

Data Warehouse Platforms

A SearchDataManagement.com Directory

2009 EDITION

This product brought to you by:



SearchDataManagement.com

Welcome!

[INTRO](#)

[SELECTING THE
RIGHT DATA
WAREHOUSE
PLATFORM
FOR YOUR
ORGANIZATION](#)

[INDEX](#)

[1010DATA](#)

[ASTER DATA
SYSTEMS](#)

[IBM](#)

[ILLUMINATE
SOLUTIONS](#)

[GREENPLUM](#)

[INFORMATICA](#)

[KALIDO](#)

[KOGNITIO](#)

[MICROSOFT CORP.](#)

[NETEZZA](#)

[ORACLE
CORPORATION](#)

[PARACCEL](#)

[PERVASIVE](#)

[SAND TECHNOLOGY](#)

[SUN MICROSYSTEMS](#)

[SYBASE](#)

[TERADATA](#)

[VERTICA](#)

[ABOUT
SEARCHDATA-
MANAGEMENT
.COM](#)

[METHODOLOGY](#)

Welcome to SearchDataManagement.com's Data Warehouse Platforms Directory. This directory is designed to be a valuable resource for those getting started with research or evaluating optimized data warehouses. Inside, you'll find basic information about vendors and the platforms they sell. Each listing is accompanied by a short description and a long description, including limited information about functionality and product use. You'll find products for businesses of all sizes as well as products that can be deployed on-demand and on-premise. Use this list to get started with the evaluation process. For more information about any of the products or to speak to a sales representative, please visit the vendor website or product website.

[SearchDataManagement.com](#) will launch a series of directories throughout the year to address unique segments of the data management market. Want to see your product listed in one of our directories? Go here to [submit a product](#). Need to update product or pricing information? [Email us here](#). For questions for the editors or to make suggestions for improving the directory, write to us at editor@search-datamanagement.com.

Happy shopping!

Selecting the right data warehouse platform for your organization

[INTRO](#)

[SELECTING THE RIGHT DATA WAREHOUSE PLATFORM FOR YOUR ORGANIZATION](#)

[INDEX](#)

[1010DATA](#)

[ASTER DATA SYSTEMS](#)

[IBM](#)

[ILLUMINATE SOLUTIONS](#)

[GREENPLUM](#)

[INFORMATICA](#)

[KALIDO](#)

[KOGNITIO](#)

[MICROSOFT CORP.](#)

[NETEZZA](#)

[ORACLE CORPORATION](#)

[PARACCEL](#)

[PERVASIVE](#)

[SAND TECHNOLOGY](#)

[SUN MICROSYSTEMS](#)

[SYBASE](#)

[TERADATA](#)

[VERTICA](#)

[ABOUT SEARCHDATA-MANAGEMENT.COM](#)

[METHODOLOGY](#)



As companies consider their ability to manage the asset that is information, the data warehouse platform (and its database management system, DBMS) is the most important decision point. The platform is the foundational component of the tool selections, the consultancy hires, the architecture, etc. In short, it defines your information culture.

However, in selecting the platform to support the data warehouse, organizations are faced with an exponentially higher number of variations and distinct departures from the traditional online transactional processing (OLTP) database management systems than ever before.

Over time, data warehouse data volumes will continue to soar as history accumulates, syndicated data is collected and new sources with more detailed data are added. Furthermore, the community consuming the data continues to grow, expanding well beyond company boundaries to customers, supply chain partners and even the Internet. Companies need to make sure they choose a proven platform not just for the initial, known requirements but also with the ability to scale to future, to-be-determined requirements.

Data warehouses, to be successful, need to provide:

- Quality data that is available.
- Architectures that provide low long-term total cost of ownership.
- Good query performance that results in increased interactive usage.
- The ability to get to real-time feeds.
- A platform to support mixed, unpredictable workloads.
- A scalable path forward as data needs grow.

CRITERIA FOR DATA

WAREHOUSE PLATFORM SELECTION

The decision process should go well beyond the usual consideration of the operational DBMS vendor. Today, that decision can be nuanced along several potential requirements, including:

- Active loading of data and immediate access to loaded data.
- A mixed workload upon the data.
- The immediate availability of information.
- Cross-functional complexity.
- The level of query concurrency.
- The scalability needs of the platform.
- The functionality of the DBMS.

Given the state of the marketplace, the technical architecture for a data warehouse platform should be:

- [INTRO](#)
- [SELECTING THE RIGHT DATA WAREHOUSE PLATFORM FOR YOUR ORGANIZATION](#)
- [INDEX](#)
- [1010DATA](#)
- [ASTER DATA SYSTEMS](#)
- [IBM](#)
- [ILLUMINATE SOLUTIONS](#)
- [GREENPLUM](#)
- [INFORMATICA](#)
- [KALIDO](#)
- [KOGNITIO](#)
- [MICROSOFT CORP.](#)
- [NETEZZA](#)
- [ORACLE CORPORATION](#)
- [PARACCEL](#)
- [PERVASIVE](#)
- [SAND TECHNOLOGY](#)
- [SUN MICROSYSTEMS](#)
- [SYBASE](#)
- [TERADATA](#)
- [VERTICA](#)
- [ABOUT SEARCHDATA-MANAGEMENT.COM](#)
- [METHODOLOGY](#)

Scalable: The data warehouse platform should be scalable in both performance capacity and incremental data volume growth. Make sure the proposed platform scales in a near-linear fashion and behaves consistently with growth in all of database size, number of concurrent users, and complexity of queries. Understand the additional hardware and software required for each of the incremental uses.

