FORRESTER®

For: Application Development & Delivery Professionals

The Forrester Wave™: Enterprise Data Warehouse, Q4 2013

by Noel Yuhanna and Mike Gualtieri, December 9, 2013

KEY TAKEAWAYS

Next-Generation EDW Delivers In-Memory, Cloud, and Hadoop Integration

EDW is expanding to support higher scalability, higher performance, deeper Hadoop integration, and more automation. Large EDW vendors have started to offer more scalable platforms that use solid-state drives and offer distributed in-memory technologies to process large amounts of data faster, delivering real-time analytics and predictive analytics.

EDW Appliances And Data Warehousing-as-a-Service Are Gaining Momentum

EDW appliances have been gaining significant momentum thanks to their stronger integration of hardware and software, improved automation, single-vendor support, higher performance and scale, and reasonable prices. DWaaS provides a low-cost alternative to on-premises EDW solutions.

Enterprise Maturity And Innovation Distinguish The EDW Leaders

The Leaders we identified offer mature, high-performance, scalable, secure, flexible, and robust EDW solutions. The Strong Performers have turned up the heat as high as it will go on the incumbent Leaders, with innovations that many customers find compelling.

The Forrester Wave™: Enterprise Data Warehouse, Q4 2013

by Noel Yuhanna and Mike Gualtieri
with Holger Kisker, Ph.D. and Sarah Bookstein

WHY READ THIS REPORT

Although enterprise data warehouse (EDW) is a mature technology, vendors continue to innovate and enhance its real-time analytics, distributed in-memory, advanced compression, and scalable appliance features and its integration with Hadoop, data virtualization, and cloud to support new and emerging use cases. Today's most demanding EDW environments support petabytes of aggregated data, billions of records, complex mixed-query workloads, and subsecond queries. In Forrester's 43-criteria evaluation of EDW vendors, we evaluated 11 solutions from Actian, Amazon Web Services (AWS), HP, IBM, Kognitio, Microsoft, Oracle, ParAccel, Pivotal, SAP, and Teradata. We scored factors like performance, scalability, integration with Hadoop and data services, security, high availability, mixed workload management, and hardware optimization to give you the decision tools to create the right shortlist for your particular environment and scenarios. This report details our findings about how well each solution fulfills the criteria and where they stand in relation to each other.

Table Of Contents

Notes & Resources

Forrester conducted product evaluations in Q2 2013 and interviewed 10 vendor companies, including Actian, Amazon Web Services, HP, IBM, Kognitio, Microsoft, ParAccel, Pivotal, SAP, and Teradata. We also interviewed two reference customers for each vendor.

Related Research Documents

Information Fabric 3.0
August 8, 2013

Future Of Customer Data Management
March 6, 2013

Forrester Wave™: Enterprise Data Warehousing Platforms, Q1 2011
February 10, 2011

EDW IS GETTING BIGGER AND FASTER AND OFFERING NEW USE CASES

EDW enables organizations to deliver actionable, timely, and trustworthy intelligence to a broader range of business users and operational systems. EDW organizes and aggregates historical analytical data from functional domains — such as customer, manufacturing, finance, and human resources — that align with key processes, applications, and roles. EDW serves as the repository for a substantial amount of an organization's operational history. It offers in-database analytics, predictive models, and embedded business algorithms to drive business decisions. EDW is a robust, secure, and proven ecosystem that supports integration with data models and security frameworks, real-time analytics, automation, and a broad range of business intelligence (BI) and visualization tools. It is the foundation for BI to support timely reports, ad hoc queries, and dashboards and supply other analytics applications with trusted and integrated data. Most organizations that have a DW are either enhancing it or considering new business initiatives that require faster processing and larger storage requirements. Forrester defines EDW as:

> *An enterprise data warehouse is a repository of information that is used for reporting and analytics. It includes key data management functions such as concurrency, security, storage, processing, SQL access, and integration.*

Organizations that buy and implement data warehouses usually do so for one or all of the following reasons: They need to consolidate data from several application transactional databases to support operational analytics, analytics, and predictive analytics; they need much faster processing and support for ad hoc queries; and they need to store increasing volumes of structured and unstructured data.

Next-Generation EDW Focuses On In-Memory, Cloud, Real-Time, And Hadoop

Data warehouses have been around for decades, but vendors continue to deliver innovative solutions to support new and growing data requirements. The next-generation EDW is expanding to support higher scalability, higher performance, deeper integration, real-time analytics, stronger security, and more automation (see Figure 1). Recent EDW innovations and product enhancements have been around:

- **Integrating with in-memory architectures.** Data stored in memory can be accessed orders of magnitude faster than that stored on disk. However, until recently, putting data in memory was not an option because memory was prohibitively expensive compared with disk space. Large EDW vendors have started to offer more scalable platforms that can use a mix of solid-state drives (SSDs) and offer distributed in-memory technologies to process large amounts of data faster, delivering real-time analytics and predictive analytics.

