G00252553

Hype Cycle for Business Intelligence and Analytics, 2013

Published: 31 July 2013

Analyst(s): Kurt Schlegel

This Hype Cycle will help business intelligence and analytics leaders prioritize their investments in the emerging technologies that will make their BI and analytics initiatives more precise, transparent and decisive.

Table of Contents

List of Tables

List of Figures

Analysis

What You Need to Know

Business intelligence (BI) and analytics leaders should use this Hype Cycle to help prioritize their investments in BI and analytics initiatives along with its four siblings:

"Hype Cycle for Performance Management, 2013"

"Hype Cycle for Information Infrastructure, 2013"

"Hype Cycle for Analytic Applications, 2013"

"Hype Cycle for Big Data, 2013"

A small number of technologies are intentionally repeated in multiple Hype Cycles if they are considered important from more than one perspective. Technologies on the left-hand side of the Hype Cycle, pre- trough, should be looked at as opportunities to gain a competitive advantage or solve a particular business problem by leveraging an emerging technology. Those on the right-hand side of the Hype Cycle, post trough, should be viewed as almost automatic investments.

The Hype Cycle

The BI and analytics space, although very mature, keeps reinventing itself. While the core technologies have been around for a few decades, there are always new capabilities. Ten new entries were added to the Hype Cycle this year. Most of these are focused on advanced analytic technologies, such as Data Science and Model Managers, or capabilities that enable us to uncover relationships hidden in large datasets, such as Graph Databases or Entity Resolution and Analysis. Many of the new technologies, such as Operational Intelligence Platforms, Prescriptive Analytics, Intelligent Business Operations and Complex-Event Processing, are focused on making BI and analytics more actionable. Some, such as Quantified Self, introduce entirely new ways of measuring and classifying our personal lives.

Gartner's Nexus of Forces is a useful guide to the dominant trends across all of IT in general and BI and analytics in particular:

- **Information** is obviously at the core of everything BI and analytics represents. Many of the entries in this Hype Cycle, such as the Information Capabilities Framework, Big Data, Enterprise Metadata, and Logical Data Warehouse, introduce creative ways to manage information that balance governance with agility.
- **Mobile** represents the new consumerization trend of delivering BI to tablets and smartphone devices. In addition to the Mobile BI entry on this Hype Cycle, Search-Based Data Discovery Tools and Interactive Visualization provide a much more consumer-focused BI and analytics experience than traditional tools. Mobile also creates new ways to collect and contextualize information.
- **Cloud** indicates the growing, but still rather embryonic, adoption of BI and analytics in the SaaS delivery model — particularly for packaged analytic applications. One of the more interesting cloud trends in the BI and analytics space is the increasing tendency for end-user organizations to provide analytics to their own customers, partners and suppliers through a hosted service.
- **Social** has a double impact on the BI and analytics space. First, social provides new subject areas to report and analyze. For example, marketers can now analyze sentiment analysis of relevant brands by monitoring social media sites. Second, social augments existing BI and

analytics technologies with more collaboration (such as Collaborative Decision Making), making it more actionable and decisive.

Overview

In the last few years, we have seen organizations supporting emerging reporting and analysis requirements using a variety of new information management architectures such as the Logical Data Warehouse and data discovery (see "The Future of Data Management for Analytics is the Logical Data Warehouse" and "Market Trends: The Collision of Data Discovery and Business Intelligence Will Cause Destruction").

Increasingly, data managed for analytics won't only come from a consolidated and modeled repository (a traditional data warehouse). Sometimes, the traditional data warehouse will coexist with distributed data in various source systems with a combination of formal data models and distributed processing (the Hadoop Distributed File System and MapReduce, for example) — that is, the Logical Data Warehouse. At other times, the data will be centralized but in a non-modeled repository (Data Discovery). In addition, information management will invest heavily in text analytics, NoSQL, and search capabilities to analyze unstructured information (documents) and semi-structured (log files) information in conjunction with more structured data (relational).

In addition to the information management changes discussed above, the BI and analytics space will increasingly focus on advanced capabilities to build sophisticated analytics models that go well beyond the slicing and dicing of multidimensional data. In particular, the role of Data Science will emerge as companies invest in overcoming their lack of analytical skills and innovate ways to apply analytical insights to solve business problems.

The technologies and services discussed in this Hype Cycle will result in three profound changes in how business is conducted:

1. **More precise and actionable systems of measurement and classification.** The ability to handle increasing volume (using Column-Store DBMS, In-Memory Analytics) and variety (using Text Analytics or Search-Based Data Discovery, for example) of data enables analysts to build more exact systems of measurement and to classify customers and other business entities with more precision than traditional BI systems that typically relied on summarized data. This precision will lead to more personalized business processes (See "Precision Is the Future of Analytics").
2. **Improved transparency to customers, partners, and suppliers**. Business analytics delivered via SaaS is one of the technologies that are increasingly being leveraged by companies to deliver externally facing analytics to customers, partners and suppliers. Coupled with other consumer-friendly technologies, such as Interactive Visualization, Mobile BI, and Dashboards, information that was traditionally reported internally will become much more transparent to external stakeholders (see "Make Customer-Facing Analytics Part of Your Business Model").
3. **Decisive analytical systems that go beyond typical performance measurement by embedding decision support into information systems**. Improved decision making is perennially listed as a key reason that companies invest in BI and analytics. Several of the entries in this Hype Cycle, including Collaborative Decision Making, Prescriptive Analytics,

Complex-Event Processing, and Real-Time Decisioning Platform, are squarely focused on delivering decision support systems (see "Advanced Analytics Enables Real-Time Business Optimization").

Technology Trigger

Over half of the technologies and services covered in this Hype Cycle are bunched into the Innovation Trigger segment. Most of these are not widely known in the market. Some, such as Collaborative Decision Making, have been around a few years, and have been promoted by Gartner, but haven't really gained much traction in the market. Others, like Graph Databases — which is new to the Hype Cycle this year, are only known in very technical circles but will likely become more widely known as the demand for associative data models becomes more prevalent. Quantified Self is a newly minted term but has the potential to become mainstream very quickly, particularly with the plethora of personal monitoring devices (including health monitors) now available.

Figure 1. Hype Cycle for Business Intelligence and Analytics, 2013

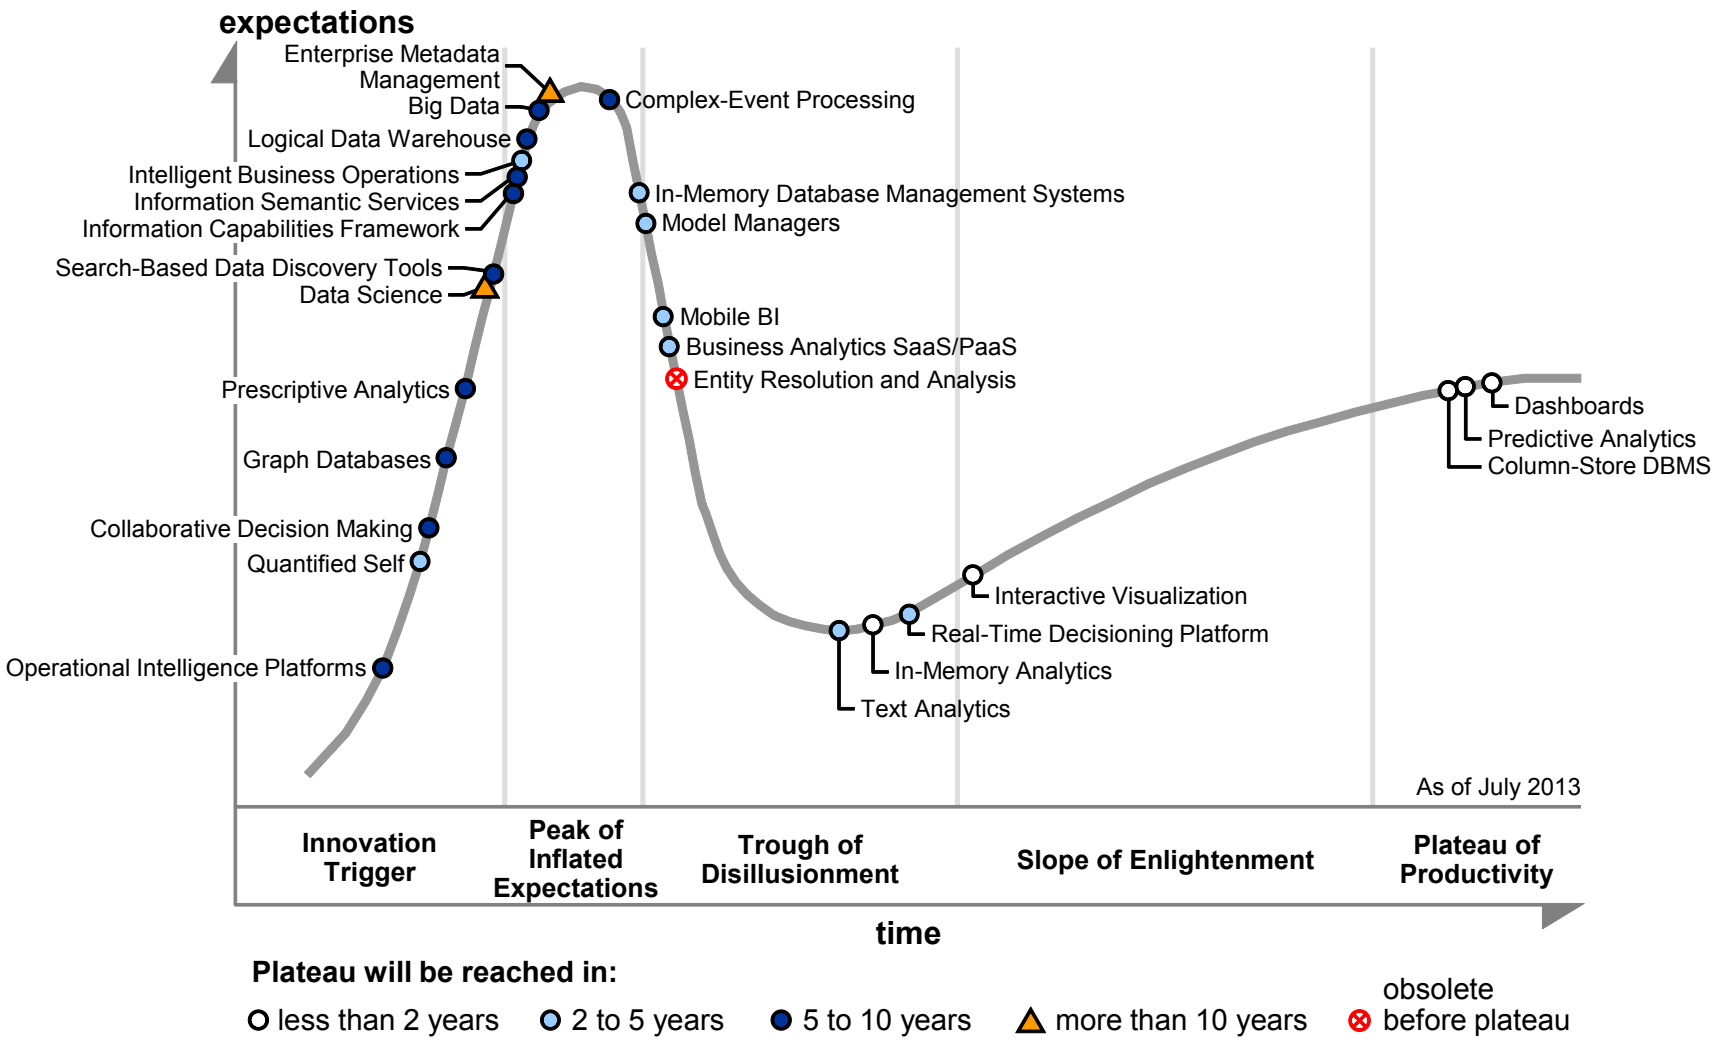

BI = business intelligence; DBMS = database management system; SaaS = software as a service; PaaS = platform as a service

Source: Gartner (July 2013)

The Priority Matrix

The Priority Matrix shows technologies in the following categories:

- **Transformational benefit, less than two years to mainstream adoption.** These technologies are ready to be deployed today and provide instant benefits. There are no technologies in this category.
- **Transformational benefit, two to five years to mainstream adoption.** These technologies are fundamentally changing the way BI is consumed, in particular by increasing the speed at which information is delivered. In-Memory Database Management Systems have the potential to impact virtually every IT application. In particular, the ability to recalculate variables (such as dynamic pricing) based on new information will change business processes making them more agile. Intelligent Business Operations also has the potential to significantly change business processes by embedding BI directly into the business process.
- **Transformational benefit, five to 10 years to mainstream adoption.** These technologies are aimed at changing the way organizations interact with information. Because of a required cultural change, Collaborative Decision Making or Text Analytics will take a while before broad adoption, but will have a transformational influence on how users consume, share and analyze information. While it will take time, several technologies including Big Data, Collaborative Decision Making, Complex-Event Processing, Graph Databases, and the Information Capabilities Framework are poised to have a transformational impact on certain business processes.
- **Transformational benefit, more than 10 years to mainstream adoption.** These technologies can enable a transformational change to an organization's BI initiatives. Because of the high complexity of delivering information infrastructures, investments in this area are more strategic in value than tactical. This Hype Cycle has no technologies in this area.
- **High benefit, less than two years to mainstream adoption.** These technologies are, without exception, mature and on, or approaching, the Plateau of Productivity. These technologies, such as Dashboards, In-Memory Analytics, Interactive Visualization and Predictive Analytics, have reached maturity and promise a very quick business benefit.
- **High benefit, two to five years to mainstream adoption.** These technologies, such as Text Analytics, Real-Time Decisioning Platform and Quantified Self, will take time to mature. Of the three, Text Analytics is probably the most mature. We are seeing Real-Time Decisioning Platform in very specific use cases, particularly in the financial services industry. Quantified Self is a largely unknown term that has just emerged on the market.
- **High benefit, five to 10 years to mainstream adoption.** These technologies are more strategic and require long-term planning and incremental investment. Technologies in this group enable improvements to broad BI and analytics initiatives in various domains. They include new concepts such as the Logical Data Warehouse, Prescriptive Analytics, Operational Intelligence Platforms and Search-Based Data Discovery Tools.

- **High benefit, more than 10 years to mainstream adoption.** These technologies are very strategic in nature, require long-term thinking and are beyond most organizations' planning horizons. Enterprise Metadata Management and Data Science are found in this group.
- **Moderate benefit, less than two years to mainstream adoption.** These technologies are widely used, popular among end users, but not necessarily highly impactful. There are no technologies on this Hype Cycle in this category.
- **Moderate benefit, two to five years or five to 10 years to mainstream adoption.** These technologies add incremental benefits to existing BI deployments. Depending on the time frame, they can be implemented with existing tools (for example, Mobile BI) or a new delivery model (for example, Business Analytics SaaS). Further out on the horizon we expect Information Semantic Services to reach higher adoption.
- **Low benefit.** Technologies in this group are not viewed as large contributors to BI and analytics initiatives. Therefore this Hype Cycle has no technologies in this area.

To help organizations prioritize investments in relation to their level of impact we provide a Priority Matrix (see Figure 2), although impact is not the only consideration when selecting vendors and products — applicability, budget, time to implement and receive payback, and return on investment are also important considerations. Vendors are enabling BI with a wide variety of technologies, in different ways and at different rates. The Priority Matrix shows the degree of benefit attainable relative to a technology's progression along the Hype Cycle. It is intended as a general guide because the benefits from, and the maturity of, any technology depend partly on industry conditions and the organization's ability to use the technology effectively.

Figure 2. Priority Matrix for Business Intelligence and Analytics, 2013

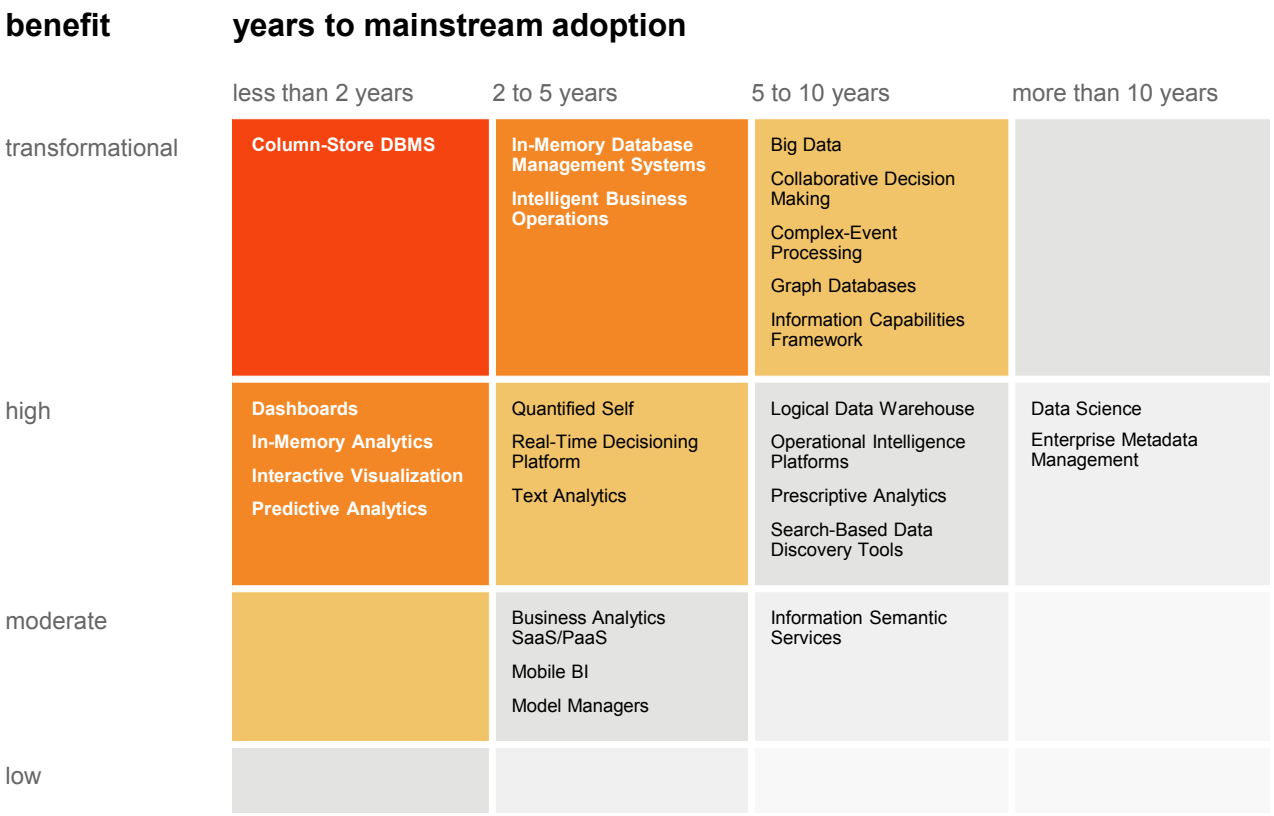

BI = business intelligence; DBMS = database management system; SaaS = software as a service; PaaS = platform as a service

Source: Gartner (July 2013)

Off The Hype Cycle

Open Source BI was removed from the Hype Cycle because this is increasingly less relevant to the BI and analytics leader. Most of the BI and analytics platform vendors don't even market themselves as open-source vendors, but are more focused on the growing need for embedded BI and analytics.