Powerful: The data warehouse platform should be designed for complex decision support activity in a multi-user mixed workload environment. Check on the maturity of the optimizer for supporting every type of query with good performance and to determine the best execution plan based on changing data demographics. Check on conditional parallelism and what the causes are of variations in the parallelism deployed. Check on dynamic and controllable prioritization of resources for queries.

Manageable: The platform should be manageable with minimal support tasks requiring DBA/System Administrator intervention. It should provide a single point of control to simplify system administration. You should be able to create and implement new tables and indexes at will.

Extensible: The data warehouse platform should provide flexible database design and system architecture that keeps pace with evolving business requirements and leverages existing investment in hardware and applications. Ask such questions as: What is required to add and delete columns? What is the impact of repartitioning tables?

Available: The platform should support mission-critical business applications with minimal downtime. Check on “hot pluggable” components, understand system downtime requirements and any issues that might deny or degrade service to end users. These can include batch load times, software/ hardware upgrades, severe system performance issues, and system maintenance outages.

Interoperable: The platform should allow for integrated access to the Web, internal networks and corporate mainframes.

Affordable: The proposed technology (hardware, software, services, required customer support) should allow for low total cost of ownership (TCO) over a multi-year period.

Flexible: The platform should be able to provide optimal performance across the full range of normalized, star and hybrid data schemas with large numbers of tables. Look for a proven ability to support multiple applications from different business units, leveraging data that is integrated across business functions and subject areas.

Robust in-database management systems, features and functions: The platform should include DBA productivity tools, monitoring features, parallel utilities, robust query optimizer, locking schemes, security methodology, intra-query parallel implementation for all possible access paths, chargeback and accounting features, and remote maintenance capabilities.

Referenceable: There may not be one

- [INTRO](#)
- [SELECTING THE RIGHT DATA WAREHOUSE PLATFORM FOR YOUR ORGANIZATION](#)
- [INDEX](#)
- [1010DATA](#)
- [ASTER DATA SYSTEMS](#)
- [IBM](#)
- [ILLUMINATE SOLUTIONS](#)
- [GREENPLUM](#)
- [INFORMATICA](#)
- [KALIDO](#)
- [KOGNITIO](#)
- [MICROSOFT CORP.](#)
- [NETEZZA](#)
- [ORACLE CORPORATION](#)
- [PARACCEL](#)
- [PERVASIVE](#)
- [SAND TECHNOLOGY](#)
- [SUN MICROSYSTEMS](#)
- [SYBASE](#)
- [TERADATA](#)
- [VERTICA](#)
- [ABOUT SEARCHDATA-MANAGEMENT.COM](#)
- [METHODOLOGY](#)

single reference that matches your environment exactly, but you should be able to see a consistent trend across a wide range of references that enforces what you are looking for in a data warehouse platform.

Organizations should also consider vendor viability, especially in these days of vendor consolidation. The vendor's financial stability, the importance of data warehousing to their overall business strategy and their continued research and development in these areas toward a well-developed and relevant vision are all key components of a vendor's viability in this critical decision.

DATA WAREHOUSE ARCHITECTURAL CONSIDERATIONS

To enable the above criteria, companies have some new options in the types of data warehouses available, many of which have been created in recent years owing to the high cost of ownership of the traditional approaches.

Today, typical data management environments have several restrictions they've "learned to live with" in seeking acceptable analytical performance. Indexes, summary tables, cubes, utility executions like reorganizations and various de-normalizations are often introduced into the environment. There's a point at which overcoming these challenges ceases to be "easy" and affordable and therefore ceases to get done. But perhaps more important than these very real restrictions are the restrictions in possibilities thinking and information exploitation. The higher the untuned performance of the environment is, the better.

ORIENTATION

A row-oriented database is an implementation of relational theory. Data is stored as bytes, with all columns for a row stored in order. These bytes are grouped by the several thousand—from 4,000 to 64,000—into data blocks. This is the unit of input/output, an exception being the "pre-fetch" capa-

“Organizations should also consider vendor viability, especially in these days of vendor consolidation.”

—WILLIAM MCKNIGHT

bility of row-oriented database management systems (DBMS) sensing a pattern in the reads. A columnar DBMS is an implementation of the relational theory, but with a twist. The data storage layer contains a grouping of columns. For example, all the column 1 values are physically together, followed by all the column 2 values, etc. The data is stored in record order, so the 100th entry for column 1 and the 100th entry for column 2 belong to the same input record.

In a columnar database, all the same data—your data—is there. It's just organized differently (automatically, by the DBMS). The main benefits include the ability to highly compress the data. A byproduct of the compression is the ability to perform columnar operations—like MIN, MAX, SUM, COUNT and AVG—very rapidly.