- **Leveraging Hadoop to support larger and more complex data sets.** Hadoop, an open source initiative under the Apache License 2.0, delivers a distributed, scalable data processing platform

to support big data. It supports batch processing of data through the parallel processing of large sets of data on clusters using commodity servers. Many large EDW vendors now offer the ability to integrate their solutions with Hadoop to store and process large amounts of structured and unstructured data. Forrester sees many enterprises already using an "extract-Hadoop-load" approach to extract data from transactional systems, load it into Hadoop, aggregate it using MapReduce, and then load it into the EDW to support integrated analytics and predictive analytics. This approach reduces the storage and processing required on EDW platforms, optimizing costs in the long run.

- **Delivering EDW as a service in the public cloud.** Although EDW in the public cloud is not new, delivering data warehousing-as-a-service (DWaaS) is new and emerging. DWaaS allows organizations to store, process, and access data in the public cloud. Enterprises view cloud platforms as the fastest, most flexible way to deliver new capabilities and services, and DWaaS fits the mold. It automates the provisioning, administration, backup, recovery, availability, security, and scalability of the DW repository without the need for a database administrator and delivers economies of scale through elastic compute capacity, automation, and standardization.

- **Integrating with data virtualization platforms to simplify ingestion.** Data virtualization integrates disparate data sources in real time or near real time to support analytics, predictive analytics, customer personalization, and real-time integrated analytics.[1] It integrates with EDW solutions, ingesting derived data sources after they have been transformed, cleansed, and integrated. Leading vendors supporting data virtualization solutions include Composite Software, Denodo Technologies, IBM, Informatica, Microsoft, Oracle, Red Hat, and SAP.

- **Supporting advanced data compression to manage larger data sets more efficiently.** Although data compression in databases and DWs isn't new, recent innovations have significantly improved compression ratios and reduced system overhead. Enterprises are reporting 2x to 50x compression ratios from leading EDW solutions. Vendors are also offering in-memory data compression and adaptive compression that enables faster analytics by processing larger data sets in memory.

- **Enabling in-database analytics to process complex functions.** Companies are increasingly deploying EDWs at the heart of the next-generation enterprise application platform, executing and integrating analytics and transactional computing functions. The current best-of-breed EDW platforms support these application integration scenarios through features and interfaces such as MapReduce, in-database function pushdown, embedded statistical algorithm libraries, predictive modeling integration, decision automation, and mixed workload management.[2]