Business activity monitoring (BAM) has moved off the Hype Cycle in 2013 because of the declining use of the term. BAM originally referred to the general notion of real-time and near-real-time monitoring of any business activity, although some people thought of it in a narrower sense: meaning monitoring a business process that was being orchestrated by a workflow or process-orchestration tool. The number of companies that are monitoring their business activities in real time or near real time is increasing rapidly, but users and vendors are now more likely to use the general term "operational intelligence" or a purpose-specific label such as business process monitoring, call

center monitoring, supply chain visibility, manufacturing operations intelligence or some other term to describe the activity. The most appropriate term for the general category of tools for monitoring, alerting and adaptive decision making is now "Operational Intelligence Platform."

Excel as a BI Front End was also removed from the Hype Cycle this year. Gartner rarely receives inquiries about this topic and, while it is still a useful capability for certain use cases, there isn't a lot of innovation in this area. We decided to remove it from the Hype Cycle to make room for the plethora of new technologies that fit the story around big data and advanced analytics of this Hype Cycle.

Data Mining Workbenches was also removed because it matured off the Plateau of Productivity. It had been on the edge of the plateau for years.

Business Intelligence Platforms was removed because it also matured off the Plateau of Productivity. It had been on the edge of the plateau for years.

On the Rise

Operational Intelligence Platforms

Analysis By: W. Roy Schulte

Definition: An operational intelligence platform is a set of development and runtime tools that enables applications that monitor, alert or support adaptive decision making, based on current conditions. These platforms have adapters to ingest event data; context stores to hold context data; logic to find patterns and detect anomalies and other threats and opportunities; decision management capabilities, such as rules or analytics; interactive dashboards; alerting facilities; and capabilities to trigger responses in applications, devices or workflow tools.

Position and Adoption Speed Justification: Operational intelligence platform is a new entry on the Hype Cycle. It is at the Technology Trigger stage and still relatively unknown. It will take several years to reach the Peak of Inflated Expectations and up to 10 years to reach the Plateau of Productivity, at which point over half the potential users will be using this technology.

The term "operational intelligence platform" was coined in 2012 because there was no general label for software that supports applications for continuous (or regularly recurring) monitoring, alerting or adaptive decision making (see "Commercial Operational Intelligence Platforms Are Coming to Market"). However, the concept is not entirely new as some kinds of operational intelligence platform have been used in particular industries or applications for many years. They are known by labels specific to their purposes, which include supply chain visibility, business process monitoring, business activity monitoring, manufacturing operations intelligence (see "Hype Cycle for Manufacturing Product Life Cycle and Operations Management, 2012"), customer contact center monitoring and truck fleet management. "Operational intelligence platform" is a broad category that encompasses those established purpose-specific monitoring systems, as well as newly emerging monitoring systems for other industries and applications. It also encompasses general-purpose products that can be tailored to virtually any business scenario for which there is no suitable commercial, off-the-shelf, purpose-specific monitoring and alerting system.

Operational intelligence platforms are often, but not always, used for real-time or near-real-time decisions. Real-time and near-real-time decisions are those in which at least some of the data used for making decisions has been generated within the previous 15 minutes. Other decisions are not real-time because all the data is more than 15 minutes old (it may be hours or days old). When operational intelligence platforms are used for real-time or near-real-time purposes, and for decision automation rather than decision support, they are adopting the role of a Real-Time Decisioning Platform.

User Advice: Architects and analysts should use operational intelligence platforms to provide a new layer of oversight and insight into processes and operations. The platform is sometimes used to enable panoramic (or "360 degree") monitoring applications that cover end-to-end processes or multiple aspects of an operation that previously had no continuous (or regularly recurring) monitoring, or where the monitoring was limited to narrow "keyhole" views of individual application systems and devices. Operational intelligence platforms should do more than provide visibility; they should also be used to sense variances from a baseline of historical, expected or desired activity. They should send alerts to people (as part of decision support) about business threats or opportunities as they occur. In some cases, these platforms will trigger responses automatically (decision automation).

Architects and analysts should work with business managers and users to identify what data needs to be maintained in the context store, and how long it should be held. A context store can integrate information from multiple applications, sensors or other event sources inside and outside a company. In some cases, it is used to track each instance of a business process through its life cycle, even if the process is not managed by a workflow, business process management (BPM) or business process management suite (BPMS) tool. Context information is continuously (or at least frequently) refreshed so that decisions made by people or application systems reflect the most current and complete view of conditions.

Operational intelligence platforms are one of the technologies well suited to implementing intelligent business operation strategies. They are relevant when retrofitting monitoring, alerting or adaptive decision-making capabilities on top of processes or operations that will not be replaced or fundamentally modified, or when implementing monitoring systems across heterogeneous operations and processes that span multiple business units or disparate application packages. However, architects should consider intelligent BPMS (iBPMS) products, instead of operational intelligence platforms, when implementing all-new business processes or making fundamental changes to processes. iBPMS products also have monitoring, alerting and decision-making capabilities, so they overlap operational intelligence platforms. iBPMSs tend to have superior workflow, adaptive case management and structured process orchestration capabilities, but inferior monitoring and alerting capabilities, although these are generalizations and do not apply to all products (see "Commercial Operational Intelligence Platforms Are Coming to Market").

Business Impact: The benefits of operational intelligence platforms differ, depending on the industry and business function to which they are applied. In general, the platforms give business people broader, more holistic and more current views into their operations. Traditional business operations are managed using daily, weekly or monthly reports without an operational intelligence platform (or iBPMS or custom-built monitoring and alerting system). Business people have no way

to monitor the up-to-the-minute status of end-to-end business processes, and have only limited visibility into conditions in adjacent parts of the business or the outside world. Traditional business intelligence systems are set up only to report results, not to detect anomalies or other threats and opportunities.

Operational intelligence platforms operationalize objectives such as reducing customer churn, expanding sales and reducing the cost of manufacturing or delivery. They help people share a common operating picture, which improves their ability to collaborate when making business decisions. They improve the quality of decisions by providing more contextual information and offloading mathematically based business logic from people and application programs, so that decisions are better informed and more precise.

Benefit Rating: High

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Sample Vendors: Access Intelligence; Aha Software; Axway; BusinessPort; C3global; ClearPriority; FeedZai; Greenlight Technologies; Intelligent InSites; JackBe; Kofax; Lavastorm Analytics; Oversight Systems; Rockshore; SAP; Software AG; Splunk; Systar; Vitria; West Global

Recommended Reading: "Commercial Operational Intelligence Platforms Are Coming to Market"

"Use Intelligent Business Operations to Create Business Advantage"

Quantified Self

Analysis By: Frank Buytendijk; Whit Andrews; Svetlana Sicular

Definition: Originating in San Francisco, Quantified Self is a movement promoting the use of self-monitoring through a wide variety of sensors and devices. It uses mobile apps and wearable or portable devices to collect data about a user's activities, biometrics, environment and other personal experiences. Analysis of this data allows individuals to gain a better understanding of their experiences and improve their wellbeing. Integration with social media allows users to connect with peers, share information, gain community support and learn from others.

Position and Adoption Speed Justification: The Quantified Self movement, currently expanding into Europe with conferences and new chapters, is a trigger for the socialization of new types of technology. However, it will take two to five years before these are adopted by the mainstream.

Although there are multiple types of applications, the most successful commercial implementations can be found in sports and health. There are thousands of health-related tools and supporting apps in smartphone app stores. Commercially-available motion-tracking devices include Nike+, Fitbit, Amiigo, Basis, Withings, BodyMedia and Jawbone Up. Although the areas of application are quickly expanding, self-monitoring currently focuses on motion trackers and vital-sign monitoring (blood pressure and heart rate), but some manufacturers and developers are already talking about mood monitoring.

The range of devices itself is developing quickly as well. Currently based on wristbands and reported through smartphones, various high-tech companies including Apple, Google and Samsung have announced smart watches and devices that can record visual and audio, while displaying information on displays built into glasses. Numerous startup initiatives are focusing on wearable computing with sensors in clothing. There are many other objects that are being turned into monitoring devices. One bra is already fitted out with sensors that detect the first signs of breast cancer. Or consider slippers with balance sensors and carpet sensors that track movement, or Dr. Toilet that monitors stool and urine samples, just to name a few. The sudden popularity of these devices, and the immaturity of the technology, can sometimes cause stability and quality issues.

User Advice: The Quantified Self movement isn't mainstream yet, but the number of personal devices that collect data and provide feedback to users is increasing. While people under the age of 30 are interested in sharing this information, people over 40 are seeking self-awareness and medical insight. The ability to tap into the uncharted depths of customer intimacy through offering personal analytics is also attracting the interest of marketers across all industries. Marketing has never had the opportunity of being so personal and urgent.

Marketers should, however, be careful what they wish for as users take their analytics perhaps a little more personally than you'd like. With the advent of personal analytics, consumers feel that they own the data they collect, whether that data is subject to T&Cs or not. If you violate that sentiment and the data is used for any reason other than those with direct consumer benefit, protests and subsequent reputation damage may occur. Over the past two years several cases (involving Internet providers, retailers, telecom providers, banks and consumer device makers) have made it to the newspapers.

Marketers that do it right, and focus on personal analytics with a promise to never sell that data, will get more than a loyal customer; they will get their whole network. When the data benefits the person that collects it, users will be likely to seek out like-minded companions in an effort to make their lives better. They will find people who share their interests and they will become ambassadors for the products and services by suggesting that others join; not just for fun, but to create even richer comparisons for themselves, to share again with others and, not incidentally, with any company they feel adds value.

Business Impact: As more people use mobile and social technologies to collect and assemble data about themselves and their immediate surroundings, business opportunities emerge to facilitate the process or exploit the collected data. Enterprises may take various steps to catalyze or benefit from such trends.

They may choose to:

- Create new devices or applications that generate revenue streams through subscriptions or advertising.
- Create devices and applications that provide increased affinity between their core products and the user.

- Seek to create incentives (or even imperatives) that encourage or require employees to apply such analytics to measure performance or honesty, or to track employees in hazardous environments for health and safety reasons.

Benefit Rating: High

Market Penetration: Less than 1% of target audience

Maturity: Emerging

Sample Vendors: Fitbit; Jawbone; Nike

Recommended Reading: "Analytics Gets Personal with the Quantified Self"

Collaborative Decision Making

Analysis By: Rita L. Sallam

Definition: Collaborative decision making (CDM) platforms combine business intelligence (BI) and other sources of information used for decision making with social and collaboration capabilities and decision support tools and methodologies, to help knowledge workers make and capture higher-quality decisions. CDM brings the right decision makers and information together with decision tools and templates to examine an issue, brainstorm and evaluate options, agree on a course of action and then capture the process to audit and mine for best practices.

Position and Adoption Speed Justification: CDM is early on the Hype Cycle because, while there are a limited number commercial offerings that comprehensively deliver on a broader decision optimization vision, a number of products have emerged over the past three years (including: Decision Lens, SAP StreamWork, IBM Lotus Connections, Hexigo, D-Sight, Lyza [from Lyzasoft], Tibco tibbr, ThoughtWeb, salesforce.com, Chatter and Calinda Software) that integrate multiple pieces and could be (and, indeed, are being) enhanced to enable a broader CDM vision. Moreover, most organizations are evolving culturally to support more transparent, fact-based decision-making and social business concepts, which will be a key requirement for widespread adoption.

We have seen these vendors enhance their collaboration and social capabilities, add decision tools and, in the case of SAP StreamWork, provide integration with SAP's enterprise applications and the SAP BI platform. As a stepping stone to CDM, a number of BI and analytics vendors, such as QlikTech, Tibco Spotfire, Yellowfin, and Panorama Software have added and continue to enhance the social and collaboration capabilities of their BI platforms. This allows them to facilitate collaboration among decision makers around specific BI content or collaborative BI (such as the ability to collaborate on the development or the result in a dashboard or a specific analysis), with more BI vendors that are planning to release similar capabilities in the future. While many vendors already provide the capability for users to tag comments to specific reports, analysis or dashboards, collaborative BI may also include social capabilities to find the right people (based on social profiles), to include in a discussion thread to comment, rate and blog on a specific result and then dynamically recommend additional analysis based on past user behavior and similar content. In collaborative BI, the focus is really on the BI artifact, such as a performance measure or query result. In collaborative decision making, the focus is on the decision itself, such as "do we change

our pricing model?" or "should we enter a new market?" or "are we in a position to hire more employees?" In addition to BI vendors, some corporate performance management vendors such as Cogniti, Portfolio Decisions, Decisyon and Actuate are enhancing the quality and transparency of planning and performance management processes by adding collaboration capabilities, value-based planning optimization tools and closed-loop monitoring of key performance indicators tied to a plan of their platforms. We characterize these capabilities as collaborative performance management (PM). In collaborative PM, the focus is on decisions specific to the planning and performance management process. Collaborative BI and PM capabilities are purpose built and provide pieces of what a broader CDM platform would enable. They could therefore be (and often are) the first step in a road map that provides broad and deep support for CDM initiatives within an organization.

A complete environment for CDM today is likely to require some custom integration and development. However, technology is not the primary barrier to CDM. Most large IT organizations could customize a social software environment with some basic templates for decision making that visually depict a decision, options, pros and cons and tags to relevant information. The real challenge to CDM adoption is cultural. Gartner believes that CDM will continue to ascend the Hype Cycle during the next 18 months because much of the technological foundation to support this use case is in place. Additionally, high-profile decision-making failures in the public and private sectors have acted as a catalyst for improved decision-making quality and transparency.

User Advice:

- Find a senior business executive willing to sponsor cultural change in support of fact-based, transparent decision making. This champion should excel at collaboration and the use of BI and analytics in decision making. The BI and analytics competency center is a logical place from which to spearhead a CDM initiative.
- Demonstrate the value of CDM through pilot projects, decision audits and simulations. Linking decisions to performance metrics, training decision makers in decision-making best practices, and using CDM in trade-off and optimization decisions will further demonstrate the value of CDM and build cultural acceptance of decision optimization as a core competency.
- Focus initially on definable departmental, line-of-business or process-specific decision processes — such as vendor selection, portfolio optimization, strategy management or forecasting — where the benefit of higher-quality and transparent decision making is easy to measure.

Provide incentives for decision collaboration and transparency to help reduce resistance.

Business Impact: CDM platforms are helping to solve the perennial challenge that — despite significant investments in BI and analytics made in the name of improving decisions — the vast majority of organizations make thousands of increasingly collaborative decisions, often with poor outcomes, without insight into how those decisions are made or their effectiveness. CDM platforms can dramatically improve the quality and auditability of decision making by alerting decision makers to events and changing patterns that indicate an early need to act.

These benefits include:

- Visually depicting what decision needs to be made, what the options are, what the weighted criteria are and what information is relevant to the decision.
- Identifying and bringing together the right people, information and analysis tools.
- Alerting decision makers to events and changing patterns that indicate the need to make a decision.
- Allowing participants to discuss an issue, assess and capture assumptions, brainstorm and evaluate options and agree on a course of action, thus enabling a new style of consensus-driven leadership.
- Capturing details of the collaboration and the information and assumptions used to make decisions.
- Providing decision tools, engines and methodologies to optimize decisions.
- Reducing the risk posed by personal bias, group thinking, failure to consider contrary views and blindness to the secondary effects of a decision.
- Improving the transparency and audit trail of decisions by capturing the details of the decision-making process and recording the "who, what, when, where, why and how" of a decision, including all inputs and assumptions.
- Linking the decision-making process to the actual outcome of the decision itself, so that it can be measured and mined for decision-making best practices and patterns that could provide leading indicators of changes in the business environment as well as templates for future decision making.

Benefit Rating: Transformational

Market Penetration: Less than 1% of target audience

Maturity: Emerging

Sample Vendors: Calinda Software; Decision Lens; Hexigo; IBM; Lyzasoft; Panorama Software; SAP; ThoughtWeb

Recommended Reading: "Who's Who in Collaborative Decision Making"

"Tutorial for Collaborative Decision Making"

"The Rise of Collaborative Decision Making"

"Overcoming the Gap Between Business Intelligence and Decision Support"

Graph Databases

Analysis By: Mark A. Beyer

Definition: Graph databases store information in a structure that records the direct relationship between any two adjacent elements. Nodes have properties and connect to other nodes at "edges." There are key value pair databases, but also network graphs and triple-stores, which are specialized forms of Graph databases that store chaining or multistep relationship sequences.

