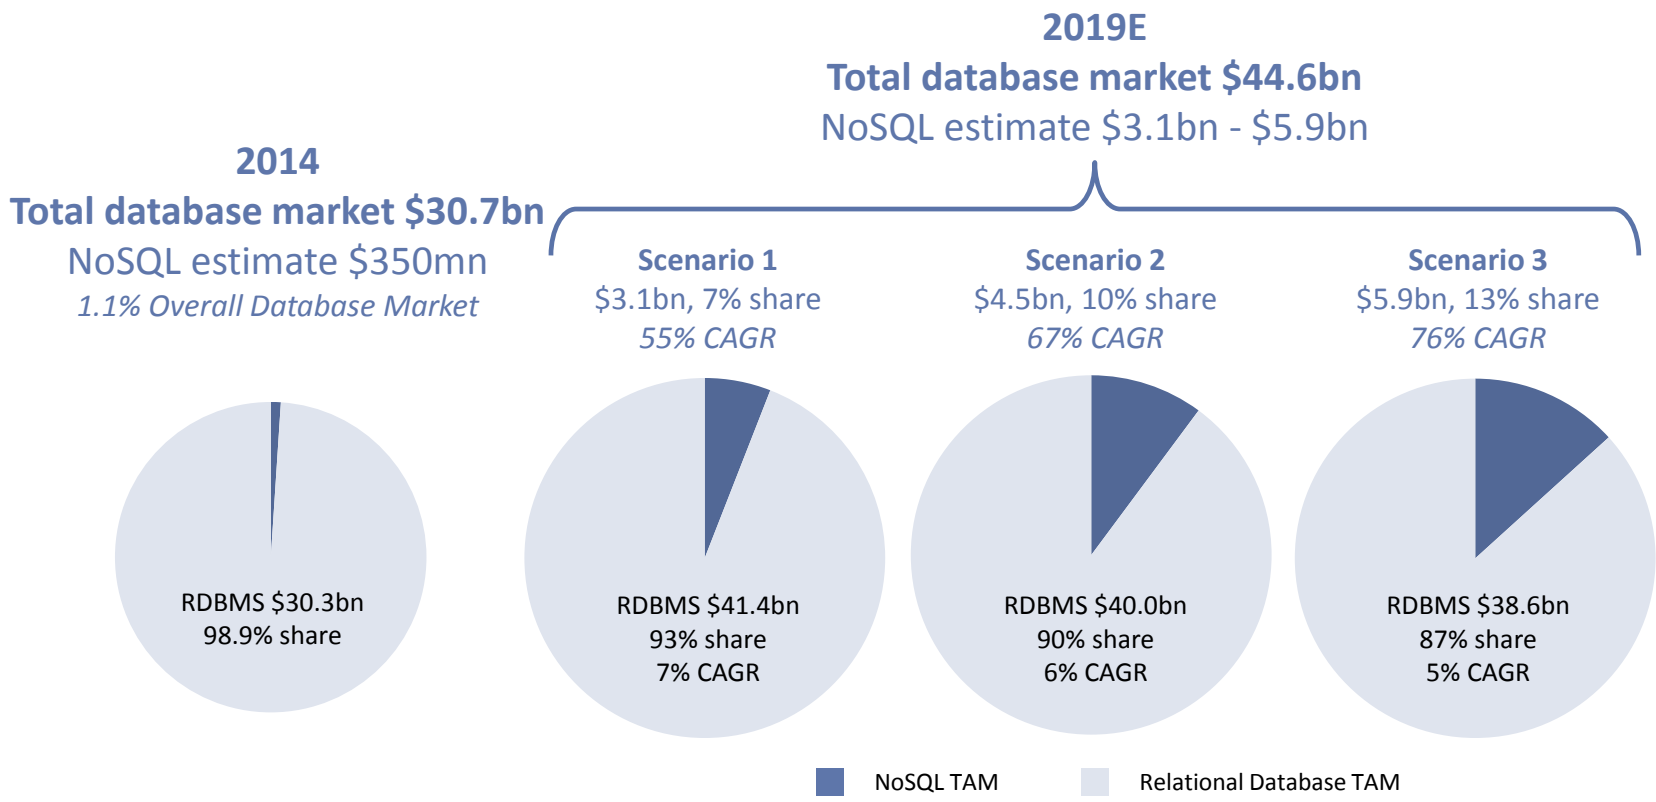

Source: Goldman Sachs Global Investment Research. AM reflects Gartner's 1Q15 forecast for databases (excludes pre-relational). $350mn NoSQL TAM estimate for 2014 reflects Gartner's December 2014 report.

Types of NoSQL Databases: Document, Key Value, Column or Table Style and Graph

There are four types of NoSQL databases, each optimal for different purposes. We estimate document databases have the largest share of the market today, led by MongoDB, MarkLogic and IBM at just over 50% share of the NoSQL market (~$185mn in revenue in CY14), followed by key value offerings which represented roughly 15% of the NoSQL market, followed by column or table offerings at just over 10% share. We summarize key differences between the types of databases in Exhibit 9, and in more detail on the following pages:

Exhibit 9: Summary of NoSQL Types

Type	Description	Organizations Using Each Type	High Performance	High Scalability	High Flexibility	Low Complexity
Document	Data is stored in documents (like JSON and XML)	ADP, Conde Nast, Foursquare, MetLife, MTV Networks, Under Armour	✓		✓	✓
Key Value	Data is stored as "values" on a set of "keys"	LinkedIn, McGraw Hill, Turner Broadcasting, UK National Health Service	✓	✓	✓	✓
Column or Table	Data is stored in flexible columns or tables	Apple, eBay, Facebook, Netflix	✓	✓		✓
Graph	Processes many-to-many connections	CrunchBase, Pitney-Bowes, TomTom, Walmart			✓	

Source: Goldman Sachs Global Investment Research. The four right columns is a summary from Planet Cassandra.

Document databases

- **Data Model/Structure:** As the name suggests, data is stored in documents instead of rows. Document databases store data into a document using JSON (JavaScript Object Notation), XML (Extensible Markup Language) or other formats. With popular and flexible formats such as JSON, users can quickly get up and running, adding new keywords (similar to the "FIRST" and "LAST" in the following Exhibit) as needed instead of defining them upfront. This structure offers high flexibility, high performance and low complexity.
- **Use Case Example:** Document databases are often optimal for "read" heavy applications and rapid application development. A "read" is when an application accesses data from the database. For example, if a user goes onto their

smartphone and looks at a friend's profile on Facebook to see their birthday, behind the scenes, the Facebook app accesses a database, where the birthday date is stored. A "read" is when the application reads the birthday date from the database.