Figure 1 The Next-Generation EDW Platform

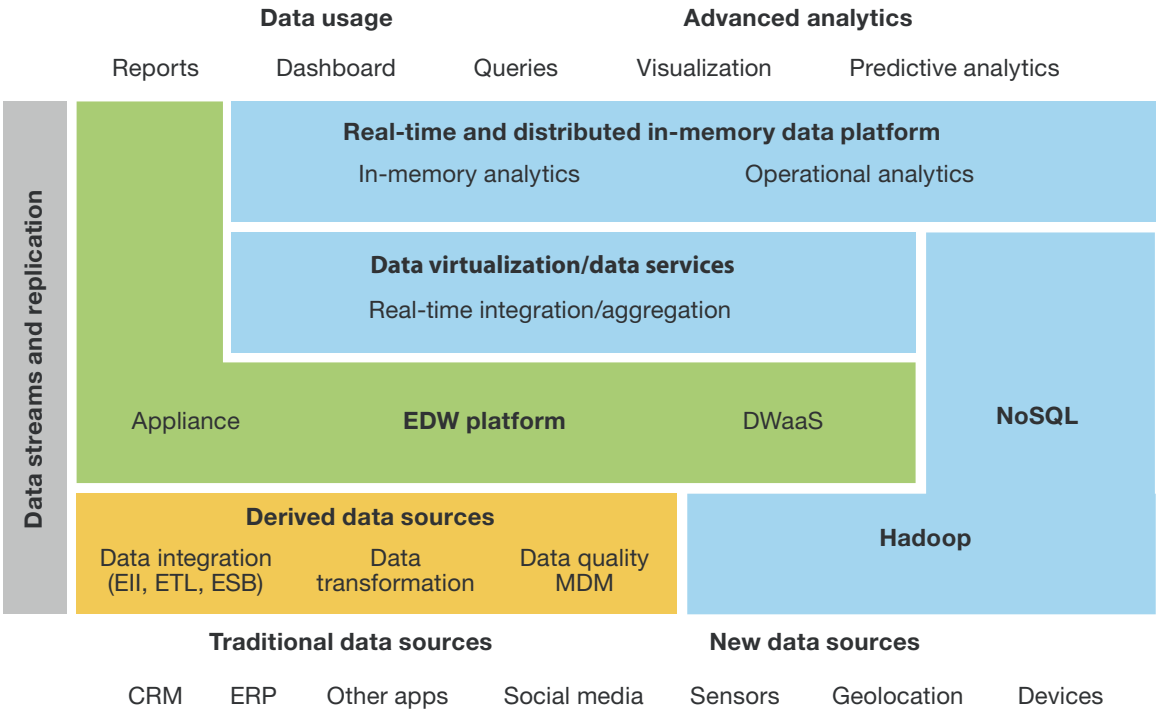

86621 Source: Forrester Research, Inc.

EDW Solutions Are Now Available In Many Form Factors — Not Just Software

EDW appliances have been gaining significant momentum thanks to their stronger integration of hardware and software, improved automation, single-vendor support, higher performance and scale, and reasonable prices. Top EDW vendors offer enterprise-grade appliance-based EDW solutions either directly or through hardware partners. Leading EDW appliance vendors include HP, IBM, Microsoft, Oracle, Pivotal, SAP, and Teradata. EDW solutions are available in four categories:

- **Specialized EDW software that manages large data sets more efficiently.** Vendors like Pivotal, SAP, and Teradata have created specialized EDW software to support large, complex data sets. This category includes EDW software optimized for analytics workloads such as columnar optimized EDW solutions and in-memory analytics processing engines.

- **Traditional DBMS product enhancements that support mixed workloads.** Traditional general-purpose database management system (DBMS) software vendors like IBM, Microsoft, and Oracle have extended their database products over the years to support enterprise-grade

EDW requirements. Using common DBMS software for transactional and analytics can reduce data movement and processing requirements.

- **EDW appliances that optimize software and hardware to deliver automation and scale.** Vendors like HP, IBM, Microsoft, Oracle, Pivotal, SAP, and Teradata offer EDW appliances that optimize servers, storage, memory, and software to deliver faster analytics and predictive analytics. They also automate the administration, tuning, and scaling of DWs.

- **DWaaS offers new economical and deployment options.** These new cloud offerings, which are a combination of the traditional software-as-a-service (SaaS) and infrastructure-as-a-service markets, provide a low-cost alternative to an on-premises EDW solution by automating provisioning, configuration, security, tuning, optimization, scalability, availability, and backup. Vendors offering DWaaS include 1010data, AWS, bitYota, Microsoft, SAP, Teradata, and Treasure Data.

ENTERPRISE DATA WAREHOUSE EVALUATION OVERVIEW

To assess the state of the EDW market and see how the vendors stack up against each other, Forrester evaluated the strengths and weaknesses of top EDW vendors. The EDW market is extremely competitive, because it's become a critical category for enterprise software companies, leading pure-play DW vendors, and big data startups gunning for a piece of the huge pie. Customers will benefit as the pace of innovation increases and the cost per terabyte goes down.

Evaluation Criteria Focus On Scale, Performance, Cloud, And Appliances

After examining past research, user need assessments, and vendor and expert interviews, we developed a comprehensive set of evaluation criteria. We evaluated vendors against 43 criteria, which we grouped into three high-level buckets:

- **Current offering.** To assess the breadth and depth of each vendor's EDW product set, we evaluated each solution's architectural and operational functionality.

- **Strategy.** We reviewed each vendor's strategy to assess how each vendor plans to evolve its EDW solution to meet emerging customer demands. We also evaluated each vendor's go-to-market approach, commitment, and direction strategies.

- **Market presence.** To establish each EDW product's market presence, we evaluated each solution provider's company financials, adoption, and partnerships.

Evaluated Vendors Meet Functional, Architectural, And Market Presence Criteria

Forrester included 11 vendors in the assessment: Actian, Amazon Web Services, HP, IBM, Kognitio, Microsoft, Oracle, ParAccel, Pivotal, SAP, and Teradata. Each of these vendors has (see Figure 2):

- **An enterprise-class data warehouse offering.** We included vendors that offer the following core DW functional components, tools, and features: 1) Support for core EDW features and functionality, including high availability, security, performance, scalability, and management; 2) support for analytical data storage for persistence, integrity, and access; 3) integration with one or more vendors to support BI, reporting, and other analytical solutions; 4) native tools or integration with third-party vendors to support data loading, unloading, transformation, and cleansing of analytical data sets; 5) support for multiple concurrent analytical SQL queries, aggregated data sets, and integrated data access; 6) the ability to deploy on-premises, in the public cloud, or both; and 7) tools that database administrators, data architects, and DW analysts can use to manage the EDW platform.