Position and Adoption Speed Justification: Graph databases are considered to have reached the Plateau of Productivity when they have penetrated at least 15% of the analytics market (which is currently many years away). Graph databases were considered a component of NoSQL solutions in 2012, but now enter the Hype Cycle as emerging rather than embryonic because Graph analysis is not a new science and has gained significant maturity. Graph analysis has been traditionally completed in advanced statistical analysis packages and can also be accomplished with open-source "R" statistical programming language.

The development of best practices for how to establish Graph nodes, directed edges and then properties that define what is in a node are already being used in statistical analysis and it is possible to adapt these lessons learned to Graph databases for feature/functional design. The preferred language and skill set for Graph is SPRQL, which requires greater familiarity in the market to gain adoption, but these are not the only skill sets or programming paradigms and all will need more experience in the market before adoption increases.

There are currently open-source and commercially-licensed Graph databases and some mainstream vendors also include Graph analytics and Graph node management in their DBMS products (for example IBM and Oracle). Many Graph databases are triple-stores using resource description framework (RDF) data model approaches, but they do not have to be triple-stores. This debate over the best approach to Graph databases in the market will resolve over the next five to seven years and lengthen the time for the solution reaching the Plateau of Productivity.

There are also layered functionality and specific features in differing Graph databases that lend advanced data science functionality to their use (for example, the R2DF framework, which utilizes RDF weights to the path ranking). Some front-end Graph analytics leverage configurable storage to create a Graph database (for example, OQGraph, which can use MySQL or MariaDB for storage), while others run as embedded servers (like OpenLink's Virtuoso).

User Advice: Organizations should determine the extent to which manually developed graph analysis is being performed by advanced statisticians to quantify the number of hours available for reduction by introducing Graph functionality (in other words, how big is the opportunity). The analysis should include the frequency at which predictive models are determined to be inaccurate and require correction.

Business Impact: Currently, business analysts work primarily with "list-based" queries. Importantly, these same analysts do not perceive their analysis as list-based because their tools combine many list results together to give the appearance of a "cube." Even cubes in most DBMSs are the result of this same premise. In a list-based analysis, each list must be constructed first and then combined. This means analysts must know what should be "in" the lists before they get them and have some idea of how the lists go together.

Graph databases do not rely on list queries and instead record every connection of every data point (called "nodes" in Graph databases) that is available at the time of data capture. In this way, the selection of a single data point will deliver all of the relevant lists and greatly reduce human bias in analytics. A significant benefit of Graph is that it is deployed with the intention of any node being a dimension or fact, which reduces the overhead for maintaining analytic dimensions.

Graph is being used for particularly successful analytics in social networks. Social networks are groups of people interacting in a number of ways. As a result, security analysts, fraud investigators and others identifying malfeasance use it. In a case of insurance fraud detection, a major insurer used Graph to identify not only the direct agents of the fraud, but also to identify gaps in the criminal network that helped them pursue the organizers that were previously difficult to identify.

Another good example is engineering analytics, where combining issues such as harmonics, structural integrity, tensile strength and many other aspects of physical devices requires interaction of all of the aspects in the final design. Graph analysis has an enormously high potential for privacy and civil liberties violations, so the producers suggest that legal parameters need to be developed immediately ahead of technology adoption.

Graph analysis is possibly the single most effective competitive differentiator for organizations pursuing data-driven operations and decisions after the design of data capture. This is a bold claim based on how Graph contributes to breaking the "list" analysis model, enabling the removal of human bias. Graph is also used in developing predictive analytics, thereby demonstrating its validity and alternative scenario development.

Graph analysis is the data science practice of determining how different concepts relate to each other in a chain with various degrees of "weight" or tendency at each edge. Graph databases are an approach that supports graph analysis, but they are not graph analytics.

Most analytics solutions can emulate a graph relationship, but human intervention is usually required to construct the data analysis that confirms or denies the graph and its overall strength of relationships. This means that the introduction of graph databases will reduce the level of human error from early bias and the amount of time taken to discover the relationships in the first place. Faster, more accurate, data-driven analytics is the result.

Benefit Rating: Transformational

Market Penetration: Less than 1% of target audience

Maturity: Emerging

Sample Vendors: IBM; Neo Technology; Objectivity; OpenLink; Oracle

Recommended Reading: "Social Infrastructure"

"Cool Vendors in Content and Social Analytics, 2013"

Prescriptive Analytics

Analysis By: Lisa Kart

Definition: Prescriptive analytics are a set of analytical capabilities that specify a preferred course of action by calculating expected future outcomes of alternative decision options. The most common examples are optimization methods, such as linear programming, and decision analysis methods, such as influence diagrams.

Position and Adoption Speed Justification: Although the concepts of optimization and decision analytics have existed for decades, they are now re-emerging along with greater awareness about advanced analytics and hype around big data. To date, many applications of prescriptive analytics have been in logistics, supply chain and manufacturing, but we are now beginning to see broader business applications, such as customer best next offer, pricing optimization and call center agent assignment. IBM's Watson "advisors," such as Oncology Treatment Advisor, are also an example of prescriptive analytics — they generate hypotheses, evaluate evidence and return the "best" responses with corresponding confidence levels.

In the past, the focus of this technology was on optimization algorithms to solving for the "best answer," which met specified business objectives and constraints. The increased adoption of prescriptive analytics includes more focus on the decision-making process, where the insights come less from finding the best answer and more from the ability to do "what if" simulation, explore alternative decisions and understand trade-offs. For example, Decision Lens assists in prioritizing strategic decisions that involve multiple objectives and trade-offs.

A recent Gartner survey among business intelligence (BI) users found that fewer than 3% of companies are using prescriptive analytics; however, that number is higher in specific industry segments — for instance, manufacturing, logistics, supply chain and airline scheduling. Because prescriptive analytics also leverages predictive methods, its adoption will be highest among companies that have built predictive capabilities. Over the next three to five years, Gartner expects prescriptive analytics to increase its hype and ultimately enter the Trough of Disillusionment as pitfalls and limitations come to fruition. It will reach a period of productivity in five to 10 years.

User Advice: The best option would be hiring individuals with optimization experience (operations research and management science). Individuals with predictive analytics experience (often with statistics — or, increasingly, analytics — degrees) are a good alternative. Otherwise, working with an experienced provider that can help you avoid pitfalls, demonstrate some initial success and learn about the process is also a common practice. However, organizations should ensure the skill transfer is effective. With prescriptive analytics, it becomes even more important to keep track of past actions and perform A/B testing to measure the impact of alternative actions. Consider prescriptive analytics and optimization tools that go beyond solving for "the answer" and provide business insight by enabling scenario comparisons and what-if analysis. An increasing number of optimization applications in the market help business decision makers do just that.

Business Impact: Prescriptive analytics can apply to strategic, tactical and operational decisions to reduce risk, maximize profits, minimize costs, or more efficiently allocate resources. Significant business benefits are common, obtained in two ways: (1) by small improvements over a large

number of operational decisions (for example, reducing inventory carrying costs in supply chain optimization or revenue yield management in air transport); or (2) by large improvements in less frequent, but complex, strategic and tactical decisions (such as acquiring a company, determining where to drill for oil, or optimizing pricing tiers). An important part of the approach is geared to make trade-offs among multiple objectives and constraints. A critical success factor is having senior business leaders closely involved such that the decision of trade-offs are aligned with achieving organization objectives.

Benefit Rating: High

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Sample Vendors: Ayata; Decision Lens; FICO; Frontline Systems; Gurobi; IBM; River Logic; SAS

Recommended Reading: "Advanced Analytics: Predictive, Collaborative and Pervasive"

"Marketing Essentials: How to Develop an Analytics Offering"

"Apply Three Disciplines to Make Business Operations More Intelligent"

Data Science

Analysis By: Douglas Laney; Roxane Edjlali

Definition: Data science is the business capability and associated discipline to model, create and apply sophisticated data analysis against disparate, complex, voluminous and/or high-velocity information assets as a means to improve decision making, operational performance, business innovation or marketplace insights.

Position and Adoption Speed Justification: Data science is a discipline that spans data preparation, business modeling and analytic modeling. Hard skills include statistics, data visualization, data mining, machine learning and database and computer programming. Soft skills that organizations frequently desire in their data scientists include communication, collaboration, leadership and a passion for data (see "Emerging Role of the Data Scientist and the Art of Data Science"). This fast-emerging capability is often also associated with big data, which Gartner defines as "information assets with volumes, velocities and/or variety requiring innovative forms of information processing for enhanced insight discovery, decision making and process automation." Unlike other business capabilities, such as CRM for example, data science does not describe the vector through which the capability delivers strategic benefit. Data science is still an emerging discipline where practices and ROI benefits are not yet established.

The increasing availability of big data, combined with the arrival of new analytics specialists called data scientists, indicates a confluence of resources and skills that can help organizations achieve transformative advances in business performance and innovation. Today, data gathering comes in many forms: from our transactional systems and social collaborative systems, but also in video and audio and from outside the enterprise in the form of complex electronic maps, syndicated data and

vast government datasets. Beyond the usual operational data, organizations are moving toward a world of instrumentation — in which sensors collect continuous detailed data on every manner of device and process. Even low-level computer system operating data logs are finding new uses (for example, in the forensic examination of record update time stamps, user behavior modeling, or preventative maintenance).

Modern analytics tools, high-level script programming languages, powerful algorithms, simple visualization tools, techniques such as video analytics and cloud sharing of datasets — when all combined — have demonstrated the potential to transform almost any organization in any sector or geography. Information-centric companies such as Google, Amazon and Facebook base far more of their decisions on complex ad hoc analysis of data. Data scientists need to be aware that in the realm of decision making the data being used is continually morphing and evolving, so that decisions made possible today may be more richly informed than those made previously (see "Toolkit: Role Description: Data Scientist").

Like many similar areas, data science is not entirely new and it has historic precursors in specialized capabilities, such as yield and revenue management, actuarial science, algorithmic trading, and informatics in various biosciences. In many ways, it extends the scope of existing business analytics with new and innovative approaches for optimizing business performance. However, the range of data types, the scale and detail of data becoming available, and the breadth of business use mark out a completely new level of capability; also, many of the tools, techniques and skills used to support data science are new. The best people are needed to make it work and they must operate within a culture that thinks differently about the way decisions are made. The traditional combination of reactive, requirements-based data warehousing and business intelligence (BI) is quickly giving way to a more proactive, opportunistic and experimental mode of execution that combines big data and data science. The term "data science" alone hints at the inclination to follow the scientific method: one of hypothesis, problem modeling, data gathering, data analysis, conclusion and retesting. However, since the term speaks more to a spectrum of analytic techniques than an overall purpose-built capability, we caution management to embrace it with business objectives in mind (see "No Data Scientist is an Island in the Ocean of Big Data").

The data scientist role is critical for organizations looking to extract insight from information assets for big data initiatives, and requires a broad combination of skills that may be fulfilled better as a team, for example:

- Collaboration and teamwork is required for working with business stakeholders to understand business issues.
- Analytical and decision modeling skills are required for discovering relationships within data and detecting patterns.
- Data management skills are required to build the relevant dataset used for the analysis.

This role is new and the skills are scarce: leading to a dearth of talent of several hundred thousand qualified professionals through to the end of the decade. The shortage is so pronounced, and demand so high, that more than 50% of organizations are now attempting to build these skills from within rather than pay extreme premiums to source them externally. While universities are

scrambling to come to the rescue, it will be a few years before they churn out data scientists with any abundance.

The data management side of data science is also giving rise to a role that is becoming more prevalent, that of the chief data officer (CDO, see "CEO Advisory: Chief Data Officers Are Foresight, Not Fad"). As information becomes an acknowledged asset, rather than just talked about as one, CDOs will emerge as the ultimate stewards of these assets. CIOs regularly contend that they are too consumed with technology-related or functionality-enabling issues to give sufficient attention to the need for improved curation, management and governance of information. The role of the CDO is to maximize value and use of data across the enterprise and to manage the associated risk. They will often focus on the places and ways in which certain information assets will have more impact on the organization. And, just as other key corporate resources have independent executive oversight and organizations (such as material assets, financial assets, human capital), information assets are also beginning to do so. As such, CIOs, CDOs and COOs (or line-of-business leaders) are starting to form a new and exciting management triumvirate.

User Advice: Catalog and consider the range of data sources available within the organization and the greater ecosystem of information assets available. Hypothesize and experiment, looking to other industries for astounding ideas to adopt and adapt. Create sandboxes for data scientists to "play" in, and don't conflate your data warehouse or BI competency center with the data science function. Then, confirm the relative economic value of findings and the organization's ability to leverage results (technically and culturally). Where could they have the most impact and is the organization ready to enact them? Recognize that data scientists are different from statisticians or BI analysts in terms of both skill set and goals. But also recognize that they are in short supply, so incubating skills internally or paying handsomely for top talent are the only options. Data science teaming arrangements that have the requisite skills in aggregate can work, but are not the same as individuals with end-to-end abilities. And, as leadership is noted as one of the key inhibitors to benefiting from big data, strive to develop deeper data literacy and acceptance throughout your management ranks about the transformative potential of data science.

Business Impact: Businesses that are open to leveraging new data sources and analytic techniques can achieve considerable competitive leaps in operational or strategic performance over those of traditional query and reporting environments. Advances in data science have yielded significant innovations in sales and marketing, operational and financial performance, compliance and risk management and new product and service innovation, and have even spawned capabilities for directly or indirectly productizing data itself. While every organization should not expect to generate quantum advances, incremental ones — with increments larger than before — are to be expected. Shifting investments and goals from hindsight-oriented descriptive analytics to more insight-oriented diagnostic analytics and foresight-oriented predictive and prescriptive analytics will hasten success.

- Risk: Moderate — Failure to evolve from basic BI is evident in many current business results.
- Technology intensity: High — IT core concept origins, unusual additional IT investment.
- Strategic policy change: Moderate — Expansion of analytics and institution of fact-based execution.

- Organization change: Moderate — Some new specialists and perhaps a new support function needed.
- Culture change: Substantial — Belief that data (as much as experience) should drive decisions can be very hard to instill.
- Process change: Low — Some decision processes will change, but core production and administration processes will not.
- Competitive value: New industries, radical product and service innovation, cost reduction and quality improvement, employee and compliance risk reduction.

Industry disruption: Three broad categories of industries (as shown below), split on a combination of related characteristics — the physical nature of the work product (for example, mining versus banking), the associated capital asset intensity of the business, the service and knowledge value-add in the business models, and the relative information intensity of the industry.

- Weightless (for example, insurance): High — New opportunities to transform enterprise performance.
- Mixed (for example, retail and consumer product goods): Moderate — Long-term competitive rankings changed.
- Heavy (for example, construction): Low — Visible benefits to a minority of players.

Benefit Rating: High

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Recommended Reading: "Emerging Role of the Data Scientist and the Art of Data Science"

"How Data Science Will Drive Digital Marketing Toward Influence Engineering"

"Big Data Strategy Components: IT Essentials"

"Predicts 2013: Information Innovation"

"Toolkit: Role Description: Data Scientist"

Search-Based Data Discovery Tools

Analysis By: Rita L. Sallam; Whit Andrews

Definition: Search-based data discovery tools enable users to develop and refine views and analyses of multistructured data using search terms. Like visualization-based data discovery tools, they have:

1. A proprietary data structure to store, model and correlate data from across structured and unstructured data, which minimizes reliance on predefined metadata.
2. A performance layer using RAM or indexing to lessen the need for aggregates, summaries and precalculations.
3. An intuitive interface enabling users to explore data with little training.

Position and Adoption Speed Justification: Search-based data discovery tools have been around for some years, but they have not been taken up as quickly as their visually driven cousins (like Tableau, Tibco Spotfire and QlikTech, for example) that have been less IT-dependent in their deployment.

Although the use of search to find pre-existing business intelligence (BI) platform artifacts (such as reports) is becoming more common, the broader use of search as a means to explore varied data types in a more free-form manner remains low. While the inclusion of keyword search of BI platform objects and object content (reports, queries, dashboards, metadata including keywords, date, time and author for example) is a valuable addition to standard BI interactions, this does not constitute search-based data discovery. Instead, it covers a broader set of use cases, specifically with respect to relating unstructured and semistructured data, such as transactions, movement, email correspondences, social media chatter and relationships to identify illegal activity or to get a more comprehensive profile of a customer.

Adoption is likely to accelerate given the rising interest in big data analytics which, by its nature, is diverse and a good fit with a search-based data discovery approach. This is in addition to greater focus on this area by megavendors, evident in Oracle's 2011 acquisition of Endeca (one of the leading lights in search-based data discovery) and IBM's acquisition of Vivisimo in early 2012.

Many search-based data discovery tools, such as Attivio, also embed advanced content analytics — such as text, social and sentiment analytics — and have both search and SQL APIs for broad user accessibility and use with existing BI and analytics tools.

User Advice: Organizations wanting to give users the chance to go beyond the analysis and reporting of structured data alone should examine the potential use of search-based data discovery tools (see "Integrating BI and Content Analytics Gives Better Results Than Using Them Separately"). In addition, those looking to make analytics more pervasive might consider using these technologies as they make it easier for those workers not accustomed to traditional, structured BI tools to find the information they need to make decisions. Organizations should, however, consider how they will fit search-based data discovery into their business analytics solution architecture and, more widely, how it relates to their enterprise search tools, by working with the IT staff (who are often remote from BI and analytics) that "own" search within the organization. It should be noted that the evaluation and adoption of search-based data discovery is something that IT must drive as, unlike visualization-based data discovery tools, these products tend not to be sold directly to individual lines of business, so this technology is unlikely to self-propagate without IT leadership.

Business Impact: The business impact of search-based data discovery technology is potentially significant, as it can extend the frame of reference beyond that commonly associated with structured reporting and help drive adoption in areas resistant to normal models of interaction with

data. Critically, search-based data discovery can unify fact and context, enabling users to explore the "what" and the "why" in one step. One example would be using search to combine a classic structured query with qualitative, often external, information (for example, "show my 50 best-performing products by revenue and region with associated online reviews and ratings"). In addition, it can assist users who would normally be too intimidated to use an ad hoc query tool, but who are quite comfortable with using a search engine to find the information they need in a structured database.