- [INTRO](#)
- [SELECTING THE RIGHT DATA WAREHOUSE PLATFORM FOR YOUR ORGANIZATION](#)
- [INDEX](#)
- [1010DATA](#)
- [ASTER DATA SYSTEMS](#)
- [IBM](#)
- [ILLUMINATE SOLUTIONS](#)
- [GREENPLUM](#)
- [INFORMATICA](#)
- [KALIDO](#)
- [KOGNITIO](#)
- [MICROSOFT CORP.](#)
- [NETEZZA](#)
- [ORACLE CORPORATION](#)
- [PARACCEL](#)
- [PERVASIVE](#)
- [SAND TECHNOLOGY](#)
- [SUN MICROSYSTEMS](#)
- [SYBASE](#)
- [TERADATA](#)
- [VERTICA](#)
- [ABOUT SEARCHDATA-MANAGEMENT.COM](#)
- [METHODOLOGY](#)

DATA WAREHOUSE APPLIANCES MODEL

Data warehouse appliances have emerged as viable short-list solutions for new or refurbishing data management efforts. Vendors now dot the landscape (with representation in analyst "quadrants" and "spectrums" and hundreds of millions in venture backing) and they are worthy of any data warehouse professional's understanding, attention and consideration.

The data warehouse appliance is a hardware/software/OS/DBMS/storage preconfiguration for data management requirements. Many utilize commodity components and some use open source DBMS for a low overall TCO. These open source DBMSs provide a starting point for basic database functionality, and the appliance vendors focus on necessary functionality enhancements. Performance of queries, especially against large volumes of data, is distinctly impressive thanks to the automatic parallelism many, but not all, appliances provide. Low TCO for a mixed-workload data warehouse environment is consequential with appliances as well.

PARALLELISM ARCHITECTURE

Parallelism is as important as columnar orientation to query performance success, and massively parallel processing (MPP) systems provide the most parallelism possible. MPP systems evolved from symmetric multiprocessing (SMP) systems, which, in turn, evolved from uniprocessor systems.

Sharing disk and/or memory across nodes creates overhead. Sharing nothing minimizes disk access bottlenecks and is preferred in MPP, which also fea-

tures an exponential uptake from SMP and clusters in the number of processors and the sophistication of the interconnect. MPP is the generally acknowledged ideal parallel architecture for analytic query, business intelligence and data warehousing.

"A true 'data warehouse in a box' has remained somewhat elusive..."

—WILLIAM MCKNIGHT

PACKAGED ANALYTICS

Finally, today, most large enterprise software vendors offer some type of pre-built reports or dashboards centered on typical business functions such as Finance, Sales and Marketing, and Supply Chain Management. In addition, a number of smaller companies have developed packaged analytics built on licensed platforms and targeted to specific industries, applications or specialized functions, or to supplement other third-party services.

While a true "data warehouse in a box" has remained somewhat elusive, analytic packages bridge the gap between off-the-shelf canned reports and a ground-up custom BI software. Such packages simplify the development process and deliver functionality with more cost certainty. Organizations must complete a careful evaluation process to determine whether a package is right for them, and they must understand the various costs and benefits of different options. With all of this

- [INTRO](#)
- [SELECTING THE RIGHT DATA WAREHOUSE PLATFORM FOR YOUR ORGANIZATION](#)
- [INDEX](#)
- [1010DATA](#)
- [ASTER DATA SYSTEMS](#)
- [IBM](#)
- [ILLUMINATE SOLUTIONS](#)
- [GREENPLUM](#)
- [INFORMATICA](#)
- [KALIDO](#)
- [KOGNITIO](#)
- [MICROSOFT CORP.](#)
- [NETEZZA](#)
- [ORACLE CORPORATION](#)
- [PARACCEL](#)
- [PERVASIVE](#)
- [SAND TECHNOLOGY](#)
- [SUN MICROSYSTEMS](#)
- [SYBASE](#)
- [TERADATA](#)
- [VERTICA](#)
- [ABOUT SEARCHDATA-MANAGEMENT.COM](#)
- [METHODOLOGY](#)

information at hand, decision makers will be better positioned to make the best choice for their organization.

DATA WAREHOUSE SELECTION PROCESS

Once the possibilities have been narrowed down according to the evaluation of the above, a proof of concept should be done with your data. Data warehouse vendors all have the ability to take in your data for an offsite proof of concept. These offsite data centers are well worth using. You can “visit your data” in person in the process, if you wish.

You need to determine, and gather, the representative data as well as determine the representative data access paths that you wish to explore in a proof of concept (POC). The POC should comprise these phases:

1. Prepare and gather data.
2. Set up environment(s) at off-site location(s).
3. Validate data access tool connectivity, performance, and load capa-

bilities, with emphasis on concurrency.

4. Compile qualitative information on each offering, using published information, consulting, and references.
5. Weigh quantitative and qualitative recommendations on each vendor.

Make sure each important criterion is covered in the POC.