- **EDW revenue.** We included EDW vendors that actively market an EDW solution. We included vendors with at least $25 million in EDW-specific revenues in the latest fiscal calendar year, where at least 80% of revenues derive from DW solutions, not from professional services.

- **A standalone EDW solution.** We included EDW products that are not technologically tied to any particular application, such as enterprise resource planning or customer relationship management; are not tied to any particular BI, business performance solution, predictive analytics, extract-transform-load, or middleware stack; and that do not require embedding in other applications. The EDW solution can stand alone.

- **Referenceable customers.** All of the participating EDW vendors — Oracle declined to participate in this Forrester Wave evaluation — provided contact information for at least two customers that agreed to speak to Forrester about their use of the EDW solution.

- **Sparked client inquiries and/or has technologies that put it on Forrester's radar.** Forrester clients often discuss the vendors and products through inquiries; alternatively, the vendor may, in Forrester's judgment, warrant inclusion in this evaluation because of technology trends or market presence.

Figure 2 Evaluated Vendors: Vendor And Product Information And Selection Criteria

Vendor name	Product name and version
Actian	Vectorwise 3.0
Amazon.com	Redshift
HP	Vertica
IBM	• PureData System for Analytics N2001 • PureData System for Operational Analytics • DB2 10.1 for Linux, Unix, and Windows • DB2 10 for z/OS with DB2 Analytics Accelerator • Informix 12 with TimeSeries Acceleration
Kognitio	Kognitio Analytical Platform 8.1
Microsoft	SQL Server 2012 PDW V2 AU1
Oracle	Oracle 11gR2
ParAccel	ParAccel Analytical Platform 4.0.3
Pivotal	Pivotal Greenplum Database
SAP	• SAP NetWeaver BW 7.3 powered by SAP HANA • SAP HANA SP6 • SAP Sybase IQ 16
Teradata	• Teradata Database 14.10 • Teradata workload specific platforms • Teradata Unified Data Architecture • Teradata Aster Database 6

Selection criteria
An enterprise-class data warehouse solution that offers all core data warehouse functional components, tools, and features.
At least $25 million in annual revenues attributable specifically to the enterprise data warehouse.
The solution can stand alone and is not technologically tied to any particular application, performance solution, or stack.
The vendor can provide the names of at least two referenceable customers that can vouch for the quality of its offering.
Significant interest from Forrester clients.

Source: Forrester Research, Inc.

ENTERPRISE MATURITY AND INNOVATION DISTINGUISH THE LEADERS

The evaluation uncovered an EDW market in which (see Figure 3):

- **Teradata, SAP, IBM, Oracle, Pivotal, and Microsoft lead with mature, highly scalable solutions.** The EDW Leaders — Teradata, SAP, IBM, Oracle, Pivotal, Microsoft — have supported large enterprise customers' analytics needs for many years. All offer mature, high-performance, scalable, secure, flexible, and robust EDW solutions that often combine an analytics-optimized DW with query optimization, parallel loading, deep data compression, and mixed-workload management features. Teradata has a full range of EDW appliances to support any DW requirement. SAP Sybase IQ supports massively parallel columnar DW technology; coupled with SAP HANA, it delivers a distributed in-memory analytics platform. IBM has ramped up its analytics offering with its integrated PureData System appliance for analytics and operational analytics that delivers high performance, scale, and automation to support large platforms. Oracle leverages its database ubiquity to dominate the market with the most widely used EDW solution. Pivotal continues to innovate and execute well on its strategy, increasing its adoption and offering integration with its own Hadoop distribution. Microsoft is the second most widely used EDW solution; its cost-effective EDW software and appliances dominate the midmarket.

- **Kognitio, HP, AWS, Actian, and ParAccel are lapping at their heels.** This group has turned up the heat as high as it will go on the incumbent leaders, with innovations that many customers find compelling. Kognitio's entirely in-memory, distributed EDW is appealing for customers looking for fast performance on commodity hardware. HP Vertica has risen quickly into the top tier of EDW platform providers through its strategy, innovation, and commitment. Amazon, a new entrant in the EDW market, has done extremely well, gaining more than 1,000 customers since its launch earlier this year. ParAccel was acquired by Actian during the course of the research for this Forrester Wave evaluation. ParAccel's distributed EDW on commodity hardware make it an innovative choice, but new owner Actian must invest in more enterprise features if it's to rise up and compete with other vendors in the Strong Performers category, much less with the Leaders.

This evaluation of the enterprise data warehouse market is intended to be a starting point only. We encourage clients to view detailed product evaluations and adapt criteria weightings to fit their individual needs through the Forrester Wave Excel-based vendor comparison tool. Clients can also schedule an inquiry with the analysts to discuss specific needs.