Search-based data discovery fits into the "Describe" category in Gartner's Information Capabilities Framework (see "The Information Capabilities Framework: An Aligned Vision for Information Infrastructure") and we expect that, despite the relatively low level of adoption now, these types of capabilities will be one of the most critical parts of the information infrastructure of the future.

Benefit Rating: High

Market Penetration: 1% to 5% of target audience

Maturity: Adolescent

Sample Vendors: Attivio; Coveo; EasyAsk; Exalead; HP-Autonomy; IBM; Information Builders; MarkLogic; Oracle; SAP BusinessObjects

Recommended Reading: "The Rise of Data Discovery Tools"

"IBM Will Boost Analytics Ability by Buying Search Provider Vivisimo"

"Endeca Buy Extends Oracle's Ability to Support and Discover Diverse Data"

At the Peak

Information Capabilities Framework

Analysis By: Ted Friedman

Definition: The Gartner Information Capabilities Framework describes the set of information management technology capabilities needed to define, organize, integrate, share and govern an organization's information assets in an application-independent manner to support its enterprise information management goals.

Position and Adoption Speed Justification: Organizations increasingly attempt to move from architectures of tightly coupled systems and self-contained applications to modular software components, reusable services and multipurpose content. These transitions expose information management infrastructure vulnerabilities, such as poor data quality, lack of metadata transparency, inconsistent business intelligence and analytics, conflicting master data and the lack of an integrated view across the content continuum. Organizations have technologies and processes to address such challenges, but they are scattered throughout the organization.

For most enterprises, the current approaches to information management technology are heterogeneous and complex, often with information silos affecting data sources, databases and application environments, as well as legacy data. At the technology core of new approaches to information management is an information environment, a series of co-dependent services, repositories, tools and metadata management that enable the describing, organizing, integrating, sharing and governing of all types of information in an application-neutral way, giving users the information and tools they need for their specific use case.

Innovators are aware that the optimal path to adding capacity and capabilities is no longer through the simple addition of storage, applications and databases, without consideration of how the information will move throughout the supporting infrastructure and a sense of interlocking and interactive management services.

There is now a focus on more transparency and optimization (via rich metadata capabilities), as well as standardization and reusability of functions commonly required across information-intensive use cases is of key importance. However, through 2015, 85% of enterprises will fail to adapt their infrastructures for information management to align with these ideals. Emerging architectures such as the logical data warehouse require application of the principles supported by the ICF model — and early-stage adoption of these principles is already happening.

User Advice: Organizations must rethink their approaches to delivering information management infrastructure, with a focus on capabilities that are required across multiple use cases and independent of specific applications and physical representations of data. By viewing information as a peer strategic asset, alongside applications and business processes, they can develop stronger competencies in the governance of information assets and greater value from leveraging them, while also increasing consistency, shareability and reuse.

The ICF is a vision for how these goals can be achieved. Organizations should begin to work toward this by identifying opportunities to align and standardize various information management capabilities in support of closely related initiatives, while also filling in capability gaps that may already exist in their environment.

The ICF concept does not dictate specific architectural approaches, implementation tactics, or specific tools and technologies. Rather, organizations should use it as a guiding description of the set of capabilities that, when properly aligned and integrated, can enable the fulfillment of EIM principles, which includes:

- Management of information in an application-independent manner.
- Provision of type- and source-neutral views of and interaction with information assets.
- Support for a range of use cases and consistency for these capabilities across them.
- Enablement of consistent reuse and sharing and governance of information for exponential increase in value.

The information capabilities framework addresses critical components of information management, and organizations can and should adopt these principles to promote a better understanding of the

meaning and value of information assets, to expose and share them in a variety of formats and contexts and with the appropriate conformance to information governance policies.

In effect, organizations should adopt the ICF concepts as their vision for how to fulfill the information management infrastructure requirements of the enterprise in a strategic manner.

Business Impact: Organizations will move toward the vision articulated by the ICF at different speeds and in different ways. However, the evolution toward a cohesive information infrastructure is inevitable. Enterprises are beginning to recognize that information management technologies should be approached as a coherent set of capabilities that operate on the enterprise's information assets. Gartner believes that through 2015 organizations integrating high-value, diverse, new information types and sources into a coherent information management infrastructure will outperform their industry peers financially by more than 20%.

Furthermore, the gap between leading organizations in information management practices and others will expand rapidly. Those failing to adopt the principles articulated via ICF concepts will continue to fall behind. This gap will increase, ensuring the eventual dominance of top-performing organizations.

Those that apply this approach can capture a range of specific benefits, such as:

- Enabling business growth by improving the timeliness and quality of decision making through access to a more comprehensive set of information sources.
- Improving the agility of enterprise processes for introducing new context-aware products and services.
- Improving the ability to predict new opportunities or challenges through pattern seeking, matching and discovery.
- Reducing/managing risk by improving enterprise compliance with regulations and policies through improved information quality and governance.
- Reducing the cost of storing, locating and integrating information across the information continuum.

Benefit Rating: Transformational

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Recommended Reading: "Information Management in the 21st Century"

"The Information Capabilities Framework: An Aligned Vision for Information Infrastructure"

"How to Use (And Not Use) Gartner's Information Capabilities Framework"

"Predicts 2013: Big Data and Information Infrastructure"

Information Semantic Services

Analysis By: Mark A. Beyer; Frank Buytendijk

Definition: Information semantic styles are the agreements for how to govern the interdependence between application flows and repositories. Information semantic services convert that agreement from embedded code to callable services, which include taxonomic and ontological recognition of how a business process uses the data.

Position and Adoption Speed Justification: Most applications begin with a dedicated style of semantic in which a single application joins with a single repository to manage and enforce information governance. Changing to more advanced styles will necessarily take time to complete. Over time, categories of information semantic services will be deployed to support six different styles: dedicated, registry, consolidation, external, auditing and orchestration (see "Information Management in the 21st Century Is About All Kinds of Semantics").

Importantly, while the other styles are more advanced, there will remain significant and appropriate demand for dedicated semantics. Information semantics are part of the larger information capabilities framework (see "The Information Capabilities Framework: An Aligned Vision for Information Infrastructure").

Early adoption of registry style semantics is found in analytics and data service buses. A more advanced style is orchestration, which will utilize the other five styles as "children" or components of a reconfigurable composite service. That said, master data management (MDM) is possibly one of the early strategies that utilizes orchestration to blend the other five styles together. Gartner clients have begun to report the use of such component services with very good initial results and this has accelerated interest and brought information services to the Peak of Inflated Expectations.

However, some analytics solutions and MDM programs have begun to encounter resistance to wider application precisely because they are the most demanding application of using all six semantic styles if they reach full maturity. Architects and developers simply cannot accomplish the full services approach easily — avoidance is becoming a tolerated practice. This will create significantly brittle systems and limit flexibility in the information architecture.

Adoption of semantic services is driven by the increased acceptance of standards such as Resource Description Framework (RDF) and Web Ontology Language (OWL) and the adoption of semantic technologies in areas such as social media, search, business process management, analytics, security, content management and information management. In addition, the introduction of semantic interpreters will encourage extensibility to include big data.

Significant barriers to these more complex and flexible styles include legacy applications, issues with information abstraction (such as ontology and taxonomy resolution) and a reluctance by organizations to adopt formal business process modeling (to demonstrate reusable application flows). Importantly, we emphasize more business benefits this year as a fresh perspective on the technology, which requires a "demystification" of this largely metadata function.

User Advice: Organizations should not try to implement sweeping replacements of legacy systems with loosely coupled, independent services in their information management architectures. Rather,

they should create a targeted approach to experiment with each of the six semantic styles and combinations of them.

For now, continue to build the data service layer with an orientation toward the file management, structured and unstructured repositories and message queues. The proper abstraction of information assets and the appropriate management of metadata will enable a more flexible future architecture. This particular style of interfacing between information assets collected by a wide variety of systems will be particularly useful in cloud computing for analytics and data integration.

Pursue a more formal business process design and documentation standard in the organization to promote the identification of shared application processing flows. When business processes cross, their information and information processing flows also cross one another. The advice here refers to reviewing and modeling business process flows and not enterprise data objects.

Evaluate application development platforms, data management and integration business applications for their ability to share metadata, call external services, the commonality of the developer's interface and their capabilities for specifying business logic through explicit models instead of code. Interoperability of development tools should be a highly-rated selection criteria.

Business Impact: The ability to model processes by the business and have tool-based change detection processes in place to inform IT will decrease the time to delivery for new information processing demands. Similarly, by placing process modeling at the center, the true owner of information assets (the process, not people, applications or databases) begins to push requirements simultaneously to the information design and the application design. The result is that the influence of individual managers in the business will be more easily identified as solid business process or personal agendas, which may at times be contrary to the business.

Shared application flows also means shared ontology is equally important, including identifying when assumed ontological sharing is incorrect. This will force the business to identify when its processes are attempting to share data, which actually hinders the various processes from collecting new data points when needed.

This architectural approach and its incumbent design demands help businesses identify gaps in their knowledge regarding business processes and inappropriate linking, so that the design of information assets becomes a business process clarification exercise.

While the overall benefits of utilizing these types of service is very high and probably transformational, as a practice, there is a significantly high probability that this will become embedded in other technologies.

Benefit Rating: Moderate

Market Penetration: Less than 1% of target audience

Maturity: Emerging

Sample Vendors: Ontology Works; Software AG

Recommended Reading: "Information Management in the 21st Century Is About All Kinds of Semantics"

"The Nexus of Forces Is Driving the Adoption of Semantic Technologies, but What Does That Mean?"

"Stop the Madness of Ad Hoc Information Infrastructure"

"Toolkit: Improve Architectural Decisions with Business-Driven Information Infrastructure"

Intelligent Business Operations

Analysis By: W. Roy Schulte; Teresa Jones

Definition: The term "intelligent business operations" refers to a style of work in which operational intelligence technologies are integrated into the transactional systems and physical operations that run a business. Operational intelligence technologies include monitoring and alerting tools that enhance situation awareness; decision management tools, such as rule engines and predictive and prescriptive analytics, that enable better decisions; and workflow and process orchestration engines that coordinate adaptable and collaborative processes.

Position and Adoption Speed Justification: The concept of intelligent business operations is nearing the Peak of Inflated Expectations. It appears likely to be adopted in a majority of development projects for large, new operational systems within five years — which would put it on the Plateau of Productivity by 2018.

Industry observers and vendors use many labels for this concept, including IBM's "Smarter Process" and Tibco Software's "two-second advantage." Companies are able to make their operations more intelligent because of the vast increase in available real-time data and the decreasing cost of computers, memory, mobile devices, sensors and communication networks. Vendors have steadily improved the enabling technology in areas including rule management, adaptive case management, predictive analytics, optimization, business dashboards, business activity monitoring, complex-event processing, data discovery and other technologies. New kinds of software that are intended to facilitate the implementation of intelligent business operations, notably including intelligent business process management suites and operational intelligence platforms, have recently come to market (see "Magic Quadrant for Intelligent Business Process Management Suites"and "Commercial Operational Intelligence Platforms Are Coming to Market").

The concept of intelligent business operations is part of a broader movement toward increased use of analytics in business. It is directed specifically at the day-to-day and minute-to-minute operational decisions made in the course of the processes that run a business, in contrast to the more traditional, offline use of business intelligence and analytics to make tactical and strategic decisions. Intelligent business operations apply to information-centric business activities, such as Web-based retailing and back-office banking operations, and to operational technology (OT) activities. OT is hardware and software that detects or causes a change, through the direct monitoring and/or control of physical devices, processes and events in an enterprise. OT uses sensors to detect presence (for example, via RFID tags), location (for example, via GPS units),

temperature, vibration and other data. It uses actuators to run machines, dispense medicine and control other devices. One of the primary techniques employed in intelligent business operations is to combine information from the Internet of Things with business applications to support holistic, integrated IT/OT approaches to business.

User Advice: Analysts and architects should work with business managers and subject matter experts to develop an understanding of their processes and the decisions that will be made in the course of their operations. The observe, orient, decide and act (OODA) loop is a good framework to use to analyze operations:

- If an operation requires enhanced situation awareness, analysts and developers should apply some form of event management technology to implement real-time or near-real-time monitoring (*observe* and *orient*).
- If the operation requires faster, better or more consistent decisions, analysts and developers should use decision management technologies such as rule engines, or predictive or prescriptive analytics (*decide*).
- If the operation requires flexible and better coordinated processes, analysts and developers should employ business process management (BPM) flow management tools that support adaptive case management and collaboration. In some cases, the response to a situation should be automated (the *act* phase may be manual or automated).

Where possible, companies should use real-time intelligence to detect leading indicators that warn of situations before they materialize, so that people or systems can be proactive. Where this is not possible, real-time intelligence can still improve the outcome by reducing the lag between events and responses. As companies ramp up the intelligence of their business operations, they should continue to support and evolve their offline strategic and tactical business intelligence, performance management and advanced analytics programs.

Organizations that pursue the BPM approach to implementing systems are among the most successful adopters of intelligent business operations. The BPM discipline encourages analysts to identify where operational intelligence tools should be used, and to define the relevant goals, constraints, business processes, key performance indicators, alerts, visual displays, business rules, patterns and responses.

Business Impact: The impact of intelligent business operations is immediately apparent to business people because it changes the way they do their jobs in obvious ways. The most dramatic change is an increase in situation awareness. Individual contributors and managers can see how the company is running, and what is happening in its processes and external environment. Operational decisions are better informed, more precise and consistent, and can be made more quickly. The level of collaboration between people within a department or in separate business units or companies is improved. The net result is that companies allocate staff time, equipment and other resources more efficiently; capitalize on cross-selling opportunities more effectively; pay less for goods and services; avoid some losses from fraud; provide higher levels of customer service; and perform better in many other respects.

Intelligence in business operations is relative, not a binary choice of unintelligent versus intelligent. Companies will gradually ramp up the level of intelligence in their operations over many years. Thousands of vendors support some aspects of intelligent business operations. Those listed below are just examples of vendors that have focused on this concept and provide products that support multiple aspects of it.

Benefit Rating: Transformational

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: Access Intelligence; Aha Software; Appian; Aurea; Axway; Bosch Software Innovations; C3global; Cordys; Fujitsu; IBM; Intelligent InSites; JackBe; Kofax; Lavastorm Analytics; Microsoft; OpenText; Oracle; Oversight Systems; Pegasystems; Provenir; Rockshore; SAP (Sybase); Software AG; Splunk; Systar; Tibco Software; Vitria; WestGlobal; Whitestein

Recommended Reading: "Commercial Operational Intelligence Platforms Are Coming to Market"

"Use Intelligent Business Operations to Create Business Advantage"

"Magic Quadrant for Intelligent Business Process Management Suites"

"Apply Three Disciplines to Make Business Operations More Intelligent"

Logical Data Warehouse

Analysis By: Mark A. Beyer

Definition: The logical data warehouse (LDW) is a new data management architecture for analytics combining the strengths of traditional repository warehouses with alternative data management and access strategies (specifically, federation and distributed processing). It also includes dynamic optimization approaches and multiple use-case support.

Position and Adoption Speed Justification: The LDW will form a new best practice by the end of 2015. In early 2013, Gartner clients began reporting a significantly increased interest in the LDW and its many forms (see "Understanding the Logical Data Warehouse: the Emerging Practice") for background and previous adoption rates. The discussions are taking multiple forms, including revelations that many large and midsize organizations are pursuing the approach already, but are now using the nomenclature. There is also new interest from organizations seeking some methodological approaches to combining big data technologies, federation and centralized data warehouses.

Additionally, vendors have incorporated the naming convention into their marketing and messaging. Finally, new vendors able to offer federation/virtualization solutions seek the opportunity to become major contenders for the semantic data access layer for analytics. This brings the LDW to the Peak of Inflated Expectations. LDW discussions now include the concepts of MapReduce and Graph

analysis (completed in server clusters outside the warehouse with results being loaded to the warehouse or available via federation as a services call).

From 1H13, Gartner data warehouse inquiries now include the LDW in some form between 15% to 20% of the time. Early adopters are still advanced data warehouse and analytics practitioner organizations and implementers, but database and business intelligence vendors are at least addressing the LDW in their offerings. During the next two years, new organizations will begin to encounter the more difficult issues of managing SLAs for each delivery types possible under the LDW and many will fail to manage the still prevalent performance and availability issues of virtualization and batch distributed processes running on server clusters external to the warehouse. This will lead the LDW into the Trough of Disillusionment and at this time, it is possible that the LDW will fail to emerge (just as distributed warehouses failed during 2008 to 2009).

Gartner believes that as LDW architects become more practiced in using services-oriented approaches they will be able to introduce more dynamic metadata-driven services engines and this will advance the LDW from the Trough of Disillusionment and upward along the Slope of Enlightenment by late 2015 or early 2016.

User Advice:

- Conduct query analysis of existing analytics to determine how the current system performs (the current warehouse, mart or federated views) and what data/information is included in these queries. Use the results of this analysis to identify when users are leaving the warehouse to obtain data from other information resources.
- Determine if your current warehouse software systems are capable of managing external data access and managing external processing clusters in terms of specifying jobs, initializing jobs and monitoring or managing the job flows as they complete, or if new technology should be evaluated.
- Start small with ontology and taxonomy, for example deploy and populate a business glossary/ data dictionary, then move onto the more advanced technologies that perform wider metadata management and ontology/taxonomy rationalization.
- Organizations should identify a single area of analytics requiring a combination of three information access and management approaches:
  - Traditional repository style approaches
  - Real-time access to operational systems (to pilot virtualization)
  - An embedded use of distributed processing (such as, MapReduce of large datasets and graph analysis of networks of information or content analytics)
- Build a pilot analysis, collecting inputs from all three information access and management approaches.