CONCLUSION

These days, the best answer to your data warehouse platform selection needs may not be as evident as it was in the past. Given the high relevance of data and information, investment and innovation continue in this area and there are a surprisingly large number of options. Many platforms may work, but at high short- and long-term prices. Making a selection consistent with your goals as an organization and with a broad and detailed view of the landscape is paramount to success. ■

ABOUT THE AUTHOR



William McKnight, President, McKnight Consulting Group

William functions as strategist, lead enterprise information architect and program manager for complex, high-volume, full lifecycle implementations worldwide—utilizing the disciplines of data warehousing, master data management (MDM), business intelligence (BI), data quality and operational BI. McKnight is a Southwest Entrepreneur of the Year finalist and a frequent best-practices judge; he has authored more than 150 articles and white papers and given more than 150 international keynotes and public seminars. His team’s implementations from both IT and consultant positions have won Best Practices awards. He is a former IT VP of a Fortune company, a former engineer of DB2 at IBM, and holds an MBA. He can be reached at william@williammcknight.com.

Index at a Glance

Click on the product name at left to jump to a longer description.

	PAGE	VENDOR	PRODUCTS	SAAS OR SERVICES	ON PREMISE SW
INTRO					
SELECTING THE RIGHT DATA WAREHOUSE PLATFORM FOR YOUR ORGANIZATION					
INDEX					
1010DATA	10	1010data	1010data	●	◆
ASTER DATA SYSTEMS	10	Aster Data Systems	Aster nCluster	●	◆
IBM	11	IBM	InfoSphere Warehouse		◆
ILLUMINATE SOLUTIONS	11	illuminate Solutions	iLuminate	●	◆
GREENPLUM	12	Greenplum	Greenplum Database 3.3	●	◆
INFORMATICA	12	Informatica	PowerCenter		◆
KALIDO	13	Kalido	Kalido Dynamic Information Warehouse (DIW)		◆
KOGNITIO	13	Kognitio	WX2 analytical database	●	◆
MICROSOFT CORP.	14	Microsoft Corp.	SQL Server 2008		◆
NETEZZA	14	Netezza	AppNexus		◆
ORACLE CORPORATION	15	Netezza	Data Warehouse Appliances		◆
PARACCEL	15	Oracle Corporation	HP Oracle Database Machine		◆
PERVASIVE	16	ParAccel	The ParAccel Analytic Database (PADB)	●	◆
SAND TECHNOLOGY	16	Pervasive	Pervasive Data Warehouse		◆
SUN MICROSYSTEMS	17	SAND Technology	SAND/DNA Access	●	◆
SYBASE	17	Sun Microsystems	MySQL	●	◆
TERADATA	18	Sybase	Sybase IQ 15		◆
VERTICA	18	Teradata	The Teradata Purpose-Built Platform Family		◆
ABOUT SEARCHDATA-MANAGEMENT.COM	19	Vertica	Vertica Analytical Database 3.0	●	◆
METHODOLOGY	19	Vertica	Vertica Analytic Database for the Cloud	●	
	20	Vertica	Vertica Virtualized Analytical Database	●	

Vendor: Vendor/developer of product at directory press time; ● **SaaS or services** indicates technology available as SaaS, hosted, on-demand, ASP and Web services; ◆ **On-premise** indicates software or systems on premise; ⓘ Descriptions were written by the SearchDataManagement.com editorial team based on information gathered from vendor websites.

1010DATA

1010data

1010data's data warehousing service provides a high-performance alternative to DW and conventional OLAP technologies, empowering organizations with a fast, reliable, cost-effective and easy-to-use capability to get the most out of data. ● ◆

COMPANY WEBSITE: www.1010data.com

FOUNDED: 2000

SUMMARY: 1010data's service provides a unique platform that combines front-end usability with back-end functionality. Delivered as a managed solution, 1010data allows analysts to quickly construct complex, sophisticated queries on very large datasets and get results in seconds, allowing organizations to make quick, critical decisions. 1010data's service model provides a low-risk, fully managed and supported database solution that does not require additional personnel, has a short time to delivery and is relatively inexpensive. Since it is so quick and easy to integrate, enterprises experience high speed-to-value and very rapid ROI.

PRICING: 1010data's business intelligence service starts at \$50,000; price is based on amount of data and the number and type of users to ensure proper service levels.

ASTER DATA SYSTEMS

Aster nCluster

Aster nCluster is built as a high-performance analytic SQL database for frontline data warehousing and analytics. ● ◆

COMPANY WEBSITE: www.asterdata.com

FOUNDED: 2005

SUMMARY: Aster nCluster is a high-performance analytic SQL database for frontline data warehousing. It's designed to provide an always-on, always-parallel MPP architecture with the first-ever In-Database MapReduce programming framework. "Always parallel" means that Aster can deliver maximum system performance and efficiency through network-optimized MPP architecture. "Always on" means that Aster can provide hands-free 24/7 availability for data-driven applications. And In-Database MapReduce is designed to provide in-database analytics done right—the analytic expressiveness and flexibility of MapReduce combined with the rich functionality and support of SQL.

PRICING: Aster prices are based on amount of user data. Pricing starts at \$100,000. Depending on the amount of data being analyzed, pricing can go from low six-figures to the seven-figure range.