Figure 3 Forrester Wave™: Enterprise Data Warehouse, Q4 '13

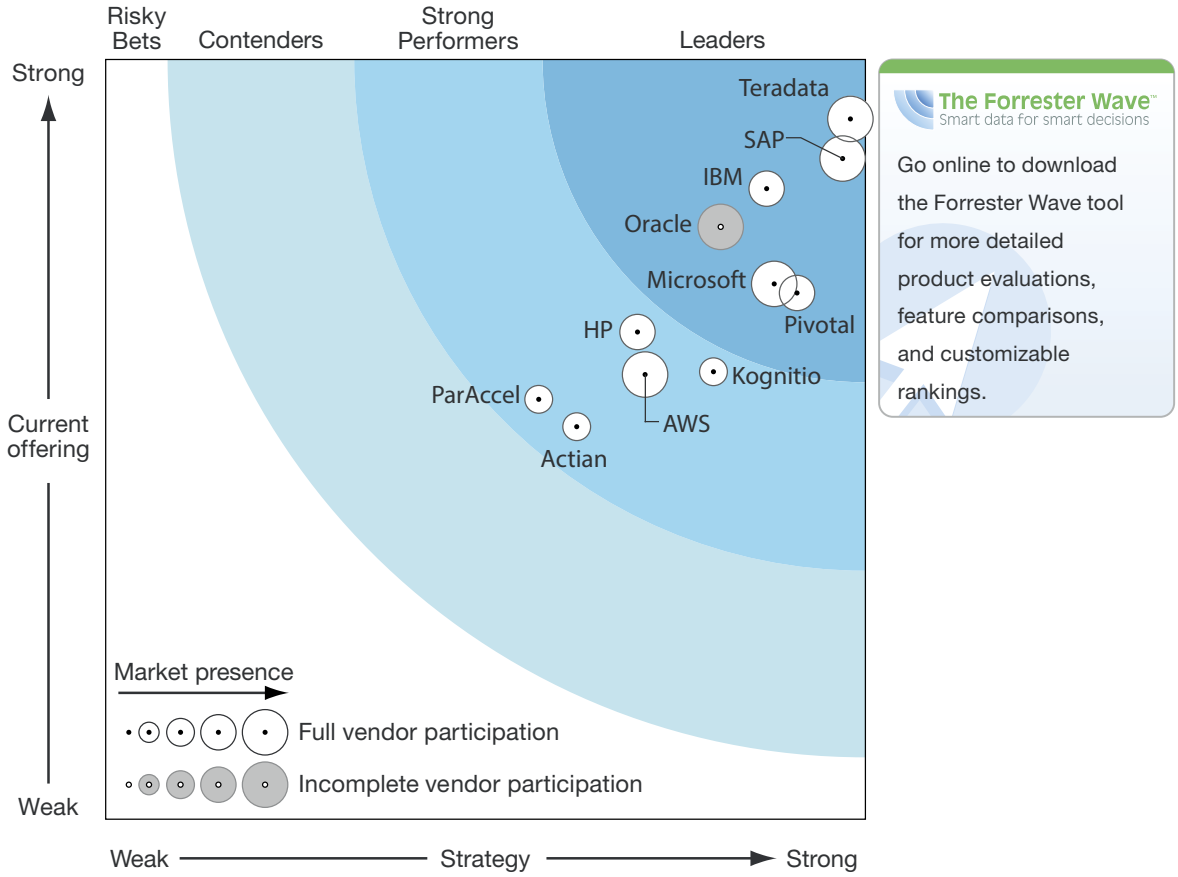

Source: Forrester Research, Inc.

Figure 3 Forrester Wave™: Enterprise Data Warehouse, Q4 '13 (Cont.)

	Forrester's Weighting	Actian	AWS	HP	IBM	Kognitio	Microsoft	ParAccel	Pivotal	SAP	Teradata
CURRENT OFFERING	50%	2.58	2.88	3.16	4.10	2.95	3.52	2.76	3.41	4.35	4.56
Architecture	40%	2.78	2.28	3.10	4.50	2.85	4.00	3.00	3.60	4.75	4.75
Operations	60%	2.46	3.29	3.20	3.83	3.02	3.20	2.60	3.29	4.09	4.43
STRATEGY	50%	3.10	3.55	3.50	4.35	4.00	4.40	2.85	4.55	4.85	4.90
Go-to-market	25%	3.00	3.20	4.20	3.20	3.40	3.40	2.40	4.00	4.40	5.00
Commitment	25%	3.40	5.00	3.80	4.20	2.60	4.20	3.00	4.20	5.00	4.60
Product road map	50%	3.00	3.00	3.00	5.00	5.00	5.00	3.00	5.00	5.00	5.00
MARKET PRESENCE	0%	2.80	4.41	3.57	3.97	2.95	4.60	2.52	3.78	4.86	4.34
Company financials	40%	3.00	4.00	4.00	4.00	3.00	4.00	3.00	4.00	5.00	4.00
Adoption	30%	2.84	4.36	3.08	3.40	2.84	5.00	1.40	2.36	4.52	4.12
Third-party vendors	30%	2.50	5.00	3.50	4.50	3.00	5.00	3.00	4.50	5.00	5.00

All scores are based on a scale of 0 (weak) to 5 (strong).