Business Impact: An LDW has the potential to eliminate the constant level of compromise between comprehensive data needs, performance optimization and time to delivery cycles. By introducing

virtualization and distributed processes as peers to the repository, it is now possible to select the deployment architecture based on the driving service-level expectation instead of defaulting to existing practices. As such, traditional data warehouse vendors are supporting aspects of the LDW to varying degrees (IBM, Oracle, Teradata and SAP, for example).

Other smaller vendors also have an opportunity specifically in the area of multiple use cases for the data warehouse data.

In addition, many information assets deployed in a repository-only style warehouse are forced to follow a single ontologic/taxonomic pairing. In the LDW, a semantic layer can contain many combinations of use cases, which are effectively deployed as new variations of data taxonomies and ontologies. Many business definitions of the same information are the result. From a technical perspective, the LDW also needs a query normalization interface.

Every DBMS requires a unique set of SQL or data access interface technologies (for example, search and API access). This capability opens up many options for end-user tools and removes the current design requirement to closely coordinate the business intelligence tools with the database management system choice (see "Decision Point for Logical Data Warehouse Implementation Styles").

Early adopters have reported that the LDW does not have to follow the 80/20 analytics rule (specifically, 80% of analysis needs can be met by 20% of the data and thus designed easily into an optimized layer of the data warehouse). Instead, the LDW is using an 80/10/5 rule in which 80% of analytic needs are met by the repository, but 10% of the demand is met by virtualized solutions and 5% is met via distributed analytics (such as content analytics, MapReduce or Graph).

This leaves a remaining 5% and LDW practitioners are pleased to acknowledge that the remaining 5% will never be solved, by even this advanced infrastructure, and will remain in end users' specific control. These same customers report that 5% of needs are never met in a fixed architectural choice, and acknowledge the credibility of undefined access using any combination of the three choices or creating their own direct access and extract outside of the LDW.

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Emerging

Sample Vendors: BMMSoft; Composite Software; Denodo Technologies; IBM; Informatica; Kognitio; MarkLogic; Teradata

Recommended Reading: "Decision Point for Logical Data Warehouse Implementation Styles"

"The Logical Data Warehouse Will Be a Key Scenario for Using Data Federation"

"Understanding the Logical Data Warehouse: The Emerging Practice"

"The Future of Data Management for Analytics Is the Logical Data Warehouse"

Big Data

Analysis By: Mark A. Beyer; Sid Deshpande

Definition: Big data is high volume, velocity and variety information assets that demand cost-effective, innovative forms of information processing for enhanced insight and decision making.

Position and Adoption Speed Justification: Big data is almost at the Peak of Inflated Expectations. It will become an embedded and state-of-the-art practice by 2018, and it is more likely that big data management and analysis approaches will be incorporated into a variety of existing solutions in existing markets (see "Big Data Drives Rapid Changes in Infrastructure and $232 Billion in IT Spending Through 2016").

Notably, organizations have begun to indicate that existing analytics will be modified and enhanced by big data and not replaced (only 11% of data warehouse leaders indicated they would consider replacing the warehouse with a NoSQL or big data solution as of November 2012, down from just over 20% in 2011). Practices are diverging at this point, with confusion starting to emerge regarding exactly what constitutes big data and how it should be addressed. Some very traditional vendors that have not been considered for big data solutions should be considered, and this confusion may be their entry point into the debate about which tools to use. Other vendors will simply relabel their existing products as big data and not actually offer anything new.

Beginning late in 2014 and through the end of 2015, big data will descend into the Trough of Disillusionment as conflicting concepts of what it is and how organizations can benefit from its management and analysis multiply.

There are two significant facts that will drive it into the trough.

- Tools and techniques are being adopted ahead of learned expertise and any maturity/ optimization, which is creating confusion.
- The inability to spot big data opportunities by the business, formulate the right questions and execute on the insights.

MapReduce continues to persist as the "darling" of big data processing. Even with new additions or wider use of the Hadoop project (such as HCatalog) it remains a batch solution and so has to be combined with other information management and processing technologies. Hadoop implementations require expert-level staff or system implementers.

As anticipated in 2011, attempts to combine MapReduce with Graph have followed and inadequate attempts to address other big data assets, such as images, video, sound and even three-dimensional object modeling, will drive big data into the trough. Some big data technologies represent a great leap forward in processing management, especially relevant to narrow but deep (many records) datasets, such as those found in operational technology, sensor data, medical devices and mobile devices, among others. Big data approaches to analyzing data from these technologies represent the potential for big data solutions to overtake existing technology solutions when the demand emerges to access, read, present or analyze any data.

The larger context of big data refers to the wide variety and extreme size and count of data creation venues in the 21st century. Gartner clients have made it clear that big data must include large volumes processed in streams, as well as batch (not just MapReduce) and an extensible services framework deploying processing to the data or bringing data to the process, spanning more than one variety of asset type (for example, not just tabular, or just streams or just text). Importantly, different aspects and types of big data have been around for more than a decade — it is only recent market hype around legitimate new techniques and solutions that has created this heightened demand.

User Advice:

- Identify existing business processes that are hampered in their use of information because the volume is too large. There are many information gaps that could be filled by new information types (variety) or the velocity will create processing issues. Then identify business processes that are currently attempting to solve these issues with one-off or manual solutions.
- Review existing information assets that were previously beyond existing analytic or processing capabilities (referred to as "dark data") and determine if they have untapped value to the business, making them a first or pilot target of your big data strategy.
- Plan on utilizing scalable information management resources, whether public cloud, private cloud or resource allocation (commissioning and decommissioning of infrastructure), or some other strategy. Do not forget that this is not just a storage and access issue. Complex, multilevel, highly correlated information processing will demand elasticity in compute resources, similar to the elasticity required for storage/persistence.
- Extend the metadata management strategies already in place and recognize that more is needed to enable the documentation of big data assets, their pervasiveness of use and the fidelity or assurance of the assets by tracking how information assets relate to each other and more.

Business Impact: There are three principal aspects to big data — success will be limited unless all are addressed. The quantitative aspects of big data generally do not emerge one by one. Volume, variety and velocity most often occur together. The second aspect is that innovation must be cost-effective both in costs to deploy and maintain and in terms of time to delivery — solutions that arrive too late are useless, regardless of cost.

Finally, the focus must be on increased insight by the business into process optimization from immediate automation through the development of completely new business models. Big data permits greater analysis of all available data, detecting even the smallest details of the information corpus — a precursor to effective insight and discovery.

The primary use cases emerging include leveraging social media data and combining operational technology (machine data) with back-office and business management data and further validating existing assets (increasing their "fidelity").

Perhaps the most important business benefit of big data management and analysis techniques are that analytics and decision processing can include multiple scenarios, including highly disparate

definitions and temporality of events in the data. This means that analytics can now comprise many different scenarios. Each scenario could have different starting and ending points, and differing relationships within the data and circumstantial inputs. Finally, analysts would be able to attach probabilities to each scenario and monitor many of them simultaneously.

Benefit Rating: Transformational

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: Cloudera; EMC-Greenplum; HortonWorks; IBM; MapR; Teradata-Aster Data

Recommended Reading: "Big Data Drives Rapid Changes in Infrastructure and $232 Billion in IT Spending Through 2016"

"Big Data' Is Only the Beginning of Extreme Information Management"

"How to Choose the Right Apache Hadoop Distribution"

"CEO Advisory: 'Big Data' Equals Big Opportunity"

"The Importance of Big Data: A Definition"

Enterprise Metadata Management

Analysis By: Andrew White; Roxane Edjlali

Definition: Gartner defines enterprise metadata management (EMM) as the business discipline for managing metadata for the information assets of an organization across information initiatives in support of enterprise information management (EIM). Metadata is defined as "information that describes various facets of an information asset to improve its usability throughout its life cycle" (see "Gartner Clarifies the Definition of Metadata").

Position and Adoption Speed Justification: Metadata is pervasive throughout an organization in the form of data about information assets such as structured and content, models, designs, processes, workflows, applications and others. For most organizations, metadata is managed at the project level (for example, in support of application integration, or a data warehouse enabling BI, or master data management [MDM]). To increase the rate of ROI for information assets (specifically, improve EIM maturity), metadata needs to align across these silos and EMM is the discipline best capable for this task.

Few organizations have a formal EMM discipline in place that links and governs the metadata across disparate programs, leading to an enterprise view of the central metadata linking core business process, applications and analytics. We are seeing increased interest from Gartner customers looking to implement "enterprisewide metadata management programs" to "align and leverage all the information assets across siloed programs."

Typically individual IM programs have their own metadata — programs spanning MDM, information governance, e-discovery, information governance, risk and compliance (GRC), ERP and information life cycle management (ILM).

Newer areas driving interest in leveraging metadata include social, mobile and cloud computing and big data opportunities and threats involving disparate forms of new types of data and content. As these related disciplines become more commonplace and organizations seek to improve their EIM maturity, we expect organizations to realize that EMM is a prerequisite for such progress.

EMM is a technology-enabled discipline to help govern and sustain a single, consistent view of metadata shared by the various information programs. EMM is part of an EIM program that helps align and link the individual programs, but does not exist on its own or in isolation to the other programs. EMM does not replace current uses of metadata and in support of a program, it may reuse many metadata management tools to document and actively manage metadata (in terms of governance, compliance, security and collaborative sharing) and perform analysis (such as change impact analysis and gap analysis) using the metadata.

Repositories can also be used to publish reusable assets (such as application and data services) and browse metadata during life cycle activities (design, testing, release management and so on). However, implementing the technologies capable of managing the enterprisewide variety, volume, velocity and complexity of metadata about vital information assets, can be cost prohibitive for many/most organizations — generally far more than $1 million.

Most organizations find the greatest challenge in trying to use EMM across the organization lies in dealing with the cultural and political issues involved in the sharing of metadata and information assets. It can be a real drain on analytical and end-user resources to set up and maintain an EMM program, which might be better used to address other business opportunities and threats. Hence the need to have a justifiable ROI set of metrics on whatever EMM implementation strategy is selected (and why this discipline will move slowly on the Hype Cycle, rather than speed its way to the Plateau of Productivity, for most organizations).

It is also why the emergence of initiatives such as MDM and information governance will be needed to justify the related necessary investments in EMM. Increasingly, the need for EMM is being identified as a way to respond to short term reporting and privacy compliance mandates from government agencies.

User Advice: Only explore EMM when you have multiple, disparate IM programs (each with their own metadata and management) that are not aligned, or leveraging consistent information between them. Use EMM to help govern the metadata and information assets between these programs (see "Defining the Scope of Metadata Management for the Information Capabilities Framework," "Metadata Management is Critical to MDM's Long-Term Success" and "Understanding the Logical Data Warehouse: The Emerging Practice").

EMM is only sought when the organization is seeking to align its information management programs into a more mature EIM framework. If you want to manage metadata to support a data warehouse for BI, do this instead and don't seek EMM. The variety and volume of assets across these information systems can span different types (plans, processes, applications, data, business rules

for example), across different architectures (business, technology, information and application/ solution), at different levels of abstraction (conceptual, logical and physical), supporting cross-organizational (internal and external) roles and describing the context in which assets are used and how they relate to each other (specifically, for change impact analysis and understanding across "pace layers").

EMM won't be adopted overnight, it will form as metadata is shared between two IM initiatives, to improve their maturity and increase the positive influence on business outcomes, for example, linking MDM to an ERP migration and renewal program. EMM will grow in scope and maturity as more and more information silos are aligned (and then aligned to the data warehouse use of metadata).

As described in "CIO Critical Capabilities: Metadata Made Simple,"one of the key responsibilities of a CIO is ensuring the effective management and leveraged use of the organization's information assets via EMM. Usually, the CIO delegates the implementation and management of the discipline to others such as information managers, enterprise architects and those responsible for GRC.

This means that while someone like the CIO (or emerging chief data officer) needs to champion EMM, it requires support from those whose assets will be managed. Each organization will need to assess the current level of its readiness for EMM by considering the risks, opportunities and anticipated changes to the use of its information assets.

For many, the best place to start is to identify and publish which key information assets and metadata are currently being managed by primary stakeholders and shared with others (see "The Eight Common Sources of Metadata" and "Six Common Approaches to Metadata Federation and Consolidation").

There needs to be an EMM strategy or plan for how to improve the situation by leveraging other planned initiatives — which may involve the participation of individuals from different organizational units. The plan should include those in roles related to information management, enterprise architecture, business process management (BPM), service-oriented architecture (SOA), and GRC, along with other business stakeholders.

Business Impact: EMM enables initiatives such as MDM, BPM and SOA and makes them more sustainable because it helps to align and share common information assets between the programs. It can also help transition from the current solution architecture with its pace layers of "system of record" and "systems of difference" to incorporating new "systems of innovation" to create business competitiveness differentiation (see "Metadata Management for Pace Layering").

The discipline for a EMM transition strategy needs to account for people and process issues, but it also needs to address technology issues and choices including those related to identifying the best metadata to use (see "How to Tell Which Metadata Is Valuable"and "Toolkit: Calculate the Value of Your Metadata"), the viability of the technology housing the metadata (see " "Decision Framework for Evaluating Metadata Repositories" and " "Toolkit: Sample RFI and Vendor Rating Spreadsheet for Evaluating Metadata Repositories") and approaches to federating or consolidating metadata across technologies.

It is important to note that there are many service providers performing training and consulting in the discipline of EMM, but we have chosen to list vendors (see the Sample Vendors section) selling EMM-enabling software (specifically, metadata repositories) since they not only provide training and consulting but do so in a more coordinated and customized way with their EMM tools. Many also have large user groups with years of pragmatic experience in implementing the EMM discipline successfully.

Benefit Rating: High

Market Penetration: 1% to 5% of target audience

Maturity: Emerging

Sample Vendors: Adaptive; ASG Software Solutions; Data Advantage Group; IBM; Oracle; Software AG; SOA Software

Recommended Reading: "Gartner Clarifies the Definition of Metadata"

"How Metadata Improves Business Opportunities and Threats"

"Metadata Management for Pace Layering"

"The Eight Common Sources of Metadata"

"Six Common Approaches to Metadata Federation and Consolidation"

"Defining the Scope of Metadata Management for the Information Capabilities Framework"

"Information Management in the 21st Century Is About All Kinds of Semantics"

"Decision Framework for Evaluating Metadata Repositories"

"Toolkit: Sample RFI and Vendor Rating Spreadsheet for Evaluating Metadata Repositories"

"How to Tell Which Metadata Is Valuable"

"Toolkit: Calculate the Value of Your Metadata"

"Metadata Management is Critical to MDM's Long-Term Success"

"Magic Quadrant for Enterprise Architecture Tools"

"Gartner Assessment of Enterprise Architecture Tool Capabilities"

"Metadata Will Improve the Return on Your Video Investments"

"Decision Point for Choosing the Best Database and Metadata Format"

"Toolkit: How to Create and Maintain an ESI Inventory"

Complex-Event Processing

Analysis By: W. Roy Schulte; Zarko Sumic; Nick Heudecker

Definition: Complex-event processing (CEP), sometimes called event stream processing, is a kind of computing in which incoming data about events is distilled into more useful, higher-level and more complex event data that provides insights into what is happening. Multiple events from one or more event streams are correlated on the basis of having a common value in a key field, and patterns and trends are detected. One complex event may be the result of calculations performed on dozens, hundreds or even thousands of input (base) events.

Position and Adoption Speed Justification: CEP has progressed slightly on the Hype Cycle, putting it just past the Peak of Inflated Expectations. It attracts considerable attention because it is the technology used for many kinds of real-time analytics on big data. However, companies are adopting CEP at a relatively slow rate because its architecture is so different from conventional system designs. Although most developers are unfamiliar with this technology, it is the only way to get many kinds of insights from event streams in real time or near real time, so it will inevitably be adopted in multiple places within virtually every company. It may take up to 10 years for CEP to reach the Plateau of Productivity and be in use in the majority of applications for which it is appropriate.

The amount of real-time event data is growing rapidly. It comes from transactional application systems, sensors, market data providers, social computing (activity streams), tweets, Web-based news feeds, email systems and other sources. Companies need to tap the information in these event streams to be able to respond faster and smarter to changing conditions but conventional software architectures can't address these business requirements. Conventional applications use a save-and-process paradigm in which incoming data is stored in databases (in memory or on disk), and then queries are applied to extract the relevant subset of data needed for a particular application function (in other words, the query comes to the data). This paradigm is too slow when the volume of incoming data is high and the results of computation must be made available immediately. Architects are therefore driven to use an alternative, event-driven design in which the computation is triggered immediately as input data is received into memory. In this "process first" CEP paradigm, the application runs continuously and the data comes to the query because the query is in place before the data arrives. This is known as "data in motion."

User Advice: Companies should use CEP to enhance their situation awareness and to automate certain kinds of decisions, particularly those that must be made in real time or near real time. Situation awareness means understanding what is going on so that you can decide what to do. CEP should be used in operational activities that run continuously and need ongoing monitoring, using a sense-and-respond approach. For example, it can apply to near-real-time precision marketing (cross-sell and upsell), fraud detection, factory floor and website monitoring, customer contact center management, trading systems for capital markets and transportation operation management (for airlines, trains, shipping and trucking). In a utility context, CEP can be used to process a combination of supervisory control and data acquisition (SCADA) events and "last gasp" notifications from smart meters to determine the location and severity of a network fault, and then to trigger appropriate remedial actions.

Developers can obtain CEP functionality by custom coding it into their application or tool, or by acquiring an event-processing platform and tailoring it to their specific requirements. Prior to 2004, developers wrote custom code for CEP logic as part of their application or tool, because off-the-shelf event-processing platforms were not available. Developers still write custom CEP logic for many purposes, but a growing number of developers leverage commercial and open-source event-processing platforms to reduce the time and cost required to implement CEP-based applications (see examples of vendors listed below). Some commercial event-processing platform products have extensive analytic capabilities, such as built-in statistical functions, tools for building business dashboards, off-the-shelf adapters for packaged applications or industry message-format standards, alerting mechanisms and graphical development tools.