- **Case Studies:**
    - MetLife used MongoDB to build a single view of 100 million customers in 90 days. The application lets employees access client information from one central app instead of 70
    - ADP used MongoDB, allowing millions of users to check their paystubs
    - Conde Nast used MarkLogic for its digital asset management system. They were able to load assets into the NoSQL database without creating a structure around it beforehand

Key Value databases

- **Data Model/Structure:** This type of NoSQL database consists of "keys" which contain data or "values". The values stored are called BLOBs or binary large objects. The database does not know what is stored on each BLOB, and therefore users query the primary key, not all fields. This gives the database high performance, scalability, and flexibility, with low complexity.
- **Use Case Example:** Key Value databases are often optimal for a blend of "reads" and "writes", scale and constant streams of data. Key value databases are popular amongst gambling sites, as users place bets ("writes" to the database) and other users look at the updated odds or cards ("reads" of the database) at large scale (hundreds or millions of users).
- **Case Studies:**
    - LinkedIn used Couchbase to build a scalable index for its metrics visualization engine. And it delivers 400k operations per second on just four servers
    - The UK National Health Service is using Basho's Riak database. This is their main patient database that logs millions of citizen's health information that was running on Oracle (The Register, 9/9/14)
    - Turner Broadcasting used Basho's Riak databases as the back end systems to CNN.com, as well as for hyper-scale events such as the 2012 US presidential elections

Column or Table style databases

- **Data Model/Structure:** Conceptually, Column or Table Style databases are similar to massive Microsoft Excel spreadsheets. Data is stored in flexible columns or tables instead of rows. The data model is less structured than relational databases (which require a consistent number of columns) but slightly more structured than Key Value or Document type databases. This makes Column or Table style databases less flexible than Key Value or Document databases, but still provides high performance and is highly scalable.
- **Use Case Example:** Column or Table style databases are often optimal for "write" heavy applications, for scalability, and semi-structured data. A "write" is when data is written into a database. For example, for web sites, enterprises often track every time a user clicks on an item, such as clicking on a specific pair of shoes, the shoe color, the shoe size, adding a

product to their shopping cart, or closing the browser before the purchase is made. In this scenario, a "write" is when the click is recorded onto a database.

- **Case Studies:**
  - eBay is using Apache Cassandra to track "want", "own" and "like" click data
  - Netflix is changing the data architecture of its popular streaming app, with Apache Cassandra at its core. Netflix collects data each time a user clicks play or pause when watching content, enabling customers to stop watching on one device and continue watching it on another seamlessly. Netflix benchmarked Apache Cassandra's scalability in November 2011 and found that it could handle over one million writes per second
  - Apple has one of the largest production deployments of Apache Cassandra, with over 75k nodes storing over 10 petabytes of data.
  - Facebook chose to use Apache HBase in 2010 for its real time messaging system. At the time, the system handled over 350mn users and 15bn person-to-person messages

Graph databases

- **Data Model/Structure:** Graph databases contain connected entities (nodes) with attributes (key value pairs). Relationships are the connections between two nodes, and can have any properties. This highly flexible model is similar to the idea of a LinkedIn relationship graph (InMaps).
- **Use Case Example:** Graph databases are often optimal to examine the relationships between elements, such as displaying the best recommendations for a user.
- **Case Studies:**
  - Walmart uses Neo4j to serve personal recommendation to its web customers
  - TomTom uses Neo4j to improve its map data from collecting data from their fleet of surveying vehicles and users
  - CrunchBase runs its comprehensive database of 650k people and company profiles on Neo4j

Source: vendors and customer websites, except where noted.

Exhibit 10: Examples of different data models by type of NoSQL database

Document

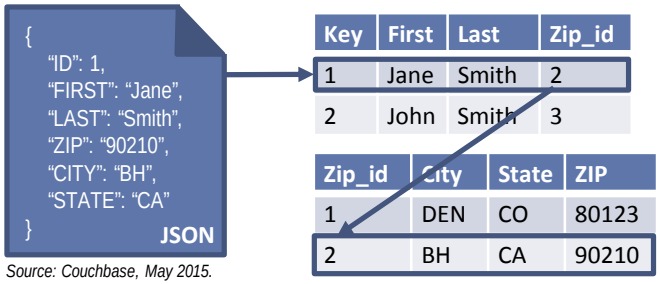

Source: Couchbase, May 2015.

Column or Table Style

User

123456	Name	Email	Phone	State
	John Smith	jsmith@gmail.com	212-555-5555	NY

Item Likes	Item IDs		
123456	12535	24546	54851
	iPhone	iPad	iWatch

Source: eBay, May 2013.

Key Value

Key: 1	ID: js	First Name: John	
Key: 2	Email: jsmith@gmail.com	Location: New York	Age: 30
Key: 3	Facebook ID: johnsm	Password: *****	Name: John

Source: Aerospike, May 2015.

Graph

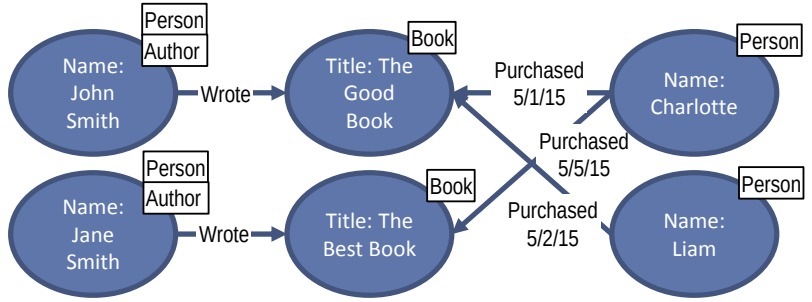

Source: Neo4j, May 2015.

Source: Cited above plus public available data, Gartner and Goldman Sachs Global Investment Research.