INTRO
SELECTING THE
RIGHT DATA
WAREHOUSE
PLATFORM
FOR YOUR
ORGANIZATION

INDEX

1010DATA

ASTER DATA
SYSTEMS

IBM

ILLUMINATE
SOLUTIONS

GREENPLUM

INFORMATICA

KALIDO

KOGNITIO

MICROSOFT CORP.

NETEZZA

ORACLE
CORPORATION

PARACCEL

PERVASIVE

SAND TECHNOLOGY

SUN MICROSYSTEMS

SYBASE

TERADATA

VERTICA

ABOUT
SEARCHDATA-
MANAGEMENT
.COM

METHODOLOGY

InfoSphere Warehouse

InfoSphere Warehouse is designed to be a comprehensive data warehouse tool, providing all of the capabilities necessary to glean maximum return from your information. ♦

COMPANY WEBSITE: www.ibm.com

FOUNDED: 1896 (as the Tabulating Machine Company)

SUMMARY: IBM InfoSphere Warehouse is designed to help organizations take the next step in innovating with information. It provides a powerful tool set delivering access to structured and unstructured information and delivering operational insights in real-time. InfoSphere is designed to reduce costs by freeing storage through a broader use of compression and by making workload management a standard feature of Departmental and Enterprise editions. It's also designed to optimize the use of industry-standard XML-formatted data by directly storing standard XML data into the data warehouse, compressing XML data to save on storage costs and improving performance when managing XML records. ☉

PRICING: Vendor declined to provide pricing.

iLuminate

iLuminate is a correlation database engine that is designed to contain all the data management tools needed to build and manage an enterprise data warehouse or data mart. ● ♦

COMPANY WEBSITE: www.illuminateinc.com

FOUNDED: 2005

SUMMARY: The iLuminate correlation database engine uses value-based storage (VBS), which stores data values only once and uses a sophisticated indexing system to logically connect all data correlations and relationships. The engine contains the tools needed to build and manage an enterprise data warehouse or data mart with automated database design and no need for users to define requirements up front. The database—actually the indexing system itself—is created on the fly during the loading process and is inherently compact and fast. During the loading process, the data itself creates its own data-generated schema and all the data values are 100% indexed. This provides fast train-of-thought analysis.

PRICING: Price is based on the number of records loaded; starts at \$41,900 for up to 35 million records. Concurrent connection fees apply and start at \$2,750 for up to four connections.

[INTRO](#)[SELECTING THE RIGHT DATA WAREHOUSE PLATFORM FOR YOUR ORGANIZATION](#)[INDEX](#)[1010DATA](#)[ASTER DATA SYSTEMS](#)[IBM](#)[ILLUMINATE SOLUTIONS](#)[GREENPLUM](#)[INFORMATICA](#)[KALIDO](#)[KOGNITIO](#)[MICROSOFT CORP.](#)[NETEZZA](#)[ORACLE CORPORATION](#)[PARACCEL](#)[PERVASIVE](#)[SAND TECHNOLOGY](#)[SUN MICROSYSTEMS](#)[SYBASE](#)[TERADATA](#)[VERTICA](#)[ABOUT SEARCHDATA-MANAGEMENT.COM](#)[METHODOLOGY](#)

GREENPLUM

Greenplum Database 3.3

Greenplum Database is software built to support the next generation of data warehousing and large-scale analytics processing. ● ◆

COMPANY WEBSITE: www.greenplum.com/
FOUNDED: 2003

SUMMARY: Greenplum Database utilizes a shared-nothing massively parallel processing architecture optimized for business intelligence and analytical processing. It allows a cluster of servers to operate as a single database supercomputer—automatically partitioning data and parallelizing queries—to achieve very fast performance times. Greenplum Database moves processing power as close as possible to the data, so processing always occurs in parallel, delivering unmatched query and load performance. With Greenplum Database, users have the power to answer complex questions, running analyses that used to take days in just seconds. ○

PRICING: Vendor declined to provide pricing.

INFORMATICA

PowerCenter

Informatica PowerCenter is a single, unified enterprise data integration platform for accessing, discovering and integrating data from virtually any business system, in any format, and delivering that data throughout the enterprise at any speed. ◆

COMPANY WEBSITE: www.informatica.com
FOUNDED: 1993

SUMMARY: Highly available, high-performing and fully scalable, PowerCenter serves as the foundation for all data integration projects and enterprise integration initiatives across the enterprise, including B2B data exchange, data governance, data migration, data synchronization and replication, enterprise data warehousing, integration competency centers (ICCs), master data management (MDM), and service-oriented architectures (SOAs). PowerCenter is designed to deliver trusted, timely data throughout the enterprise—in batch, near real-time, and real-time modes, or on demand—to meet the analytical and operational needs of the business. Comprehensive auditing and reporting capabilities enable business analysts to support governance, risk and compliance initiatives. ○

PRICING: Vendor declined to provide pricing.

INTRO
SELECTING THE
RIGHT DATA
WAREHOUSE
PLATFORM
FOR YOUR
ORGANIZATION

INDEX

1010DATA

ASTER DATA
SYSTEMS

IBM

ILLUMINATE
SOLUTIONS

GREENPLUM

INFORMATICA

KALIDO

KOGNITIO

MICROSOFT CORP.