In many cases, CEP will be acquired in a packaged application or tool, or obtained as part of a SaaS offering. Leading-edge architects and developers have become aware of event processing, event-driven architecture and CEP, and are making build-versus-buy decisions in an increasing number of projects. Event-processing platforms are sometimes used in conjunction with intelligent business process management suites or operational intelligence platform products to provide more intelligent process monitoring, or to help make business decisions on a dynamic, context-aware basis. CEP plays a key role in event-driven business process management, alongside rule engines and structured process-orchestration capabilities.

Business Impact: CEP:

- Improves the quality of decision making by presenting information that would otherwise be overlooked.
- Enables faster response to threats and opportunities.
- Helps shield business people from data overload by eliminating irrelevant information and presenting only alerts and distilled versions of the most important information.
- Reduces the cost of manually processing the growing volume of event data.

CEP adds real-time intelligence to operational technology (OT) and business IT applications. OT is hardware and software that detects or causes a change through the direct monitoring and/or control of physical devices, processes and events in the enterprise. OT goes by various names in different industries, and is often owned and operated independently of IT systems. For example, utility companies use CEP as a part of their smart grid initiatives, to analyze electricity consumption and to monitor the health of equipment and networks. Elsewhere, CEP helps to process feeds of event data such as temperature, vibration and revolutions-per-second that, when analyzed together, may predict impending equipment failure. CEP is also used in business-activity monitoring applications that have a high rate of input data (high throughput), require fast (low latency) responses or require the detection of complex patterns (especially those that are temporal or location-based).

CEP is one of the key enablers of context-aware computing and intelligent business operations strategies. Some of the more sophisticated operational intelligence platform products use CEP to provide pattern matching and situation detection capabilities. The biggest single source of future

demand for CEP may be the emerging Internet of Things. Social computing may be the second largest source of new data and demand for CEP.

Benefit Rating: Transformational

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: Apache; EsperTech; FeedZai; Grok; HStreaming; IBM; Informatica; LG CNS; Microsoft; Oracle; Red Hat; SAP-Sybase; SAS-DataFlux; ScaleOut Software; Software AG; Splunk; SQLstream; Tibco Software; Vitria; WestGlobal; WS02

Recommended Reading: "Use Complex-Event Processing to Keep Up With Real-time Big Data"

"Best Practices for Designing Event Models for Operational Intelligence"

"Cool Vendors in Analytics"

"Apply Three Disciplines to Make Business Operations More Intelligent"

In-Memory Database Management Systems

Analysis By: Roxane Edjlali; Donald Feinberg

Definition: An in-memory DBMS (IMDBMS) is a DBMS that stores the entire database structure in memory and accesses all the data directly, without the use of input/output instructions to store and retrieve data from disks, allowing applications to run completely in-memory. This should not be confused with a caching mechanism, which stores and manages disk blocks in memory cache for speed. IMDBMSs are available in both row-store and column-store models or a combination of both.

Position and Adoption Speed Justification: IMDBMS technology has been around for many years (for example, IBM solidDB, McObjects eXtremeDB and Oracle TimesTen). However, many available now were introduced within the past two or three years: VoltDB, SAP Sybase Adaptive Server Enterprise (ASE), SAP Hana and small vendors that continue to emerge (for example, ParStream and Altibase). What is new in 2013 is that all major DBMS vendors have shipped or announced support of in-memory capabilities in their relational DBMSs (RDBMSs):

- Oracle released Exalytics leveraging TimesTen as the IMDBMS.
- Microsoft announced project Hekaton as part of SQL Server 2014 and delivered xVelocity Column-store as part of SQL Server 2012.
- IBM announced BLU Acceleration for DB2 and Informix.
- Teradata announced Teradata Intelligent Memory.

This adoption from all major vendors demonstrates the growing maturity of the technology and the demand from customers looking at leveraging IMDBMS capabilities as part of their information infrastructure. While SAP Hana is leading the charge with over 1,500 customers — announced at Sapphire in May 2013 — the addition of in-memory capabilities by all major players should further accelerate adoption of IMDBMS technology in the next two years.

Many use cases are supported by IMDBMS. For example, solidDB and TimesTen were originally developed for high-speed processing of streaming data for applications such as fraud detection, with the data then written to a standard DBMS for further processing. Others such as Altibase, SAP Sybase ASE and VoltDB focus on high-intensity transactional processing. Some IMDBMSs — such as Exasol, ParStream or Kognitio — are dedicated to in-memory analytical use cases. SAP Hana, although primarily for analytics and data warehousing, is now supporting both online transaction processing (OLTP) and analytics in the same database (with the general availability of BusinessSuite on Hana announced in May 2013).

The main business value offered by an IMDBMS is that it enables new business opportunities that would not have been possible previously. One example would be the sending of SMS discount coupons to customers based on their preferences when they are in, or close to, a shop. Another example comes from online gambling, whereby computing of the handicap could occur as a match is ongoing. To support such use cases, both the transactional data and the analytics need to be available in real time. The promise of IMDBMSs is to combine, in a single database, both the transactional and analytical use cases without having to move the data from one to the other.

However, a number of factors limiting adoption have been identified:

- The level of support for high availability and disaster recovery can vary. The perceived risk involved with memory failures and lack of reliable high-availability disaster recovery, and sufficiently fast backup and recovery techniques continue to be an issue when selecting an IMDBMS. As this functionality is added often through a combined software and hardware offering including clustering, this inhibitor will decrease in importance.
- There is a limited number of packaged applications, data integration tools and business intelligence tools support.
- There is potential for development impact on applications to leverage IMDBMS capabilities.
- IMDBMS requires rethinking data management practices. Managing all of the data in-memory may not be a viable option. Using transactional data for analytical use cases does not remove data consistency or data quality issues.

The availability of skills and established practices is limited.

User Advice:

- Continue to use IMDBMS as a DBMS for temporary storage of streaming data where real-time analysis is necessary, followed by persistence in a disk-based DBMS.
- For most organizations, IMDBMS for OLTP technology remains five years away. Monitor the technology and case studies for signs of maturing adoption and use cases.

- IMDBMSs for analytic acceleration is an effective way of achieving increased performance. However, given the diverse level of maturity of the vendors, organizations should verify references and run a thorough proof of concept to ensure the technology meets their expectations.
- The single most important advancement will come as IMDBMS matures as a column-store, combined OLTP and online analytical processing model as a basis for new, previously unavailable applications, taking advantage of real-time data availability, with IMDBMS for increased performance and reduced maintenance. Organizations should monitor technology maturity and identify potential business use cases to decide when to leverage this opportunity.
- Organizations evaluating IMDBMS technology should revisit data governance, data integration and data warehouse practices and design to take advantage of the increased performance, without losing data quality or consistency.

Business Impact: Once these IMDBMSs become mature and proven — especially for reliability and fault tolerance — and as the price of memory continues to decrease, the potential for the business is transformational:

- The speed of IMDBMS for analytics has the potential to simplify the data warehouse (DW) model by removing development, maintenance and testing of aggregates, summaries and cubes. This will lead to savings in terms of administration and increased flexibility for meeting diverse workloads.
- The high performance implies that smaller systems will do the same work as much larger servers, which will lead to floor space and power savings. While cost of acquisition of IMDBMS is higher than a disk-based system, the total cost of ownership of an IMDBMS should be less over a three- to five-year period because of personnel, floor space, power and cooling cost savings.
- Column-store IMDBMSs have the potential for a combined OLTP and DW single database model that will enable an entire set of new applications. These applications were not possible before because of the latency of data moving from OLTP systems to DW. However, this use case is still in its infancy.

Benefit Rating: Transformational

Market Penetration: 1% to 5% of target audience

Maturity: Adolescent

Sample Vendors: Exasol; IBM; Kognitio; McObject; Microsoft; Oracle; ParStream; Quartet FS; SAP; VoltDB

Recommended Reading: "Who's Who in In-Memory DBMSs"

"Cool Vendors in In-Memory Computing, 2013"

"Taxonomy, Definitions and Vendor Landscape for In-Memory Computing Technologies"

Sliding Into the Trough

Model Managers

Analysis By: Gareth Herschel

Definition: Model managers help organizations create, manage and use predictive models. They are not model development platforms, but are complementary to them. Model managers serve three purposes: they act as a collaboration environment between model producers and consumers, including scheduling and monitoring of model creation; they serve as a model "library" as well as reporting on their key attributes; and they monitor the use and performance of the models, enabling assessment of their business value and prompting updates when performance deteriorates.

Position and Adoption Speed Justification: Model managers have been around for several years but their importance is growing as the scale of organizational use of predictive analytics increases. However, most of the tools in the market are focused on only one or two of the core functions of a model manager. This makes this market high profile but still relatively immature.

User Advice: Organizations with multiple model development environments, model developers responsible for supporting multiple business units, or organizations that can reuse the same models in multiple lines of business should invest in a model management application. The model manager offered by the data mining platform vendor should be the default option. However, other vendors' solutions should be evaluated if the default choice is not a perfect fit. This is because model managers should be platform independent.

Business Impact: Model managers enable better prioritization of which models should be developed and updated, as well as ensuring that models are used appropriately across the organization.

Benefit Rating: Moderate

Market Penetration: 1% to 5% of target audience

Maturity: Adolescent

Sample Vendors: FICO; IBM-SPSS; KXEN; Pitney Bowes (Portrait Software); SAS

Mobile BI

Analysis By: Joao Tapadinhas

Definition: Mobile business intelligence (BI) is the delivery of reports, dashboards and basic analytic capabilities through tablets and — to a lesser extent — smartphones, that provide touch-based interfaces and a user experience that typically surpasses what is found on desktop BI deployments. With the rapidly increasing mobility of the workforce, mobile BI is a new — and renewed — focus area for BI vendors. BI customers are excited once again about the possibilities.

Position and Adoption Speed Justification: Mobile BI is not a new idea but has been enjoying a huge spark of interest over the last three years, with the hype building fast. Most vendors have embraced the trend. However, in some cases they have failed to live up to expectations. This is because mobile solutions that work as desktop replicas — able to deliver basic desktop-like BI content — fail to use some of the mobile devices' unique capabilities and only partially leverage touchscreen interface conventions. The focus has mostly been on basic reporting and dashboarding, with few advanced mobile use cases leveraging the GPS, camera or microphone. Also, solutions provide limited analytical functions and close to no applications with write-back capabilities to back-end platforms.

However, some solutions are mobile-optimized and can deliver an interesting new user experience with the ability to navigate around and drill into reports through touchscreen interfaces and finger gestures. Such gestures include tapping, squeezing or swiping through content. In addition, they provide location awareness using GPS or inputs through the camera or microphone — such as scanning a bar code or asking for a report through speech. This leads to a whole new set of use cases, with adoption by new roles that have traditionally not been recognized as "BI users," like store managers or parcel delivery truck drivers.

Although general interest in the topic is very high, adoption is expected to continue relatively sluggishly. Mobile BI can deliver real value to organizations by optimizing business processes with information "on-location" to support decisions, but relatively few production-scale implementations are underway. For the time being, we expect the continuation of development efforts by the vendor community trying to grab attention and position the mobile device as an attractive BI delivery mechanism. However, there will only be cautious experimentation and adoption by customers. Sending potentially sensitive data to a mobile phone or tablet, the need for BI content rework and the new skill sets required will remain challenges in many organizations, hindering deployment. At the same time, many organizations are still struggling with rather mundane BI issues and are not necessarily ready to invest in yet another idea, however "cool" it may appear.

User Advice: Organizations should try identifying use cases in which mobile users have a need for easy access to up-to-date information held in back-office applications or data warehouses to run small pilots to test mobile BI applications. The delivery of key metrics, through mobile reports and dashboards — to top and middle management or sales teams, for example — are recommended as starting use cases. This is because of their ease of deployment, potential benefits and visibility in the organization. However, the main infrastructure — consisting of the data warehouse, BI platform, data integration and, of course, data quality — must be in good shape before mobile BI becomes a priority. There is no point supplying modern devices to a mobile workforce so it can get access to BI anywhere when the submitted data is wrong.

Over time, mobile BI should evolve to be used for use cases meant to optimize business processes — at the production plant, warehouse, delivery truck or store — reaching out to new constituencies that can benefit from information at the point of decision. Information latency will have to be reduced and data quality must be at its best, to comply with more stringent requirements of operational usage. However, the positive impact of these solutions will likely compensate for the extra effort and investment.

Business Impact: Decision making is not something only for the back office or the boardroom. Decisions are being made on the road and in warehouses, supermarkets, client meetings or airport lounges, where the decision maker often needs quick access to a few key metrics shown on a mobile device. This can reduce decision bottlenecks and latency, increase business process efficiency and enable broader input into the decision at hand. Mobile BI can deliver all this through engaging BI outputs, capable of attracting new users — from top management to field service personnel.

Benefit Rating: Moderate

Market Penetration: 5% to 20% of target audience

Maturity: Emerging

Sample Vendors: Antivia; IBM; Information Builders; MeLLmo; MicroStrategy; Oracle; QlikTech; SAP (BusinessObjects); Tableau Software; Tibco Software (Spotfire); Transpara

Recommended Reading: "Best Practices to Deploy Mobile Business Intelligence"

"Innovation Insight: Mobile BI Innovation Expands Business Analytics Boundaries"

"Critical Capabilities for Mobile BI"

"Who's Who in Mobile BI"

Business Analytics SaaS/PaaS

Analysis By: Neil Chandler

Definition: Software or platform as a service (SaaS or PaaS)-based business analytics (BA) enables customers to deploy one or more components of BA without significant IT involvement or the need to deploy and maintain an on-premises solution. Offerings include analytic applications with prepackaged functionality for specific solutions; BA platforms that provide a development/ integration environment, information delivery and analysis; and information management infrastructure that supports data architecture/integration to store, cleanse and transform data.

Position and Adoption Speed Justification: Initially, the on-demand vendors focused on CRM applications, such as salesforce.com and Siebel CRM On Demand; travel and expense management, such as Concur Technologies; and Web analytics software, such as Webtrends and Google Analytics. SaaS-based BA enables business users — particularly small or midsize businesses (SMBs) and enterprise departments — to quickly and easily implement solutions. These products can provide front-end reporting and analysis, as well as back-end integration and data management, without significant IT involvement. Lower startup costs, easier support without an IT contribution and easier deployments, rather than sophisticated functionality, are still driving the adoption of SaaS/PaaS-based BA. Additionally, the overall rise in cloud-based business applications drives BA SaaS/PaaS as clients don't want to bring cloud data on-premises for BA activity.

During the past few years, there has been considerable growth in the number of vendors offering BA capabilities as a service, including 1010data, Adaptive Planning, Birst, Domo, GoodData, Host Analytics, Pervasive Software, PivotLink and SAP BusinessObjects. Over one third of companies Gartner surveyed as part of the 2013 "Magic Quadrant for Business Intelligence Platforms"indicated they already have, or plan to have, a SaaS/PaaS BA product in 2013. The percentage of revenue from products that offer BA as a service continues to grow significantly, yet it remains small compared with the overall BA platform, performance management and analytic applications market. There remains a conflict between the compelling value proposition that BA provides — the ability to create business intelligence (BI) solutions in a scalable manner without an upfront investment — and the two primary barriers to SaaS-based BA:

- Trusting data to a third party
- Overcoming the inertia involved in moving to a new architecture

SaaS-based BA is over the Peak of Inflated Expectations and slipping toward the Trough of Disillusionment, and this is reflected in customer sentiment citing major migration concerns and lack of configurability as barriers to more widespread adoption. However, we continue to see a considerable increase in the number of independent and mainstream vendor cloud offerings, and in client acceptance of this form of pricing and delivery model. Partly, this is fueled by the increased demand for analytics for big data initiatives, analysis of Web content such as social analytics, and support for more unstructured data. Service providers such as Deloitte, IBM, Accenture and others have organized themselves to provide BI and analytics as a service. Their value proposition — specifically, best practice, intellectual property and technology support — plays well with business buyers seeking specific outcomes rather than focusing on the mechanics of business analytics. Data providers — those that deliver critical information by industry and role (for example, Nielsen for consumer packaged goods data, IMS for prescription statistics, and so on) — are also stepping up to deliver more SaaS BA with solutions built around their core asset: the data.

Early adopters have been departments or lines of business of large enterprises, as well as SMBs. The size and sophistication of implementations continues to expand to become more enterprisewide, but these have yet to match the breadth and scalability of on-premises alternatives. Many enterprise BA platform providers have introduced SaaS/cloud offerings, but they are largely independent offerings and should be evaluated accordingly.

User Advice: Enterprises with more straightforward requirements should consider using BA as a service. SaaS-based BA tends to do better for more narrowly defined use cases, as opposed to broad deployments (for example, as a replacement for the entire BI/data warehouse infrastructure or those needing to use analytics for big data and unstructured data). Those with, typically, more complex requirements should evaluate traditional vendor offerings and consider consuming on-premises software, but on a hosted basis.

On-demand BA poses integration challenges for enterprises that need to export and extract data (and metadata) from the service provider for incorporation into their BI infrastructures. IT managers must ensure that the SaaS vendors they select provide assurances through adequate SLAs; govern how a SaaS-based BA applications will coexist with internal applications and infrastructures; and overcome data management, privacy and security issues. Furthermore, IT organizations should

consider contingency plans if they decide to drop their SaaS-based BI vendors — particularly with regard to how they will get their data back and make it functional again.

Business Impact: Business users are becoming increasingly frustrated with the long deployment cycles, high costs, complicated upgrade processes and IT infrastructures demanded by traditional BA solutions. SaaS-based BA offers a quick, low-cost and easy-to-deploy alternative that has proved popular, particularly in SMBs and in the individual departments of large enterprises that lack, or do not want to support, internal BA resources.