Source: Forrester Research, Inc.

VENDOR PROFILES

Leaders Have Scalable EDW And Analytics Solution Delivery

- **Teradata's razor-sharp focus is paying off.** Teradata offers the most comprehensive and scalable EDW platform. Teradata has strong revenues, a solid installed base, good momentum, and partnerships in the EDW market. It has several differentiators, including one of the broadest sets of EDW packaging, licensing, pricing, and professional services options on the market, as well as certified integration with a broad range of partner applications and middleware components. It provides sophisticated functionality for in-database analytics, caching, compression, partitioning, indexing, cost-based query optimization, and workload management. It also provides in-database analytics, logical data models, user-defined functions, and stored procedures for EDW extensibility and customization. Teradata offers connectors to import and export data to and from Hadoop and also offers SQL-H and SQL-MapReduce as two options for executing queries against Hadoop. Teradata also offers Aster Big Analytics as an option for the DW to provide advanced analytics capabilities such as path analysis. Overall, Forrester has seen that customers find Teradata suitable for any of their EDW requirements; when they consider alternative solutions, it's usually due to affinity for another vendor's transactional database.

- **SAP is finding its data mojo in SAP HANA and SAP Sybase IQ to support broader EDW use cases.** SAP's differentiating go-to-market message for EDW focuses on its multinode distributed shared-nothing architecture distributed in memory, optimized data streaming, cache-coherent programming for optimized in-memory processing, real-time analytics, columnar data format, advanced compression, and data services layer. SAP HANA has done extremely well, gaining more than 1,800 customers since its launch in 2011. Enterprises are using SAP HANA for in-memory data marts and SAP BW implementations that integrate with other DWs, including SAP Sybase IQ that has more than 4,500 customer installations. SAP provides streaming functionality to queue and ingest incoming streaming data, integrating with SAP Sybase ESP, a complex event processing solution, and supports high-speed data movement through SAP Sybase Replication Server. For enterprises that want to leverage EDW on a cloud platform, both SAP HANA and Sybase IQ are available on Amazon AWS, although most customers are small and medium-size businesses. SAP continues to invest significantly in virtualization, in-memory, real-time, columnar, ad hoc, distributed EDW and predictive analytics.

- **IBM is realigning its EDW go-to-market strategy around scalable appliances.** IBM's large installed base of databases, integrated systems and comprehensive analytics product portfolio, and professional services organization give it a competitive advantage. In late 2012, IBM significantly ramped up its EDW offering by rolling out more scalable and integrated purpose-built PureData appliances for analytics and operational analytics that directly compete against Oracle, SAP, and Teradata. These appliances also integrate with IBM InfoSphere, Optim, Cognos, WebSphere, and other portfolios to support real-time and predictive analytics. From a feature perspective, IBM EDW solutions support deep data compression, in-database analytics, real-time streaming, vertical data models, and automated resource management and have native support for MapReduce and Hadoop APIs. IBM DB2 with BLU Acceleration speeds analytics and reporting using dynamic in-memory columnar architecture, parallel vector processing, and actionable compression. Firms using BLU Acceleration require less data storage and can optimize their data access without needing database tuning, indexes, and aggregates. Overall, IBM has executed its EDW strategy well, but needs to extend its public cloud EDW offering and distributed in-memory analytics to support more use cases.

- **Oracle is leveraging its EDW appliances and transactional installed base to grow its market.** Oracle's dominant position in the database market, growing family of EDW appliance solutions, and deep portfolio of application and middleware solutions gives it a competitive advantage.[3] It's an easy choice for many enterprise customers, because Oracle is their online transaction processing DBMS provider and they can leverage their Oracle ELA licensing contract. Oracle's EDW solution is mature and proven. It offers the ability to deploy EDW on several form factors, including software, an integrated Exadata appliance, cloud/SaaS services on Amazon AWS, and virtualized offerings. Oracle Exadata uses intelligent storage nodes, hybrid columnar compression, parallel in-database analytics, and integrated flash cache to deliver a scale-out EDW platform. In addition, Oracle EDW provides optimized indexing, flexible partitioning,

efficient compression, and mature, cost-based query optimization to support a scalable, high-performance EDW. Oracle has a growing family of Exalytics analytic solution appliances packaged for specific vertical industries and business functions. Oracle continues to enhance its EDW offering by investing in intelligent data caching, storage optimization, in-database analytics, virtualization, data compression, in-memory, and SSD technologies.