NETEZZA

ORACLE
CORPORATION

PARACCEL

PERVASIVE

SAND TECHNOLOGY

SUN MICROSYSTEMS

SYBASE

TERADATA

VERTICA

ABOUT
SEARCHDATA-
MANAGEMENT
.COM

METHODOLOGY

KALIDO

Kalido Dynamic Information Warehouse (DIW)

Kalido's Dynamic Information Warehouse is designed to automate the creation and lifecycle management of enterprise data warehouses based on your business model. ♦

COMPANY WEBSITE: www.kalido.com/
FOUNDED: 1997

SUMMARY: Kalido Dynamic Information Warehouse integrates data from multiple sources and powers operational and strategic reporting. The software keeps track of history so users can do as-of and before-and-after time-based comparisons. DIW is built to handle business change, adapting to deliver consistent information even during mergers and acquisitions, reorganizations or market consolidation. Its business model-driven approach helps business and IT gather user requirements and prototype, iterate and deploy the data warehouse quickly. DIW can save organizations significant deployment time and resources compared with custom-built approaches, and it reduces the manpower required to support operational implementations. ◉

PRICING: Vendor declined to provide pricing. DIW is sold only as part of the Kalido Information Engine.

KOGNITIO

WX₂ analytical database

WX₂ is designed to be one of the fastest, most mature analytical databases on the market, producing answers to queries in just seconds, and offering a wide range of deployment options (software, appliance or DaaS). ● ♦

COMPANY WEBSITE: www.kognitio.com
FOUNDED: 2005

SUMMARY: Kognitio WX₂ is designed to help mid-sized and large enterprises gain unprecedented analytical access to, and insight from, massive amounts of data, thus allowing them to remain ahead of their competitors. Available via traditional software license, as an integrated data warehouse appliance, or through Kognitio's DaaS, WX₂ enables users to reliably gain greater benefits in less time by querying, in detail, vast amounts of granular data in just seconds. WX₂ can obtain results several times faster than competing data warehouse appliances and up to 80 times faster than typical software-only databases.

PRICING: Kognitio WX₂ offers a flexible pricing model dependent on the deployment model, starting from as low as \$30,000 per TB (software-only pricing based on a 40 TB system).

[INTRO](#)

[SELECTING THE RIGHT DATA WAREHOUSE PLATFORM FOR YOUR ORGANIZATION](#)

[INDEX](#)

[1010DATA](#)

[ASTER DATA SYSTEMS](#)

[IBM](#)

[ILLUMINATE SOLUTIONS](#)

[GREENPLUM](#)

[INFORMATICA](#)

[KALIDO](#)

[KOGNITIO](#)

[MICROSOFT CORP.](#)

[NETEZZA](#)

[ORACLE CORPORATION](#)

[PARACCEL](#)

[PERVASIVE](#)

[SAND TECHNOLOGY](#)

[SUN MICROSYSTEMS](#)

[SYBASE](#)

[TERADATA](#)

[VERTICA](#)

[ABOUT SEARCHDATA-MANAGEMENT.COM](#)

[METHODOLOGY](#)

MICROSOFT CORP.

SQL Server 2008

Microsoft SQL Server 2008 is a scalable and reliable data warehouse platform for business intelligence. Through integration with the Microsoft BI platform, it is designed to enable actionable business decisions at low cost. ♦

COMPANY WEBSITE: www.microsoft.com

FOUNDED: 1975

SUMMARY: Microsoft SQL Server 2008 is a scalable and reliable data warehouse platform for business intelligence. Through integration with the Microsoft BI platform, it enables actionable business decisions at the lowest cost. Take advantage of a comprehensive Data Warehouse and BI platform in SQL Server. SQL Server 2008 ships with a world-class relational database, analysis services, reporting services and integration services. Top new features include: data compression, back-up compression, partitioned table parallelism, star join query optimizations, resource management, change data capture, MERGE SQL statements, and scalable integration services.

PRICING: Enterprise Edition: a comprehensive data management and business intelligence platform providing enterprise-class scalability, high availability, and security for running business-critical applications; Channels: Retail, VL, OEM, ISVR, SPLA; Processor Pricing: Retail \$24,999, Example, \$23,911 Server Plus CAL Pricing: Retail \$13,969 with 25 CALs; Example, \$8,487, \$162 per additional CAL

NETEZZA

AppNexus

AppNexus is built to allow users to perform intensive processes, analytics and data warehousing faster and more easily with data warehouse in the cloud. ♦

COMPANY WEBSITE: www.netezza.com

FOUNDED: 2000

SUMMARY: AppNexus has partnered with Netezza to offer its customers a data warehouse cloud service. As a stand-alone system, Netezza appliances are designed to offer a fast and easy implementation process, and customers realize a fast time to value with Netezza implementations. The AppNexus “Netezza in the Cloud” service extends the simplicity of Netezza implementations by offering rapid provisioning, scalability on demand, and pay-as-you-go pricing based on utilization. This gives customers a compelling new and better option for their immediate, and often elastic, terabyte-scale data analysis needs. ◉

PRICING: Vendor declined to provide pricing.