Instead of disrupting the enterprise BA platform (and CPM suite) market, a more likely scenario is for SaaS-based BA to tap into new opportunities, particularly with midmarket companies that have yet to invest in BI. However, there has been limited growth in this area during the past year. The penetration of SaaS-based BA into the midmarket will be a key factor in the time taken for this technology to become mainstream.

Benefit Rating: Moderate

Market Penetration: 1% to 5% of target audience

Maturity: Adolescent

Sample Vendors: 1010data; Adaptive Planning; Birst; Domo; GoodData; Host Analytics; Pervasive Software; PivotLink; SAP BusinessObjects

Recommended Reading: "Market Profile: Business Analytics in the Cloud"

"Magic Quadrant for Business Intelligence and Analytics Platforms"

Entity Resolution and Analysis

Analysis By: Mark A. Beyer

Definition: The capability to resolve multiple labels for individuals, products or other noun classes of information into a single resolved entity, to analyze relationships among resolved entities. Advanced forms of entity analytics include the ability to reverse the resolution when appropriate and communicate that reversal to all client systems.

Position and Adoption Speed Justification: As a technology and practice, entity resolution and analysis (ER&A) forms a bridge between expectations and practices regarding the input and management of data across various systems and even different information channels (for example, social versus back-office, customer-driven portals versus operations management). With the rise of cloud deployments combined with on-premises operations applications in many large organizations, the difficulty of determining whether many instances of the same individual are a single entity is increasing. At the same time, the demand for this type of resolution is becoming even more important to business. It is also becoming more difficult to determine whether networks of individuals are composed of multiple parties or if, in fact, they are composed of fictitious personas or aliases. This situation is complicating fraud detection but also the identification of individuals and networks of people for improved communication, marketing and customer service.

In 2012, organizations continued to focus on resolution, rather than analytics, which has given rise to the false hope that current data quality or master data management (MDM) practices will resolve these issues. Data architects and information managers have begun to ask "Why isn't entity resolution part of data quality?" and "Why doesn't MDM also do this as a set of functions?" The answer is that ER&A should be a shared data quality and MDM function, and will most likely be driven by data quality practices discovering issues and MDM resolving them. Some vendors claim they accomplish ER&A simply through data quality practices and MDM on either side, but then fail to complete the workflow expectation of resolving networks into individuals, and the reverse, and then tracking those changes in all client systems.

ER&A has significant benefits. However, it will merge into data quality technology and MDM practices between 2015 and 2016. Additionally, analytics technology that extracts the identifying information — used in reconciling entities from audio, video, image and other unstructured data types — is needed for the technology to fulfill its inherent potential.

User Advice: Fraud detection and other criminal investigation techniques can use ER&A solutions to enhance their use of available information. Data quality stewards and data governance experts can use them as forensic tools. Through 2014, commercial applications of this technology will be challenged because data quality problems will diminish the interest of the analytics portion of ER&A, and only half the problem will be solved.

Use of this technology should include the adoption of data stewardship and dedicated analysts to interpret the results. It is important to understand that the greatest benefit from this class of technology emerges from both facets (entity resolution and analytics), and that the use of entity resolution tools only sets the stage for analytics. In fact, failure to include analytics (specifically the absence of network awareness) will cause this technology to drop into the Trough of Disillusionment, while entity resolution moves up the Slope of Enlightenment as a much-reduced, "battle scarred," stand-alone solution.

Business Impact: Leading organizations in all industries should begin adoption of ER&A, even though it will eventually be subsumed by other data management techniques. Consumers will become increasingly intolerant of service models that fail to recognize their identity and by "voting with their wallets" will migrate to competitors that exhibit a higher level of recognition of the overall business relationship. Identity resolution and analytics will be a critical ingredient of information-aware organizations in the global market as it contributes to customer satisfaction, new customer acquisition and the exclusion of undesirable customers. Resolving identity issues will enable organizations to experience significantly improved relationships with their clients and business partners.

Benign uses of ER&A (for example, alternative bill-of-materials analysis and consumer reference networks) will see rapid results from using this technology; because it quickly identifies data issues and suppliers, and service partners and distribution partners will be more able to communicate with the business. Additionally, after resolving individual identities and the networks of activity perpetrated by individuals, it becomes possible to detect the presence of shadow processes that may actually replace formal or approved processes in an organization. One of the potentially big benefits lies in the development and analysis of social networks.

Unfortunately, criminals will exercise their creativity in circumventing these systems. Linking this technology with voice and speech recognition solutions, as well as text analysis, will help organizations stay ahead of criminal countermeasures.

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: IBM; Infoglide; Informatica

Recommended Reading: "Magic Quadrant for Data Integration Tools"

"Magic Quadrant for Data Quality Tools"

"Incorporating the Web Into Cross-Channel Customer Analysis"

Text Analytics

Analysis By: Daniel Yuen; Hanns Koehler-Kruener

Definition: Text analytics is the process of deriving information from textual sources. It is used for several purposes, such as summarization (finding the key content in a single document or a larger body of information), sentiment analysis (identifying the nature of commentary on an issue), explication (explaining what is driving that commentary), investigation (discovery of the causes of a specific issue) and classification (of the text's subject or key content).

Position and Adoption Speed Justification: Text analytics has appeared in many Hype Cycles over the years, and there has been both renewed interest and disillusionment with the technology. However, interest in the application of analytics, particularly in several areas, has resulted in growing interest among users and more robust tools and integrations into traditional content management systems. These areas are customer sentiment, which looks at direct customer feedback (for example, from surveys) and indirect customer feedback (such as blog commentaries and Facebook and Twitter posts); fraud and public security (identifying patterns that require further investigation); and categorization and classification (which involves trying to understand how a particular text objects fit into an existing structure, or creating structures from scratch).

The emergence of speech analytics, for which text analytics is one of the key enabling technologies, will help improve adoption of text analytics in organizations. Furthermore, big data (the ever-increasing volume, velocity and variety of data) demands innovative analytical approaches to derive value from it. Although still not the "black box" or magic tool that companies hoped for, text analytics has nevertheless passed the stage of unsubstantiated hype and will become a tool used by organizations everywhere.

User Advice: There is no one universal approach to analyzing text-based content. Some text analytics tools perform language parsing and entity recognition on the basis of linguistic models with or without preset taxonomies and dictionaries to build a structured view of information. Others

rely on natural-language processing (NLP) to understand text. Text analytics can also be performed with machine learning.

Users need to spend time understanding the individual capabilities, the targeted business outcomes and where the necessary information resides before investing in a toolset. Requirements for support of multiple and non-English languages, the ability to export results into standard formats and easy-to-understand user interfaces all make the selection process a crucial exercise. Providing mixed-language support for individual documents can still be a challenge for some vendors. A proof of concept with the client organization's own data and needs is essential to achieve the best results. Organizations considering investing in this space should do so incrementally. They should consider limited pilots to ensure that the application will work effectively in the chosen domain, and to establish business processes for dealing with the results, before investing in a full-scale deployment.

Business Impact: Text analytics can be used in a wide variety of use cases in multiple parts of a company. Sentiment analysis can provide insight into customer satisfaction across multiple channels. Monitoring for particular terms or phrases can help to pinpoint areas of concern. Using text analytics to classify a body of text during ingestion or upload may improve "findability" in the future or allow for greater control in information governance. Big data analytics, especially combining structured and unstructured data sources, can provide more insights into new opportunities and missed opportunities.

There are numerous potential use cases for which text analytics, combined with any of the many other analytic capabilities, will have a significant impact on an organization only if it is aware of the questions it wants answering and is willing to act on the results received. This is something each organization should explore carefully before embarking on a text analytics initiative. In many cases, the largest impact of text analytics will come from allowing it to combine extracted unstructured data with traditional structured data to provide a more complete view of an issue that could be obtained using traditional data mining or business intelligence tools.

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: Attensity; Autonomy; Clarabridge; IBM; IxReveal; Megaputer; Microsoft; Nice Systems; OpenText; Oracle; SAP (BusinessObjects); SAS; Temis

Recommended Reading: "Cool Vendors in Content and Social Analytics, 2013"

"Who's Who in Text Analytics"

"How BI Leaders Can Get Started With Text Analytics"

"Text Analytics Guidance: Building a Text Analytics Program"

In-Memory Analytics

Analysis By: Kurt Schlegel

Definition: In-memory analytics is an alternative business intelligence (BI) performance layer in which detailed data is loaded into memory for fast query and calculation performance against large volumes of data. This approach obviates the need to manually build relational aggregates and generate pre-calculated cubes to ensure analytics run fast enough for users' needs.

Position and Adoption Speed Justification: Declining memory prices, coupled with widespread adoption of 64-bit computing, continue to prime the market for in-memory analytics. Most BI vendors are now positioning in-memory analytics as a key component of their BI offerings and the use of DRAM and NAND Flash memory to speed up analytics will soon be ubiquitous as part of vendor platforms. Forward movement along the Trough of Disillusionment reflects the acceptance of in-memory analytics with more use cases and clear performance benefits becoming apparent.

In-memory analytics is no longer a fringe technology, but it is increasingly becoming the dominant performance layer for BI and analytic application architectures. The time taken to reach the Plateau of Productivity was less than two years (previously this technology remained two to five years away from the plateau for several years). Increasing levels of adoption mean it is likely to peak at the plateau by 2015.

User Advice: For response-time issues and bottlenecks, IT organizations should consider the performance improvement that in-memory analytics can deliver, especially when run on 64-bit infrastructure. Users should be careful to use in-memory analytics as a performance layer and not as a substitute for a data warehouse. In fact, users considering utilizing in-memory analytics should also be aware of how their requirement for speedier query processing and analysis could be addressed by the use of in-memory processing in the underlying databases feeding BI or via in-memory databases or data grids.

BI and analytic leaders need to be aware that in-memory analytics technology has the potential to subvert enterprise-standard information management efforts through the creation of in-memory analytic silos. Where it is used in a stand-alone manner, organizations need to ensure they have the means to govern its usage and that there is an unbroken chain of data lineage from the report to the original source system, particularly for system-of-record reporting.

Finally, it is becoming apparent as the scale of in-memory analytics deployments grows, there is still a need for performance tuning, either by the return of some aggregation at data load, or by managing application design against user concurrency requirements and the sizing of hardware and available RAM.

Business Impact: BI and analytic programs can benefit broadly from the fast response times delivered by memory-based processing and this in turn can improve the end-user adoption of BI and analytics. The reduced need for database indexing and aggregation enables database administrators to focus less on the optimization of database performance and more on value-added activities. Additionally, in-memory analytics by itself will enable better self-service analysis because there will be less dependence on aggregates and cubes built in advance by IT.

However, from an analyst user perspective, faster queries alone are not enough to drive higher adoption. In-memory analytics is of maximum value to users when coupled with interactive visualization capabilities or used within data discovery tools for the highly intuitive, unfettered and fast exploration of data.

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Adolescent

Sample Vendors: IBM; Microsoft; MicroStrategy; Oracle; QlikTech; SAP; SAS; Tibco Software

Recommended Reading: "Need for Speed Powers In-Memory Business Intelligence"

Real-Time Decisioning Platform

Analysis By: Gareth Herschel

Definition: Real-time decisions combine predictive analytic and decision capabilities to identify the optimal next action to take in a real-time process. The analysis can be based on a variety of approaches, but the solution must include an arbitration (rule) capability to select the optimal next action (based on the enterprise's strategy) of several possible options.

Position and Adoption Speed Justification: The concept behind this technology is not new. Credit decisions and some complex machines (such as those for aircrafts) have made automated (rule-based) decisions, and some of these decisions have been made on the basis of analysis. The current position on the Hype Cycle is based on the movement of vendors to create more integrated solutions, using more sophisticated predictive analytics tools and more robust rule management capability that allows for more complex decision logic. The adoption of this capability into a wide variety of business processes (such as insurance claims processing, field service dispatch, warranty analysis or customer retention) has been slow, but it is very likely to accelerate as the concept is popularized across organizations through applications such as "best next offer" for customer interactions.

User Advice: This concept can be built internally using combinations of stand-alone components (primarily a predictive analytics capability and a rule management system), but these solutions will be increasingly available as a platform to shorten the time to deploy for organizations. This concept can be applied to any process where the appropriate course of action will vary, depending on the changing context of the process, or the attributes of the subject of the process (for example, customer or claim).

Business Impact: Early adoption of this concept has primarily focused on issues such as insurance claims or credit approvals. More recent adoption has focused on the contact center, turning a purely service-oriented interaction into a blend of marketing, sales and service. Typical business benefits felt by this type of investment are:

- More appropriate decisions for the current circumstances, resulting in less waste (reducing the frequency of always taking an action because it is sometimes needed).
- Increasing the value of the process (taking an action because it would be appropriate in this case but not cost-effective or appropriate to take it every time).
- Increasing the accuracy of a decision (taking an action based on better analysis of the situation than would have been possible with a rule engine alone).

Benefit Rating: High

Market Penetration: 5% to 20% of target audience

Maturity: Early mainstream

Sample Vendors: FICO; IBM; Oracle; Pegasystems; SAS

Recommended Reading: "How to Achieve Real-Time CRM"

"Select Customer Data Mining Vendors Based on Focus and Vision"

"Automating Decisions With Intelligent Decision Automation"

Climbing the Slope

Interactive Visualization

Analysis By: Kurt Schlegel

Definition: Interactive visualization technology enables the exploration of data via the manipulation of chart images, with the color, brightness, size, shape and motion of visual objects representing aspects of the dataset being analyzed. Interactive visualization products provide an array of visualization options that go beyond those of pie, bar and line charts, including heat and tree maps, geographic maps, scatter plots and other special-purpose visuals. These tools enable users to analyze data by interacting with visual representations of it.

Position and Adoption Speed Justification: Representing complex or multidimensional data on a 2D screen and giving users the chance actively to explore the data visually helps them understand and assimilate it more effectively than rows, columns, figures and static charts ever could. Techniques like brushing, trellising, map overlays and time-series animation all help data come alive and are compelling to end users. Interactive visualization techniques are being adopted by government and commercial organizations — survey data gathered for Gartner's 2013 "Magic Quadrant for Business Intelligence and Analytics Platforms"found widespread use of these capabilities. Gartner analysts are seeing increased evaluation and adoption of these techniques as part of business intelligence (BI) solution architectures, particularly as these capabilities become available from large BI vendors (such as Microsoft, SAP and SAS). As such, interactive visualization should reach the Plateau of Productivity in less than two years, with a key driver being the growing adoption of HTML5.

Note that there is a difference between interactive visualization and visualization-based data discovery tools. Interactive visualization is a component of these data discovery tools, but the terms are not synonymous.

User Advice: Interactive visualization capabilities are now offered by the majority of BI megavendors and large independent BI vendors — this is the major development in this area in recent years. BI leaders should evaluate the offerings of their incumbent BI platform vendor in relation to interactive visualization and compare these to the products of independent specialists.

BI competency centers must ensure that any stand-alone visualization technologies integrate well with their BI architectures and are easy for the enterprise to use. In the same way as IT groups are struggling to rein in "spread marts," IT organizations must ensure that visualization applications are supported and developed as part of their portfolio of BI standards.

Business Impact: Interactive visualization technologies deliver information more effectively, enabling users quickly to ascertain trends and anomalies in data, based on visual cues. They also enable a greater degree of user self-service — probably its most impactful characteristic.

Adopting interactive visualization tools will also enable organizations to leave analysis of complex sets of data to subject matter experts who have strong contextual indications of the importance of this data for making decisions. This should improve the accuracy and speed of decision making.

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Sample Vendors: Advizor Solutions; GroupVisual.io; Microsoft; MicroStrategy; Panopticon; QlikTech; SAP (BusinessObjects); Tableau; Tibco Software

Recommended Reading: "Who's Who in Interactive Visualization for Analysis and Dashboarding"

Entering the Plateau

Column-Store DBMS

Analysis By: Donald Feinberg

Definition: A column-store DBMS indexes each column of a table, storing them in lieu of row data, unlike traditional relational DBMSs using a row-store, where data is stored in rows, with indexes optional. In addition, most column-store DBMSs include additional optimization techniques (such as compression and tokenization) to further compress the data — using less storage and increasing input/output (I/O) performance. We do not include row-store DBMS engines that offer columnar forms of compression only.

Position and Adoption Speed Justification: Like many technologies, column-store DBMS technology is a mix of mature products (such SAP Sybase IQ) and newer products (such as those from 1010data, Exasol, Infobright, ParAccel, SAP Hana and Vertica). There are currently thousands of installations of column-store DBMSs used for multiple purposes, from small data marts to full data warehouses (DWs).

With the growing tendency to use data marts combined with increasing use of analytics, the column-store DBMS has gained market share as an analytic engine, leading to a growing number of new products (see "Magic Quadrant for Data Warehouse Database Management Systems"). Most analytic applications require a fewer number of columns rather than the entire row when retrieving data and this, coupled with the compression abilities of the column-store DBMS, leads to a lower amount of I/O and higher performance.

An additional driver of the maturing position of this technology is that these DBMS engines are SQL DBMSs and for the most part, atomicity, consistency, isolation, durability (ACID)-compliant. Those that are not ACID-compliant fit the base model and are still fully consistent models. Support for full SQL and consistency implies that they can be used interchangeably with row-store DBMSs without changes or the need to support an eventually consistent model.

During the past few years, we have seen increased use of column-store DBMS solutions as the DW, primarily due to vendors enabling more sophisticated workload management software. With workload management and the high compression available, the column-store can handle more complex workloads found in current DW environments. Several vendors have added column-store technology to their DBMS products for this reason, notably EMC/Greenplum, IBM DB2, Microsoft (SQL Server 2012) and Teradata.

The other major technical breakthrough has come with creative methods of loading data into the column-store — traditionally an issue due to the high number of disk writes required. Many vendors have achieved much faster loading techniques, allowing for near-continuous loading of data.