NoSQL vendors are expanding their reach, embracing multiple types (multi-modal)

We have seen a trend where NoSQL databases are expanding their capabilities into other areas, becoming "multi-modal", for example, MongoDB started as a document database and expanded their reach to support key value functionality. Couchbase and Amazon Web Services DynamoDB started as key value and then adopted document models. Additionally, DataStax, originally a table style database, recently acquired Aurelius to expand into the graph database space. We believe this trend will continue.

Exhibit 11: Commercial NoSQL databases by type

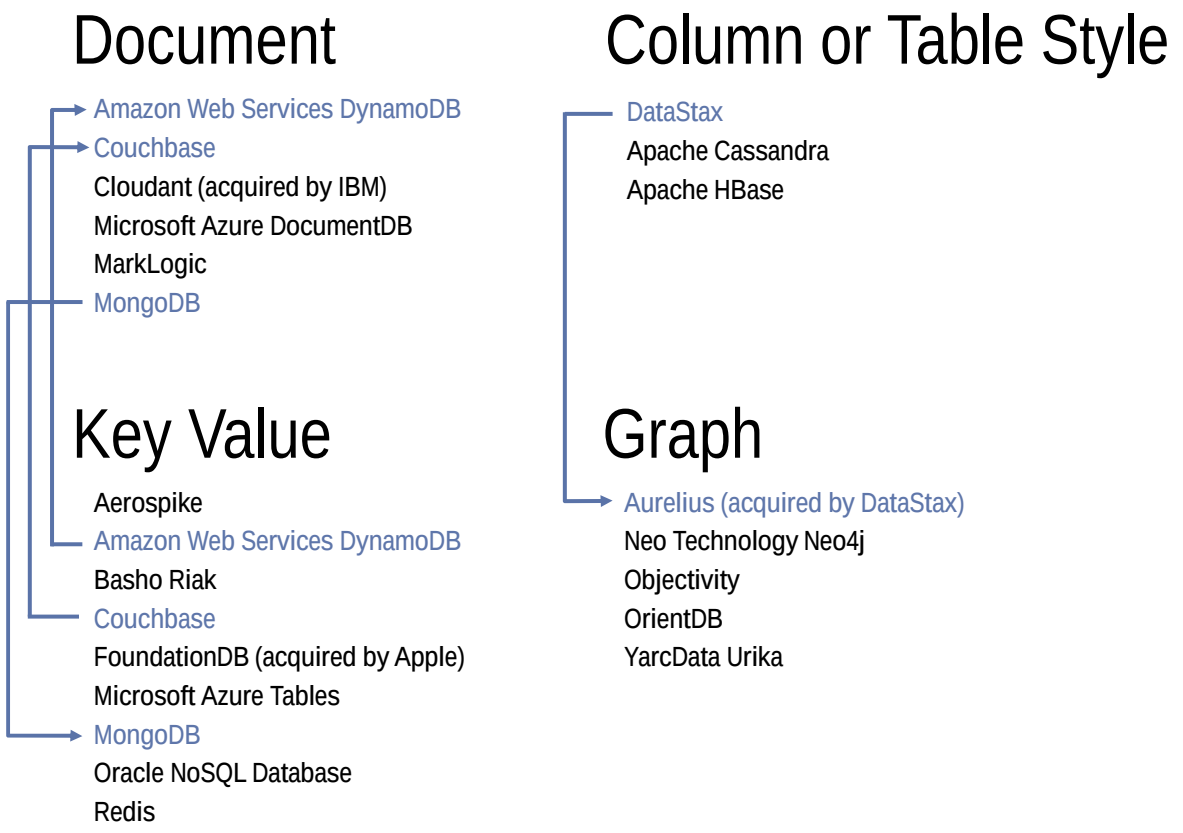

Source: Publicly available data, Gartner (combined from February 2014 and December 2014 NoSQL reports), compiled by Goldman Sachs Global Investment Research. Note: this list is not comprehensive of every vendor that has a multi-modal approach. Source files for Amazon and Couchbase: http://www.allthingsdistributed.com/2014/10/document-model-dynamodb.html http://blog.couchbase.com/key-value-or-document-database-couchbase-2-dot-0-bridges-gap

Public cloud helps level the playing field for NoSQL vendors

As more workloads move to the public cloud, leading database vendors such as Oracle, Microsoft, and IBM run the risk of customers becoming more vendor agnostic. Customers are now able to choose a NoSQL database such as Amazon's DynamoDB, MongoDB or Basho Riak, just as easily as they can choose a relational database such as Oracle or Microsoft SQL Server. As mentioned earlier, Netflix decided to use NoSQL database Apache Cassandra on Amazon AWS because it allows them to run on multiple data centers, and scale almost infinitely.

We believe NoSQL adoption will mostly impact new applications being built in the public cloud, rather than existing workloads migrating to the public cloud, as we are seeing companies who are migrating use the same infrastructure software as they are using on premise.

Exhibit 12: Cloud Platform Vendors – database presets/featured options

Amazon Web Services	Microsoft Azure	Google Cloud Platform	IBM Softlayer
MySQL	SQL Server	Google Cloud Datastore	Microsoft SQL Server
PostgreSQL	Oracle Database	Google Cloud SQL	MySQL
Oracle	DataStax Enterprise	Cassandra	Cloudera Hadoop
Microsoft SQL Server	Microsoft SQL Database (Database as a Service)	MongoDB	MongoDB
Amazon Aurora	Microsoft HDInsight		Basho Riak
Amazon DynamoDB			
Amazon RedShift			

Source: This chart was published in our Cloud Platforms Volume 1 report, titled "Riding the Cloud Computing Wave" on January 13, 2015. Data from Amazon, Microsoft, Google and IBM websites as of December 2014. These are the presets given by each vendor in the order they are presented on the website. These are not the only databases available by each vendor. Customers can access more operating systems by searching within each vendor's website or marketplace. For Google, Cassandra and MongoDB are "click to deploy software packages" options.

Additional NoSQL stats

Exhibit 13: NoSQL stats

Subscription

Per Node or Per Core

Pricing is typically on a subscription basis, per node or per core, unlike traditional relational databases that are on a license/maintenance basis

Open Source

Unlike traditional databases, users can try open source versions of NoSQL databases for free at first. Once they start using the product in a live environment, they often look for enterprise support and start to pay one or multiple NoSQL vendors.