INTRO
SELECTING THE
RIGHT DATA
WAREHOUSE
PLATFORM
FOR YOUR
ORGANIZATION

INDEX

1010DATA

ASTER DATA
SYSTEMS

IBM

ILLUMINATE
SOLUTIONS

GREENPLUM

INFORMATICA

KALIDO

KOGNITIO

MICROSOFT CORP.

NETEZZA

ORACLE
CORPORATION

PARACCEL

PERVASIVE

SAND TECHNOLOGY

SUN MICROSYSTEMS

SYBASE

TERADATA

VERTICA

ABOUT
SEARCHDATA-
MANAGEMENT
.COM

METHODOLOGY

NETEZZA

Data Warehouse Appliances

To eliminate the need for constant tuning and the technology bottlenecks caused by slow disk transfer I/O rates, slow network transfer rates and inefficient caching, Netezza developed a system designed specifically for analytical processing on extremely large amounts of data. ♦

COMPANY WEBSITE: www.netezza.com/

FOUNDED: 2000

SUMMARY: Netezza's massively parallel architecture takes a different approach to processing queries than traditional database architecture. Where standard computing architectures target operations on individual data elements, Netezza architecture is all about "streaming" processing. Rather than shuttling data between disk and memory for processing once a query comes in, data streams off the disk and through query logic loaded into a field programmable gate array (FPGA). The FPGA and processor, together with 400 GB of disk storage, reside on each of the massively parallel nodes called snippet processing units (SPUs). Queries are optimized across the SPUs for maximum performance and power efficiency. ◉

PRICING: Vendor declined to provide pricing.

ORACLE CORPORATION

HP Oracle Database Machine

Oracle serves database and data warehouse markets with products including Oracle Database 11g and the HP Oracle Database Machine. ♦

COMPANY WEBSITE: www.oracle.com

FOUNDED: 1978

SUMMARY: The HP Oracle Database Machine is designed as a complete system, including software, servers, networking and storage designed to run multi-terabyte data warehouses 10 times faster than conventional data warehouse systems. At the heart of this system is the HP Oracle Exadata Storage Server, which has smart storage software that offloads data-intensive query processing from database servers closer to the data. The Database Machine runs Oracle Database 11g and Real Application Clusters on Oracle Enterprise Linux, and also includes complete Infiniband infrastructure. It's designed to be simple and fast to implement and cost-effective. The Database Machine can linearly scale storage capacity, processing power and network bandwidth as your data warehouse grows. ◉

PRICING: Pricing can be found at <http://www.oracle.com/corporate/pricing/exadata-pricelist.pdf>
Vendor declined to provide additional information.

INTRO

SELECTING THE
RIGHT DATA
WAREHOUSE
PLATFORM
FOR YOUR
ORGANIZATION

INDEX

1010DATA

ASTER DATA
SYSTEMS

IBM

ILLUMINATE
SOLUTIONS

GREENPLUM

INFORMATICA

KALIDO

KOGNITIO

MICROSOFT CORP.

NETEZZA

ORACLE
CORPORATION

PARACCEL

PERVASIVE

SAND TECHNOLOGY

SUN MICROSYSTEMS

SYBASE

TERADATA

VERTICA

ABOUT
SEARCHDATA-
MANAGEMENT
.COM

METHODOLOGY

PARACCEL

The ParAccel Analytic Database (PADB)

PADB is a purpose-built data warehousing and analytic database management system that boasts load-and-go simplicity. PADB's innovative, MPP-columnar design makes it the fastest data warehousing platform on the planet, period. ●◆

COMPANY WEBSITE: www.paraccel.com

FOUNDED: 2005

SUMMARY: The ParAccel Analytic Database is a robust, relational and fully-transactional DBMS that brings transformational speed to data warehousing, analytic and business intelligence (BI) applications. Its impressive speed is built-in so that high-maintenance tuning structures like indexes, materialized views and summary tables are not necessary. This versatile performance is not dependent on a star schema design—our schema-neutral approach lets you also unleash the power of ParAccel for complex and ad hoc queries even on normalized and de-normalized schemas. Performance features include columnar data storage, massively parallel processing, patent-pending SAN integration, a patent-pending analytic query optimizer and scalable data loading. ☉

PRICING: Vendor declined to provide pricing.

PERVASIVE

Pervasive Data Warehouse

Pervasive Data Warehouse is based on Pervasive Data Integrator and is designed to quickly scope, plan and deliver an analytical data warehouse to meet your business intelligence objectives. ◆

COMPANY WEBSITE: www.pervasiveintegration.com

FOUNDED: 1982

SUMMARY: The Pervasive Data Warehouse delivers thorough documentation and a detailed design model for your data warehouse. An optional requirements-gathering phase can be used to define and design the integration processes necessary for the ETL portion of the project. The design can then be implemented by your developers or by a Pervasive consultant. Based on your objectives, data warehouse consultants will be able to design the analytical model for the data warehouse and relevant mart, identify the necessary source data elements required to meet business objectives, and determine the necessary SQL scripts to create the analytical model. Pervasive Data Warehouse requires Pervasive Data Integrator (see [company website](#) for details). ☉

PRICING: Vendor declined to provide pricing.