- **Pivotal could catapult over other EDW solutions.** In addition to its Greenplum EDW solution, Pivotal also offers its own Hadoop distribution, including a massively parallel processing (MPP) Hadoop SQL engine called Hawq, to offer MPP-like SQL performance on Hadoop. With many traditional EDW customers also looking for Hadoop solutions, Pivotal could offer both solutions and potentially disrupt the market by delivering a lower cost per terabyte than other vendors. Most of the Leaders in this Forrester Wave evaluation have a Hadoop strategy too, but Pivotal was early to take its MPP expertise to its own Hadoop distribution. Pivotal supports streaming functionality through integration with SQLFire and GemFire to queue and asynchronously update the DW. It has large deployments, some of which run from hundreds of terabytes to petabytes.

- **Microsoft EDW adoption is growing and challenging the top EDW players.** Microsoft's significant presence in the database, online analytical processing (OLAP), BI, spreadsheet (PowerPivot), collaboration, and browser markets gives it a strong competitive advantage. Microsoft is already leveraging these positions to deliver a growing EDW solution stack, which in the future will include expanded cloud-based EDW services, into opportunities across diverse market segments. Microsoft offers a range of EDW appliances; for example, Dell Quickstart fulfills the requirements of the midmarket and the HP AppSystem and Dell PDW appliances can handle the needs of larger enterprises. Microsoft SQL Server supports diverse OLAP, BI, analytics, predictive analytics, and real-time analytics patterns. To execute its EDW strategy, Microsoft has made several acquisitions, including DATAllegro, Stratature, and Zoomix, and partnered with key hardware vendors like HP and Dell. Furthermore, Microsoft continues to innovate and enhance its EDW solutions in the areas of caching, distributed caching, in-memory, columnar, indexing, partitioning, cost-based query optimization, data compression, data integration, and real-time analytics. Most customers report that cost and ease of use are the primary reasons for choosing SQL Server for their DW initiative.

Strong Performers Are Hot On Their Heels

- **Kognitio's in-memory solution is innovative; now it must improve its enterprise features.** In-memory solutions are fast becoming normal options for many EDW solutions. Kognitio's EDW is a strong, cost-effective alternative to SAP HANA. But if Kognitio wants to become a Leader, it will have to bolster its data modeling, administration, and management tools. Kognitio is a row-based relational DBMS and was designed from the start as an MPP (distributed) in-memory RDBMS, making extensive use of RAM-based processing for maximum performance.

It supports data compression and can linearly add data capacity by adding nodes or appliances. Kognitio is typically scaled out by adding additional servers to the system, but Kognitio can also fully and automatically scale up by exploiting RAM, processor, or disk upgrades on each server. Kognitio runs on 64-bit Linux distributions and is available as software only, an appliance, or a public or private cloud service.

- **HP Vertica is becoming a substantial EDW player through ongoing platform enhancement.** HP Vertica has shown strong growth and momentum in the EDW market and continues to innovate and enhance its products in terms of scalability, performance, deployment flexibility, and TCO. HP Vertica offers advanced data compression, support for hybrid storage, shared-nothing architecture, indexing, partitioning, distributed query optimization, and workload management. HP has a big data analytics platform that includes Vertica analytics databases, Autonomy's analytics engine, and other technologies. HP Vertica provides connectors to import and export data to and from Hadoop and supports advanced analytics, including predictive analytics. HP Vertica is significantly broadening its platform in terms of data management, analytics, and deployment options in addition to enhancing the core platform.

- **Amazon Web Services is making a big splash in the EDW market.** Amazon Redshift is the fastest-growing service in the history of AWS, with more than 1,000 customers as of May 2013. Today, customers use Amazon Redshift to support small to medium-size DWs and data marts, with larger deployments likely in the coming years. AWS did not build Redshift from scratch, but licensed technology components from ParAccel that were originally built using PostgreSQL database technology. Amazon has significantly enhanced ParAccel's technology to incorporate columnar data storage, tighter integration with Amazon's server and storage services, enhanced data compression, and additional performance, scale, and security features to support multitenancy. Amazon Redshift has a MPP architecture, parallels and distributes queries across multiple nodes, and automates most common DW administrative tasks associated with provisioning, configuring, monitoring, and security. Unlike other vendors that are now offering integrated real-time and predictive analytics, Amazon relies mostly on partnerships to provide a broader range of tools for data integration, analytics, replication, middleware integration, security, and visualization.

- **Actian needs to make a more compelling case.** Actian's Vectorwise EDW solution is a natural choice for existing Ingres databases, as Vectorwise is based on legacy Ingres database technology. However, this also means that Vectorwise's scalability is limited to one node. Actian's purchase of ParAccel was a good move, because Vectorwise did not have a scale-out architecture. Actian must now rationalize Vectorwise and ParAccel technology so that the market is not confused. It must also significantly bolster its enterprise features to be more competitive.