~$850mn of Capital Raised

The five largest pure play NoSQL vendors have cumulatively raised roughly $850mn from investors.

M&A

Year to date, there were 4 NoSQL database acquisitions: DataStax/Aurelius (2/3), Apple/ FoundationDB (3/25), Amazon AWS/Amiato (4/20) and CenturyLink/ Orchestrate (4/21). We expect consolidation to continue as the space matures

Source: Company data compiled by Goldman Sachs Global Investment Research. Capital Raised from Crunchbase (May 2015)

How does NoSQL compare to Hadoop?

NoSQL and Hadoop address different segments of data management. Hadoop is a software-framework that supports large-scale computing, and is used for analytic workloads. Hadoop is seen as an alternative to data warehouses (a place where data is aggregated from multiple databases for analysis). Hadoop is negatively impacting the data warehouse industry today as users are taking data out of the data warehouse and putting it into Hadoop to analyze. Our team has found that Hadoop is typically 1/10th the cost of traditional data warehouses. This differs from NoSQL databases, which are often used as an alternative for transactional workloads or relational databases (as a back end data store for an application). Where both worlds collide is with Apache HBase and Apache Cassandra, which are known as Hadoop databases. Both HBase and Cassandra are table style NoSQL databases. Despite being known as databases for Hadoop, they are for the most part used as transactional databases (an application database), rather than analytical or as a data warehouse.

Exhibit 14: Analytical and Transactional database types

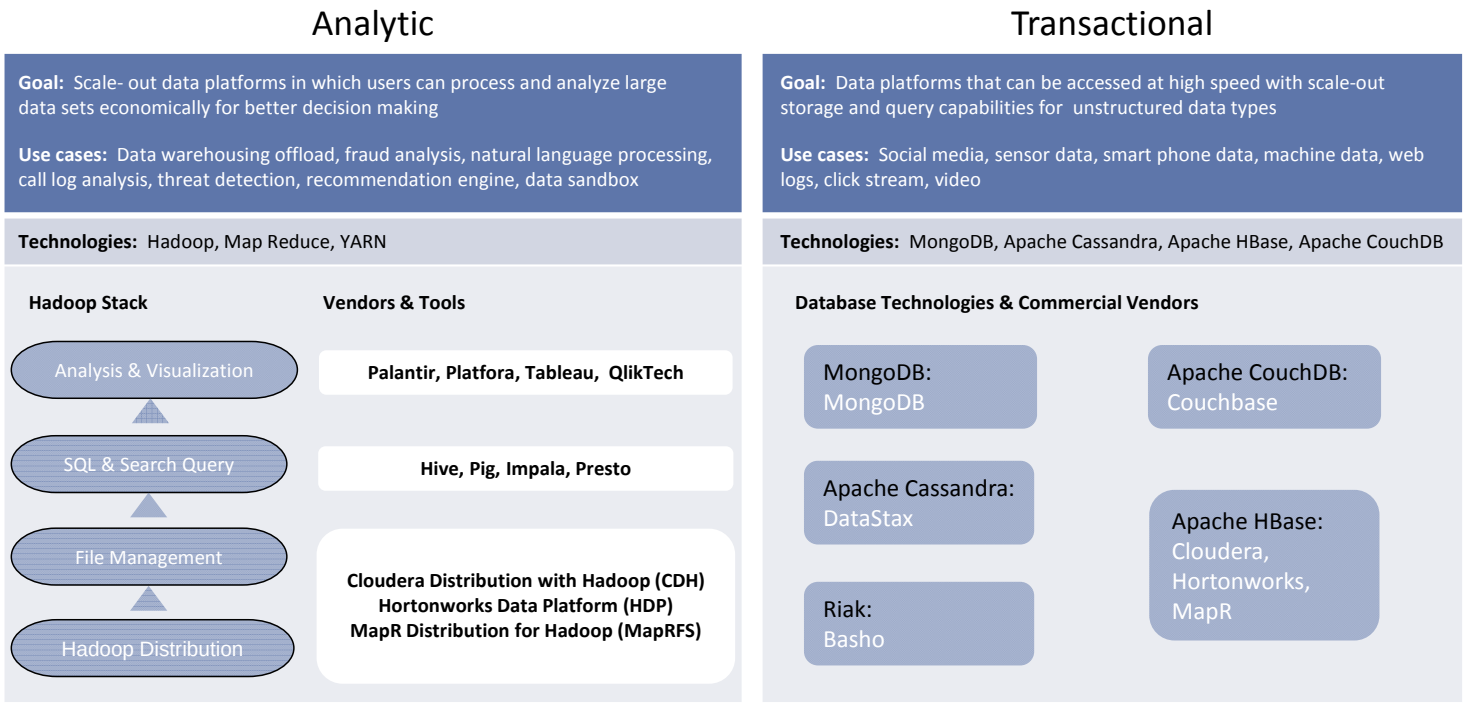

Source: Company data compiled by Goldman Sachs Global Investment Research. This image was originally published on January 28, 2014 in the Goldman Sachs report titled "Big Data: Storm clouds brewing".

Need for more developers and administrators with NoSQL skills

A top challenge for NoSQL adoption is having enough talent available to use this new technology. Compared to relational databases administrators who have decades of experience, NoSQL lacks a large pool of experienced engineers and administrators that can develop on the technology and manage day to day issues. The Exhibit below outlines NoSQL vendors with the highest mentions on LinkedIn profiles.