Finally, due to the high level of compression, the column-store has begun to be used as an in-memory DBMS, where terabytes of source data can fit into gigabytes of memory. Recently we have seen a number of new entries to the in-memory column-store DBMS sector, notably from IBM, ParStream and SAP Hana and Teradata.

With the ever-growing use of column-store DBMSs and the introduction of in-memory column-store offerings, this technology is approaching the Plateau of Productivity. Due to the number of new entries (such as IBM BLU [2013]), Microsoft SQL Server 2012 xVelocity (2012), SAP Hana (2011) and Teradata (2013), we have decided not to place it on the plateau yet. Additionally, column-store DBMS alone is not transformational. Moving the column-store to an in-memory model makes the column-store DBMS usable as a transaction DBMS.

User Advice: Column-store DBMSs should be considered:

- As primary candidates for analytic data marts — because of I/O performance gained from compression and retrieval of fewer columns (typical of analytic applications).

- In an overall archiving strategy — as one of the "near line" alternatives, because of their high compression ratios while making the data available with standard relational tools.
- For data warehouse implementations — as they continue to improve mixed-workload management capabilities and the ability to load data with much lower latencies — approaching continuous loading.
- For leading-edge organizations: begin using in-memory column-store DBMS implementations for applications requiring high performance. Many of these applications cannot be implemented on a traditional DBMS due to the resources and execution time needed.

Business Impact: As the column-store DBMS broadens its appeal within IT architecture, so its impact on the business and IT is also growing. The column-store's original function was in archiving solutions, although that is no longer the primary use case. The automatic compression achieved by storing data in column form (reported by clients to be as great as 20 to one, or more), makes column-store DBMSs a good choice for moving data out of the primary storage DBMS, while maintaining the relational structure of the data.

In addition to the obvious advantages of the reduced storage necessary for data (a clear cost saving), compression has also been shown to increase performance, due to much lower I/O. This also has implications for reducing the size of servers required for the DW — another clear cost saving.

As hardware systems with large amounts of memory become more available during the next few years and costs decline, we will see increased use of the column-store as an in-memory DBMS structure. In addition, tests are showing that in-memory, column-store technology reduces the need for complex structures such as aggregation, summaries and cubes, reducing maintenance while further increasing performance.

This development will not only allow for faster performance for many classes of application in the fields of business intelligence and analytics, but will also enable the column-store to be used as an online transaction processing (OLTP) DBMS (such as with SAP Hana). In the past, this use has been prevented by the large number of disk writes required to insert or update a row — but this is no longer an issue when handled in-memory.

The potential for the column-store to be used for both OLTP and the DW has far-reaching implications, not only for the design of systems but also in terms of major cost savings in comparison with slower disk systems with higher power and cooling requirements. As column-store technology continues to mature and its use continues to grow, it has entered the early mainstream maturity phase as it moves toward the plateau and shortly, off of the Hype Cycle.

Benefit Rating: Transformational

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Sample Vendors: 1010data; Actian; EMC; Exasol; HP; IBM; Infobright; Microsoft; ParStream; SAP; Teradata

Recommended Reading: "Magic Quadrant for Data Warehouse Database Management Systems"

"The State of Data Warehousing in 2013 and Beyond"

Predictive Analytics

Analysis By: Gareth Herschel

Definition: The term "predictive analytics" has become generally used to describe any approach to data mining with four attributes: an emphasis on prediction (rather than description, classification or clustering); rapid time-to-insight (measured in hours or days); an emphasis on the business relevance of the resulting insights; and an increasing emphasis on ease of use — thus making the tools accessible to business users.

Position and Adoption Speed Justification: The algorithms underpinning predictive analytic applications are reasonably mature. Although new techniques continually emerge from research laboratories, the 80/20 rule firmly applies with most of the commonly used algorithms (such as chi-squared automatic interaction detection [CHAID] decision trees and k-means clustering) that have been in existence for more than a decade. The applications are also approaching maturity, although the development of packaged applications to address specific business problems (compared with the generic approach of turning more-traditional data mining workbenches into predictive analytic solutions) is less mature and more diverse in its maturity. When predictive analytic applications have added project and model management capabilities with more enhancements to aid ease of use, they will have achieved maturity.

User Advice: Predictive analytics is a more user-friendly and business-relevant equivalent of data mining that is applied specifically to predictions of future behavior. Although potentially lacking some of the mechanisms to fine-tune the model performance that a traditional data mining workbench might deliver, the benefits of rapid model development and easier maintenance are appealing for most analytical initiatives. The bigger distinction is between predictive analytic solutions and packaged applications built on these solutions for specific business issues. In these cases, the selection decision should be based on the domain expertise that the vendor has been able to package into the application, versus the domain expertise the business analyst can bring to the analysis.

Business Impact: Predictive analytics can bring clarity and consistency to any situation where the likely future behavior or condition is uncertain. Common applications include understanding the future behavior of customers (Will they renew the relationship? Which products or services are they likely to buy?), the future state of customers (What will their lifetime value be to the company?) or to predict the likely performance of equipment (predictive maintenance allows the identification of at-risk components so they can be proactively replaced). By understanding likely future circumstances, organizations are better able to allocate investments to maximize returns.

Benefit Rating: High

Market Penetration: 20% to 50% of target audience

Maturity: Early mainstream

Sample Vendors: Angoss; FICO; IBM (SPSS); KXEN; Pitney Bowes Software; Revolution Analytics; SAS; StatSoft

Recommended Reading: "How to Increase the Volume of Advanced Analytics"

Dashboards

Analysis By: Kurt Schlegel

Definition: Dashboards are reporting mechanisms that aggregate and display metrics and key performance indicators (KPIs), enabling them to be examined at a glance by all manner of users before further exploration via additional business analytics tools. Dashboards help improve decision making by revealing and communicating in-context insight into business performance, displaying KPIs or business metrics using intuitive means of visualization such as charts, dials, gauges and "traffic lights" that indicate progress toward defined targets.

Position and Adoption Speed Justification: Dashboards provide an easy-to-understand presentation layer for business analytics tools and many business applications that is visually attractive and intuitive to users. Business consumers often demand that performance data be delivered to them in the context of their role or business activity, and dashboards suit this need well. Increasingly, mobile devices are the medium of choice, with cloud-based offerings also emerging. In most business analytics deployments, dashboards are ubiquitous — they are one of the primary ways in which business users receive information.

However, dashboards are sometimes implemented as tools that are not properly connected to underlying data sources and systems. In these circumstances, they fail to deliver much benefit and fall into disuse because they cannot easily adapt to changes in the business (much like the executive information systems of 20 or more years ago). Increasingly, users are realizing that dashboards are especially valuable when implemented as part of a broader business analytics strategy. But they are often a discrete buying agenda item; as such, some companies will continue to implement dashboards on a stand-alone basis. In short, for dashboards to have a business impact, users need to care about the measures.

Dashboards are a very mature technology, with virtually every business intelligence and analytics platform vendor offering them. For this reason, we have moved it further along the Hype Cycle. Moreover, user adoption has been very strong. Therefore, we expect it to reach the Plateau of Productivity in less than two years, and probably sooner rather than later.

User Advice: Evaluate use cases for dashboards in your organization, as they will be varied. Test visualization metaphors and interaction models to ensure they meet the minimum functionality that business users want. Also, determine which medium is the best "fit" for consumers, given that mobile and cloud options are viable, in addition to traditional laptop and desktop PC deployments.

Be aware that support for different media will require dedicated effort to optimize the rendering of content and interaction.

KPIs and metrics are important elements of any dashboard; do not underestimate how much effort it may take to reach agreement on metrics definitions within your company. In this case, technology is not the limiting factor — achieving consensus about the metrics that really matter to the business is. Monitor the use of dashboards and allocate time and resources accordingly in order to provide a continuous improvement program. As far as possible, tie dashboards to specific decisions in order to make them actionable.

Business Impact: Dashboards can make it easy for senior executives, managers and business users to understand quickly how their organization is performing against its business objectives.

Media tablets and smartphones are increasingly deployment targets for new constituents; these devices can be granted live connections to data, as well as equipped with a disconnected mode in which the necessary data is encapsulated in an object.

Dashboards can be deployed at any level of an organization, and are good tools for fostering discussion about action plans to achieve goals. They can also be used to promote collaboration outside an enterprise by sharing KPIs with customers, suppliers and partners.

When widely adopted in an organization, dashboards are an effective aid to leadership and to aligning people and resources to meet strategic objectives. They can have a significantly positive impact on the monitoring of corporate and operational performance and on the taking of appropriate action when certain conditions are indicated by dashboard alerts.

Benefit Rating: High

Market Penetration: More than 50% of target audience

Maturity: Mature mainstream

Sample Vendors: Domo; GoodData; IBM (Cognos); Information Builders; Microsoft; MicroStrategy; Oracle; QlikTech; SAP (BusinessObjects); SAS; Tableau; Tibco Software

Recommended Reading: "Who's Who in Interactive Vizualization for Analysis and Dashboarding"

"Scorecard or Dashboard: Does It Matter?"

"Tips for Implementers: The Basics of Good Dashboard Design"

"Magic Quadrant for Business Intelligence Platforms"

Appendixes

Hype Cycle Phases, Benefit Ratings and Maturity Levels

Table 1. Hype Cycle Phases

Phase	Definition
Innovation Trigger	A breakthrough, public demonstration, product launch or other event generates significant press and industry interest.
Peak of Inflated Expectations	During this phase of overenthusiasm and unrealistic projections, a flurry of well-publicized activity by technology leaders results in some successes, but more failures, as the technology is pushed to its limits. The only enterprises making money are conference organizers and magazine publishers.
Trough of Disillusionment	Because the technology does not live up to its overinflated expectations, it rapidly becomes unfashionable. Media interest wanes, except for a few cautionary tales.
Slope of Enlightenment	Focused experimentation and solid hard work by an increasingly diverse range of organizations lead to a true understanding of the technology's applicability, risks and benefits. Commercial off-the-shelf methodologies and tools ease the development process.
Plateau of Productivity	The real-world benefits of the technology are demonstrated and accepted. Tools and methodologies are increasingly stable as they enter their second and third generations. Growing numbers of organizations feel comfortable with the reduced level of risk; the rapid growth phase of adoption begins. Approximately 20% of the technology's target audience has adopted or is adopting the technology as it enters this phase.
Years to Mainstream Adoption	The time required for the technology to reach the Plateau of Productivity.

Source: Gartner (July 2013)

Table 2. Benefit Ratings

Benefit Rating	Definition
Transformational	Enables new ways of doing business across industries that will result in major shifts in industry dynamics
High	Enables new ways of performing horizontal or vertical processes that will result in significantly increased revenue or cost savings for an enterprise
Moderate	Provides incremental improvements to established processes that will result in increased revenue or cost savings for an enterprise
Low	Slightly improves processes (for example, improved user experience) that will be difficult to translate into increased revenue or cost savings

Source: Gartner (July 2013)

Table 3. Maturity Levels

Maturity Level	Status	Products/Vendors
Embryonic	▪ In labs	▪ None
Emerging	▪ Commercialization by vendors Pilots and deployments by industry leaders	▪ First generation High price Much customization
Adolescent	▪ Maturing technology capabilities and process understanding Uptake beyond early adopters	▪ Second generation Less customization
Early mainstream	▪ Proven technology Vendors, technology and adoption rapidly evolving	▪ Third generation More out of box Methodologies
Mature mainstream	▪ Robust technology Not much evolution in vendors or technology	▪ Several dominant vendors
Legacy	▪ Not appropriate for new developments Cost of migration constrains replacement	▪ Maintenance revenue focus
Obsolete	▪ Rarely used	▪ Used/resale market only

Source: Gartner (July 2013)

Figure 3. Hype Cycle for Business Intelligence, 2012

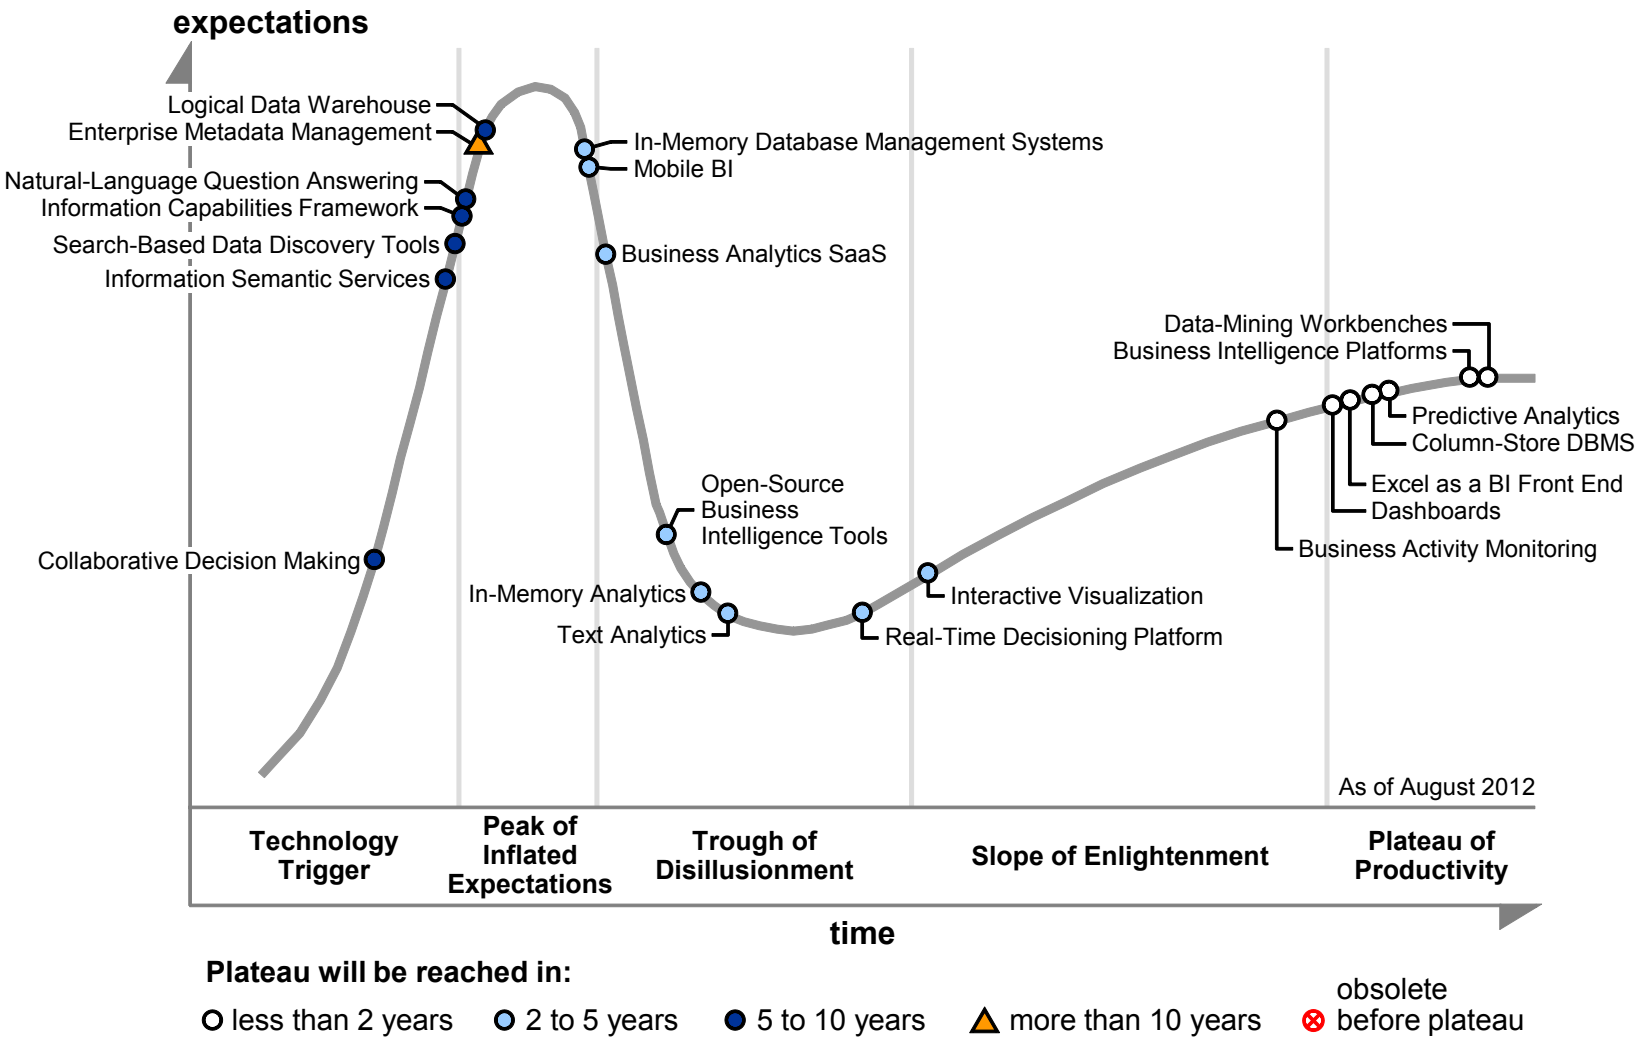

BI = business intelligence; DBMS = database management service; SaaS = software as a service

Source: Gartner (September 2012)

Recommended Reading

Some documents may not be available as part of your current Gartner subscription.

"Understanding Gartner's Hype Cycles, 2013"

"Hype Cycle for Information Infrastructure, 2013"

"Hype Cycle for Performance Management 2013"

"Hype Cycle for Analytic Applications, 2013"

"Hype Cycle for Big Data, 2013"

More on This Topic

This is part of an in-depth collection of research. See the collection:

- Gartner's Hype Cycle Special Report for 2013

GARTNER HEADQUARTERS

Corporate Headquarters
56 Top Gallant Road
Stamford, CT 06902-7700
USA
+1 203 964 0096

Regional Headquarters
AUSTRALIA
BRAZIL
JAPAN
UNITED KINGDOM

For a complete list of worldwide locations,
visit http://www.gartner.com/technology/about.jsp