- **ParAccel is innovative, but must find its enterprise sea legs.** ParAccel's biggest claim to fame is that Amazon used its EDW technology as the basis on which to build Amazon Redshift — an

impressive claim indeed. ParAccel was a natural for Amazon, because its MPP architecture is designed to run on commodity hardware. That approach can have benefits for enterprises — but, to win in the enterprise market, ParAccel must significantly enhance the operational tools, such as information life-cycle management, that enterprises expect.

SUPPLEMENTAL MATERIAL

Online Resource

The online version of Figure 3 is an Excel-based vendor comparison tool that provides detailed product evaluations and customizable rankings.

Data Sources Used In This Forrester Wave

Forrester used a combination of two data sources to assess the strengths and weaknesses of each solution:

- **Product demos.** We asked participating vendors to conduct demonstrations of their product's functionality. We used findings from these product demos to validate details of each vendor's product capabilities.
- **Customer reference calls.** To validate product and vendor qualifications, Forrester also conducted reference calls with two of each vendor's current customers.

The Forrester Wave Methodology

We conduct primary research to develop a list of vendors that meet our criteria to be evaluated in this market. From that initial pool of vendors, we then narrow our final list. We choose these vendors based on: 1) product fit; 2) customer success; and 3) Forrester client demand. We eliminate vendors that have limited customer references and products that don't fit the scope of our evaluation.

After examining past research, user need assessments, and vendor and expert interviews, we develop the initial evaluation criteria. To evaluate the vendors and their products against our set of criteria, we gather details of product qualifications through a combination of lab evaluations, questionnaires, demos, and/or discussions with client references. We send evaluations to the vendors for their review, and we adjust the evaluations to provide the most accurate view of vendor offerings and strategies.

We set default weightings to reflect our analysis of the needs of large user companies — and/or other scenarios as outlined in the Forrester Wave document — and then score the vendors based on a clearly defined scale. These default weightings are intended only as a starting point, and we encourage readers to adapt the weightings to fit their individual needs through the Excel-based

tool. The final scores generate the graphical depiction of the market based on current offering, strategy, and market presence. Forrester intends to update vendor evaluations regularly as product capabilities and vendor strategies evolve. For more information on the methodology that every Forrester Wave evaluation follows, go to http://www.forrester.com/marketing/policies/forrester-wave-methodology.html.

Integrity Policy

All of Forrester's research, including Forrester Wave evaluations, is conducted according to our integrity policy. For more information, go to http://www.forrester.com/marketing/policies/integrity-policy.html.

ENDNOTES

[1] Data virtualization is the integration of internal or external data from disparate sources in real time or near real time. See the August 8, 2013, "Information Fabric 3.0" report.

[2] In-database analytics is where the EDW and advanced analytics connect. With in-database analytics, enterprises migrate their predictive analysis, data mining, and other compute-intensive analytics functions from separate, standalone applications to execute in the EDW. See the November 12, 2009, "In-Database Analytics: The Heart Of The Predictive Enterprise" report.

[3] Forrester's February 2013 Global Database Management Online Survey indicated that Oracle is the most widely used data warehouse solution with 62% adoption, followed by SQL Server at 54%. See the June 7, 2013, "The Steadily Growing Database Market Is Increasing Enterprises' Choices" report.

FORRESTER®

About Forrester

A global research and advisory firm, Forrester inspires leaders, informs better decisions, and helps the world's top companies turn the complexity of change into business advantage. Our research-based insight and objective advice enable IT professionals to lead more successfully within IT and extend their impact beyond the traditional IT organization. Tailored to your individual role, our resources allow you to focus on important business issues — margin, speed, growth — first, technology second.

FOR MORE INFORMATION

To find out how Forrester Research can help you be successful every day, please contact the office nearest you, or visit us at www.forrester.com. For a complete list of worldwide locations, visit www.forrester.com/about.

CLIENT SUPPORT

For information on hard-copy or electronic reprints, please contact Client Support at +1 866.367.7378, +1 617.613.5730, or clientsupport@forrester.com. We offer quantity discounts and special pricing for academic and nonprofit institutions.

Forrester Focuses On Application Development & Delivery Professionals

Responsible for leading the development and delivery of applications that support your company's business strategies, you also choose technology and architecture while managing people, skills, practices, and organization to maximize value. Forrester's subject-matter expertise and deep understanding of your role will help you create forward-thinking strategies; weigh opportunity against risk; justify decisions; and optimize your individual, team, and corporate performance.

« **ANDREA DAVIES,** client persona representing Application Development & Delivery Professionals