Exhibit 15: NoSQL Skills Index March 2015

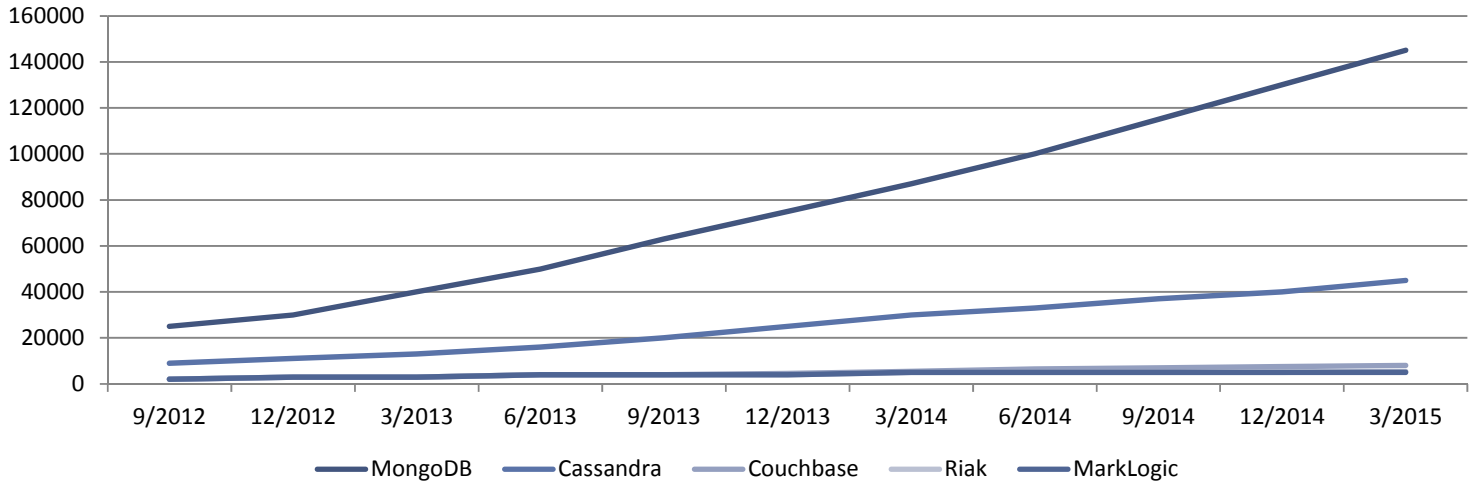

Source: http://blogs.the451group.com/information_management/2015/04/07/nosql-linkedin-skills-index-march-2015/

Capital Raised

For the five largest independent NoSQL vendors by revenue (according to Gartner), we outlined the total capital raised by each company. Four of the top five largest independent NoSQL vendors are Document and/or Key-Value type databases, and make up 0.7% of the CY14 TAM (58% of the CY14 NoSQL market) and 78% of the top five's total capital raised.

Exhibit 16: NoSQL Revenue and Capital Raised

Company	NoSQL Type	2014 Revenue ($mns)	% of total 2014 DBMS TAM	Total Capital Raised ($mns)	% of Top 5
MongoDB	Document/Key-Value	$86	0.28%	$311	37%
MarkLogic	Document	$76	0.25%	$173	20%
Datastax	Column/Graph	$38	0.12%	$190	22%
Couchbase	Key-Value/Document	$24	0.08%	$116	14%
Basho	Key-Value	$17	0.06%	$58	7%

Source: Revenue estimates from Gartner (March 2015), Capital Raised from Crunchbase (May 2015). Compiled by Goldman Sachs Global Investment Research.

Disclosure Appendix

Reg AC

We, Heather Bellini, CFA, Nicole Hayashi, Shateel Alam and Jack Kilgallen, hereby certify that all of the views expressed in this report accurately reflect our personal views about the subject company or companies and its or their securities. We also certify that no part of our compensation was, is or will be, directly or indirectly, related to the specific recommendations or views expressed in this report.

Unless otherwise stated, the individuals listed on the cover page of this report are analysts in Goldman Sachs' Global Investment Research division.

Investment Profile

The Goldman Sachs Investment Profile provides investment context for a security by comparing key attributes of that security to its peer group and market. The four key attributes depicted are: growth, returns, multiple and volatility. Growth, returns and multiple are indexed based on composites of several methodologies to determine the stocks percentile ranking within the region's coverage universe.

The precise calculation of each metric may vary depending on the fiscal year, industry and region but the standard approach is as follows:

Growth is a composite of next year's estimate over current year's estimate, e.g. EPS, EBITDA, Revenue. **Return** is a year one prospective aggregate of various return on capital measures, e.g. CROCI, ROACE, and ROE. **Multiple** is a composite of one-year forward valuation ratios, e.g. P/E, dividend yield, EV/FCF, EV/EBITDA, EV/DACF, Price/Book. **Volatility** is measured as trailing twelve-month volatility adjusted for dividends.

Quantum

Quantum is Goldman Sachs' proprietary database providing access to detailed financial statement histories, forecasts and ratios. It can be used for in-depth analysis of a single company, or to make comparisons between companies in different sectors and markets.

GS SUSTAIN

GS SUSTAIN is a global investment strategy aimed at long-term, long-only performance with a low turnover of ideas. The GS SUSTAIN focus list includes leaders our analysis shows to be well positioned to deliver long term outperformance through sustained competitive advantage and superior returns on capital relative to their global industry peers. Leaders are identified based on quantifiable analysis of three aspects of corporate performance: cash return on cash invested, industry positioning and management quality (the effectiveness of companies' management of the environmental, social and governance issues facing their industry).

Disclosures

Coverage group(s) of stocks by primary analyst(s)

Heather Bellini, CFA: America-Software.

America-Software: Adobe Systems Inc., Akamai Technologies Inc., Autodesk Inc., Citrix Systems Inc., Facebook Inc., Google Inc., Jive Software Inc., Microsoft Corp., MobileIron Inc., Oracle Corp., Red Hat Inc., RingCentral, Salesforce.com Inc., VMware Inc., Workday Inc..

Company-specific regulatory disclosures

The following disclosures relate to relationships between The Goldman Sachs Group, Inc. (with its affiliates, "Goldman Sachs") and companies covered by the Global Investment Research Division of Goldman Sachs and referred to in this research.

Goldman Sachs beneficially owned 1% or more of common equity (excluding positions managed by affiliates and business units not required to be aggregated under US securities law) as of the second most recent month end: SAP (€67.60)

Goldman Sachs has received compensation for investment banking services in the past 12 months: International Business Machines ($170.18), Microsoft Corp. ($47.23), Oracle Corp. ($43.79) and SAP (€67.60)

Goldman Sachs expects to receive or intends to seek compensation for investment banking services in the next 3 months: International Business Machines ($170.18), Microsoft Corp. ($47.23), Oracle Corp. ($43.79), SAP (€67.60) and Teradata Corp. ($38.62)

Goldman Sachs has received compensation for non-investment banking services during the past 12 months: International Business Machines ($170.18), Microsoft Corp. ($47.23), Oracle Corp. ($43.79) and Teradata Corp. ($38.62)

Goldman Sachs had an investment banking services client relationship during the past 12 months with: International Business Machines ($170.18), Microsoft Corp. ($47.23), Oracle Corp. ($43.79) and SAP (€67.60)

Goldman Sachs had a non-investment banking securities-related services client relationship during the past 12 months with: International Business Machines ($170.18), Microsoft Corp. ($47.23), Oracle Corp. ($43.79), SAP (€67.60) and Teradata Corp. ($38.62)

Goldman Sachs had a non-securities services client relationship during the past 12 months with: International Business Machines ($170.18), Microsoft Corp. ($47.23), Oracle Corp. ($43.79), SAP (€67.60) and Teradata Corp. ($38.62)

Goldman Sachs has managed or co-managed a public or Rule 144A offering in the past 12 months: Microsoft Corp. ($47.23)

Goldman Sachs makes a market in the securities or derivatives thereof: International Business Machines ($170.18), Microsoft Corp. ($47.23), Oracle Corp. ($43.79), SAP (€67.60) and Teradata Corp. ($38.62)

Goldman Sachs is a specialist in the relevant securities and will at any given time have an inventory position, "long" or "short," and may be on the opposite side of orders executed on the relevant exchange: International Business Machines ($170.18) and Teradata Corp. ($38.62)

Distribution of ratings/investment banking relationships

Goldman Sachs Investment Research global coverage universe

	Rating Distribution			Investment Banking Relationships		
	Buy	Hold	Sell	Buy	Hold	Sell
Global	32%	54%	14%	46%	37%	32%

As of April 1, 2015, Goldman Sachs Global Investment Research had investment ratings on 3,356 equity securities. Goldman Sachs assigns stocks as Buys and Sells on various regional Investment Lists; stocks not so assigned are deemed Neutral. Such assignments equate to Buy, Hold and Sell for the purposes of the above disclosure required by NASD/NYSE rules. See 'Ratings, Coverage groups and views and related definitions' below.

Price target and rating history chart(s)

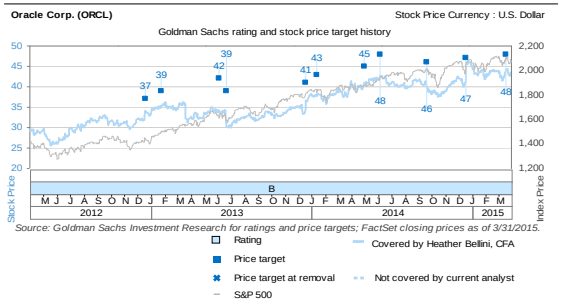

The price targets shown should be considered in the context of all prior published Goldman Sachs research, which may or may not have included price targets, as well as developments relating to the company, its industry and financial markets.

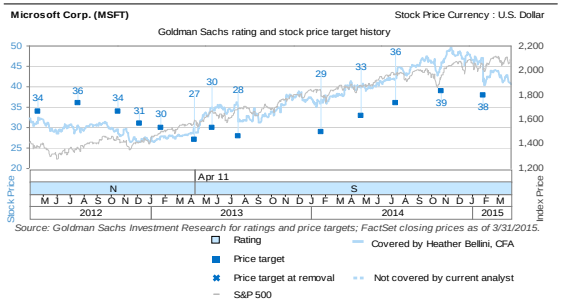

The price targets shown should be considered in the context of all prior published Goldman Sachs research, which may or may not have included price targets, as well as developments relating to the company, its industry and financial markets.

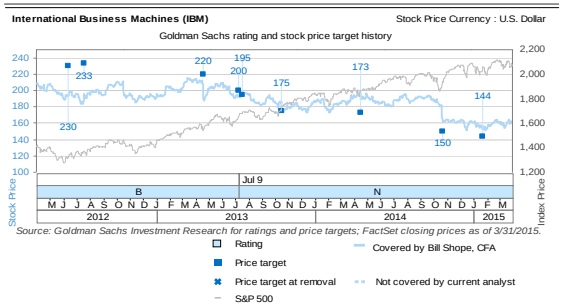

The price targets shown should be considered in the context of all prior published Goldman Sachs research, which may or may not have included price targets, as well as developments relating to the company, its industry and financial markets.

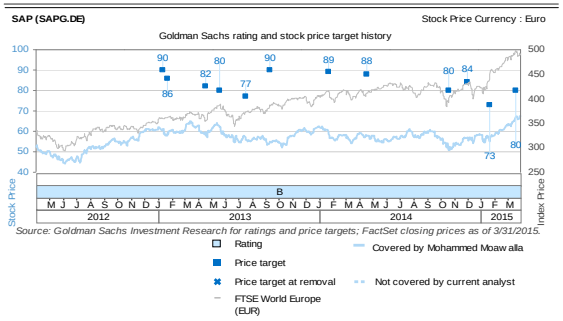

The price targets shown should be considered in the context of all prior published Goldman Sachs research, which may or may not have included price targets, as well as developments relating to the company, its industry and financial markets.

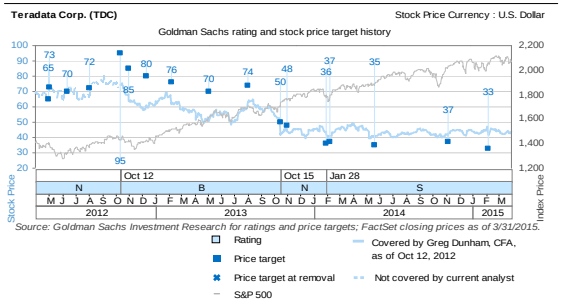

The price targets shown should be considered in the context of all prior published Goldman Sachs research, which may or may not have included price targets, as well as developments relating to the company, its industry and financial markets.

Regulatory disclosures

Disclosures required by United States laws and regulations

See company-specific regulatory disclosures above for any of the following disclosures required as to companies referred to in this report: manager or co-manager in a pending transaction; 1% or other ownership; compensation for certain services; types of client relationships; managed/co-managed public offerings in prior periods; directorships; for equity securities, market making and/or specialist role. Goldman Sachs usually makes a market in fixed income securities of issuers discussed in this report and usually deals as a principal in these securities.

The following are additional required disclosures: **Ownership and material conflicts of interest:** Goldman Sachs policy prohibits its analysts, professionals reporting to analysts and members of their households from owning securities of any company in the analyst's area of coverage. **Analyst compensation:** Analysts are paid in part based on the profitability of Goldman Sachs, which includes investment banking revenues. **Analyst as officer or director:** Goldman Sachs policy prohibits its analysts, persons reporting to analysts or members of their households from serving as an officer, director, advisory board member or employee of any company in the analyst's area of coverage. **Non-U.S. Analysts:** Non-U.S. analysts may not be associated persons of Goldman, Sachs & Co. and therefore may not be subject to NASD Rule 2711/NYSE Rules 472 restrictions on communications with subject company, public appearances and trading securities held by the analysts.

Distribution of ratings: See the distribution of ratings disclosure above. **Price chart:** See the price chart, with changes of ratings and price targets in prior periods, above, or, if electronic format or if with respect to multiple companies which are the subject of this report, on the Goldman Sachs website at http://www.gs.com/research/hedge.html.

Additional disclosures required under the laws and regulations of jurisdictions other than the United States

The following disclosures are those required by the jurisdiction indicated, except to the extent already made above pursuant to United States laws and regulations. **Australia:** Goldman Sachs Australia Pty Ltd and its affiliates are not authorised deposit-taking institutions (as that term is defined in the Banking Act 1959 (Cth)) in Australia and do not provide banking services, nor carry on a banking business, in Australia. This research, and any access to it, is intended only for "wholesale clients" within the meaning of the Australian Corporations Act, unless otherwise agreed by Goldman Sachs. In producing research reports, members of the Global Investment Research Division of Goldman Sachs Australia may attend site visits and other meetings hosted by the issuers the subject of its research reports. In some instances the costs of such site visits or meetings may be met in part or in whole by the issuers concerned if Goldman Sachs Australia considers it is appropriate and reasonable in the specific circumstances relating to the site visit or meeting. **Brazil:** Disclosure information in relation to CVM Instruction 483 is available at http://www.gs.com/worldwide/brazil/area/gir/index.html. Where applicable, the Brazil-registered analyst primarily responsible for the content of this research report, as defined in Article 16 of CVM Instruction 483, is the first author named at the beginning of this report, unless indicated otherwise at the end of the text. **Canada:** Goldman Sachs Canada Inc. is an affiliate of The Goldman Sachs Group Inc. and therefore is included in the company specific disclosures relating to Goldman Sachs (as defined above). Goldman Sachs Canada Inc. has approved of, and agreed to take responsibility for, this research report in Canada if and to the extent that Goldman Sachs Canada Inc. disseminates this research report to its clients. **Hong Kong:** Further information on the securities of covered companies referred to in this research may be obtained on request from Goldman Sachs (Asia) L.L.C. **India:** Further information on the subject company or companies referred to in this research may be obtained from Goldman Sachs (India) Securities Private Limited. Goldman Sachs may beneficially own 1% or more of the securities (as such term is defined in clause 2 (h) the Indian Securities Contracts (Regulation) Act, 1956) of the subject company or companies referred to in this research report. **Japan:** See below. **Korea:** Further information on the subject company or companies referred to in this research may be obtained from Goldman Sachs (Asia) L.L.C., Seoul Branch. **New Zealand:** Goldman Sachs New Zealand Limited and its affiliates are neither "registered banks" nor "deposit takers" (as defined in the Reserve Bank of New Zealand Act 1989) in New Zealand. This research, and any access to it, is intended for "wholesale clients" (as defined in the Financial Advisers Act 2008) unless otherwise agreed by Goldman Sachs. **Russia:** Research reports distributed in the Russian Federation are not advertising as defined in the Russian legislation, but are information and analysis not having product promotion as their main purpose and do not provide appraisal within the meaning of the Russian legislation on appraisal activity. **Singapore:** Further information on the covered companies referred to in this research may be obtained from Goldman Sachs (Singapore) Pte. (Company Number: 198602165W). **Taiwan:** This material is for reference only and must not be reprinted without permission. Investors should carefully consider their own investment risk. Investment results are the responsibility of the individual investor. **United Kingdom:** Persons who would be categorized as retail clients in the United Kingdom, as such term is defined in the rules of the Financial Conduct Authority, should read this research in conjunction with prior Goldman Sachs research on the covered companies referred to herein and should refer to the risk warnings that have been sent to them by Goldman Sachs International. A copy of these risks warnings, and a glossary of certain financial terms used in this report, are available from Goldman Sachs International on request.

European Union: Disclosure information in relation to Article 4 (1) (d) and Article 6 (2) of the European Commission Directive 2003/126/EC is available at http://www.gs.com/disclosures/europeanpolicy.html which states the European Policy for Managing Conflicts of Interest in Connection with Investment Research.

Japan: Goldman Sachs Japan Co., Ltd. is a Financial Instrument Dealer registered with the Kanto Financial Bureau under registration number Kinsho 69, and a member of Japan Securities Dealers Association, Financial Futures Association of Japan and Type II Financial Instruments Firms Association. Sales and purchase of equities are subject to commission pre-determined with clients plus consumption tax. See company-specific disclosures as to any applicable disclosures required by Japanese stock exchanges, the Japanese Securities Dealers Association or the Japanese Securities Finance Company.

Ratings, coverage groups and views and related definitions

Buy (B), Neutral (N), Sell (S) -Analysts recommend stocks as Buys or Sells for inclusion on various regional Investment Lists. Being assigned a Buy or Sell on an Investment List is determined by a stock's return potential relative to its coverage group as described below. Any stock not assigned as a Buy or a Sell on an Investment List is deemed Neutral. Each regional Investment Review Committee manages various regional Investment Lists to a global guideline of 25%-35% of stocks as Buy and 10%-15% of stocks as Sell; however, the distribution of Buys and Sells in any particular coverage group may vary as determined by the regional Investment Review Committee. Regional Conviction Buy and Sell lists represent investment recommendations focused on either the size of the potential return or the likelihood of the realization of the return.

Return potential represents the price differential between the current share price and the price target expected during the time horizon associated with the price target. Price targets are required for all covered stocks. The return potential, price target and associated time horizon are stated in each report adding or reiterating an Investment List membership.

Coverage groups and views: A list of all stocks in each coverage group is available by primary analyst, stock and coverage group at http://www.gs.com/research/hedge.html. The analyst assigns one of the following coverage views which represents the analyst's investment outlook on the coverage group relative to the group's historical fundamentals and/or valuation. **Attractive (A).** The investment outlook over the following 12 months is favorable relative to the coverage group's historical fundamentals and/or valuation. **Neutral (N).** The investment outlook over the following 12 months is neutral relative to the coverage group's historical fundamentals and/or valuation. **Cautious (C).** The investment outlook over the following 12 months is unfavorable relative to the coverage group's historical fundamentals and/or valuation.

Not Rated (NR). The investment rating and target price have been removed pursuant to Goldman Sachs policy when Goldman Sachs is acting in an advisory capacity in a merger or strategic transaction involving this company and in certain other circumstances. **Rating Suspended (RS).** Goldman Sachs Research has suspended the investment rating and price target for this stock, because there is not a sufficient fundamental basis for determining, or there are legal, regulatory or policy constraints around publishing, an investment rating or target. The previous investment rating and price target, if any, are no longer in effect for this stock and should not be relied upon. **Coverage Suspended (CS).** Goldman Sachs has suspended coverage of this company. **Not Covered (NC).** Goldman Sachs does not cover this company. **Not Available or Not Applicable (NA).** The information is not available for display or is not applicable. **Not Meaningful (NM).** The information is not meaningful and is therefore excluded.

Global product; distributing entities

The Global Investment Research Division of Goldman Sachs produces and distributes research products for clients of Goldman Sachs on a global basis. Analysts based in Goldman Sachs offices around the world produce equity research on industries and companies, and research on macroeconomics, currencies, commodities and portfolio strategy. This research is disseminated in Australia by Goldman Sachs Australia Pty Ltd (ABN 21 006 797 897); in Brazil by Goldman Sachs do Brasil Corretora de Títulos e Valores Mobiliários S.A.; in Canada by either Goldman Sachs Canada Inc. or Goldman, Sachs & Co.; in Hong Kong by Goldman Sachs (Asia) L.L.C.; in India by Goldman Sachs (India) Securities Private Ltd.; in Japan by Goldman Sachs Japan Co., Ltd.; in the Republic of Korea by Goldman Sachs (Asia) L.L.C., Seoul Branch; in New Zealand by Goldman Sachs New Zealand Limited; in Russia by OOO Goldman Sachs; in Singapore by Goldman Sachs (Singapore) Pte. (Company Number: 198602165W); and in the United States of America by Goldman, Sachs & Co. Goldman Sachs International has approved this research in connection with its distribution in the United Kingdom and European Union.

European Union: Goldman Sachs International authorised by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority, has approved this research in connection with its distribution in the European Union and United Kingdom; Goldman Sachs AG and Goldman Sachs International Zweigniederlassung Frankfurt, regulated by the Bundesanstalt für Finanzdienstleistungsaufsicht, may also distribute research in Germany.

General disclosures

This research is for our clients only. Other than disclosures relating to Goldman Sachs, this research is based on current public information that we consider reliable, but we do not represent it is accurate or complete, and it should not be relied on as such. We seek to update our research as appropriate, but various regulations may prevent us from doing so. Other than certain industry reports published on a periodic basis, the large majority of reports are published at irregular intervals as appropriate in the analyst's judgment.

Goldman Sachs conducts a global full-service, integrated investment banking, investment management, and brokerage business. We have investment banking and other business relationships with a substantial percentage of the companies covered by our Global Investment Research Division. Goldman, Sachs & Co., the United States broker dealer, is a member of SIPC (http://www.sipc.org).

Our salespeople, traders, and other professionals may provide oral or written market commentary or trading strategies to our clients and our proprietary trading desks that reflect opinions that are contrary to the opinions expressed in this research. Our asset management area, our proprietary trading desks and investing businesses may make investment decisions that are inconsistent with the recommendations or views expressed in this research.

The analysts named in this report may have from time to time discussed with our clients, including Goldman Sachs salespersons and traders, or may discuss in this report, trading strategies that reference catalysts or events that may have a near-term impact on the market price of the equity securities discussed in this report, which impact may be directionally counter to the analyst's published price target expectations for such stocks. Any such trading strategies are distinct from and do not affect the analyst's fundamental equity rating for such stocks, which rating reflects a stock's return potential relative to its coverage group as described herein.

We and our affiliates, officers, directors, and employees, excluding equity and credit analysts, will from time to time have long or short positions in, act as principal in, and buy or sell, the securities or derivatives, if any, referred to in this research.

The views attributed to third party presenters at Goldman Sachs arranged conferences, including individuals from other parts of Goldman Sachs, do not necessarily reflect those of Global Investment Research and are not an official view of Goldman Sachs.

Any third party referenced herein, including any salespeople, traders and other professionals or members of their household, may have positions in the products mentioned that are inconsistent with the views expressed by analysts named in this report.

This research is not an offer to sell or the solicitation of an offer to buy any security in any jurisdiction where such an offer or solicitation would be illegal. It does not constitute a personal recommendation or take into account the particular investment objectives, financial situations, or needs of individual clients. Clients should consider whether any advice or recommendation in this research is suitable for their particular circumstances and, if appropriate, seek professional advice, including tax advice. The price and value of investments referred to in this research and the income from them may fluctuate. Past performance is not a guide to future performance, future returns are not guaranteed, and a loss of original capital may occur. Fluctuations in exchange rates could have adverse effects on the value or price of, or income derived from, certain investments.

Certain transactions, including those involving futures, options, and other derivatives, give rise to substantial risk and are not suitable for all investors. Investors should review current options disclosure documents which are available from Goldman Sachs sales representatives or at http://www.theocc.com/about/publications/character-risks.jsp. Transaction costs may be significant in option strategies calling for multiple purchase and sales of options such as spreads. Supporting documentation will be supplied upon request.

All research reports are disseminated and available to all clients simultaneously through electronic publication to our internal client websites. Not all research content is redistributed to our clients or available to third-party aggregators, nor is Goldman Sachs responsible for the redistribution of our research by third party aggregators. For research, models or other data available on a particular security, please contact your sales representative or go to http://360.gs.com.

Disclosure information is also available at http://www.gs.com/research/hedge.html or from Research Compliance, 200 West Street, New York, NY 10282.