[INTRO](#)

[SELECTING THE RIGHT DATA WAREHOUSE PLATFORM FOR YOUR ORGANIZATION](#)

[INDEX](#)

[1010DATA](#)

[ASTER DATA SYSTEMS](#)

[IBM](#)

[ILLUMINATE SOLUTIONS](#)

[GREENPLUM](#)

[INFORMATICA](#)

[KALIDO](#)

[KOGNITIO](#)

[MICROSOFT CORP.](#)

[NETEZZA](#)

[ORACLE CORPORATION](#)

[PARACCEL](#)

[PERVASIVE](#)

[SAND TECHNOLOGY](#)

[SUN MICROSYSTEMS](#)

[SYBASE](#)

[TERADATA](#)

[VERTICA](#)

[ABOUT SEARCHDATA-MANAGEMENT.COM](#)

[METHODOLOGY](#)

SAND TECHNOLOGY

SAND/DNA Access

SAND/DNA Access provides data management software and best practices for customers with large amounts of data they want to access and store. ● ◆

COMPANY WEBSITE: www.sand.com

FOUNDED: 1983

SUMMARY: The SAND/DNA product suite is designed to complement the data lifecycle management capabilities available to organizations as their data warehouse tools grow in size. Users can move infrequently used data into a repository where it's stored in a compressed format while remaining easily accessible. Customers can quickly respond to changing business needs, regulatory compliance requests and changing market data without impacting their users and online warehouse. As a result, infrastructure costs are reduced and operational performance is improved, all without compromising users' ability to access granular historical data. SAND/DNA works with Oracle, Microsoft, DB2, SAS, SAP, Teradata and Netezza.

PRICING: DSAND/DNA Access pricing is based on the size of the flat file data stored in SAND's repository, and it is a perpetual license.

SUN MICROSYSTEMS

MySQL

MySQL database is an open source database designed for fast performance, high reliability, ease of use, and dramatic cost savings. ● ◆

COMPANY WEBSITE: www.sun.com

FOUNDED: 1982

SUMMARY: MySQL database is the world's most popular open source database because of its fast performance, high reliability and dramatic cost savings. It has also become the database of choice for a new generation of applications built on the LAMP stack (Linux, Apache, MySQL, PHP/Perl/Python). MySQL runs on more than 20 platforms, including Linux, Windows, OS X, Solaris OS, and Netware, providing flexibility. MySQL Enterprise powers a range of applications—from business-critical transactional systems to online eCommerce, data warehousing, reporting and analytics systems. Its "shared nothing" architecture doesn't require additional infrastructure investment and provides 99.999% data availability with no point of failure.

PRICING: There is no license fee for MySQL or for MySQL Enterprise. However, MySQL Enterprise offers various levels of service and support, beginning at \$599/year for Basic support. The various levels and their pricing are available on the main [MySQL website](http://www.mysql.com).

[INTRO](#)

[SELECTING THE
RIGHT DATA
WAREHOUSE
PLATFORM
FOR YOUR
ORGANIZATION](#)

[INDEX](#)

[1010DATA](#)

[ASTER DATA
SYSTEMS](#)

[IBM](#)

[ILLUMINATE
SOLUTIONS](#)

[GREENPLUM](#)

[INFORMATICA](#)

[KALIDO](#)

[KOGNITIO](#)

[MICROSOFT CORP.](#)

[NETEZZA](#)

[ORACLE
CORPORATION](#)

[PARACCEL](#)

[PERVASIVE](#)

[SAND TECHNOLOGY](#)

[SUN MICROSYSTEMS](#)

[SYBASE](#)

[TERADATA](#)

[VERTICA](#)

[ABOUT
SEARCHDATA-
MANAGEMENT
.COM](#)

[METHODOLOGY](#)

SYBASE

Sybase IQ 15

Sybase IQ 15 is a highly optimized analytics server designed to deliver faster results for mission-critical business intelligence, data warehouse and reporting tools on any standard hardware and operating systems. ♦

COMPANY WEBSITE: www.sybase.com

FOUNDED: 1984

SUMMARY: Sybase IQ 15 is a column-based analytics server that is designed to enable enterprises to turn raw data into actionable information through analytics without the constant tuning required by some row-based databases or the approach to business intelligence typically taken by hardware data warehouse appliances. Sybase IQ 15 works with diverse data—including unstructured data—and diverse data sources and is designed to deliver unsurpassed query performance for the most complex analytics tasks within an affordable, eco-friendly and manageable environment. ⦿

PRICING: Sybase IQ is available in three editions: Sybase IQ Small Business Edition starting at \$120,000, which includes four CPUs; Sybase IQ Enterprise Edition, at \$64,800 per CPU; Sybase IQ Single Application Server Edition, starting at \$2,595 per CPU

TERADATA

The Teradata Purpose-Built Platform Family

The Teradata purpose-built platform is designed for speed, scalability and affordability and is built to handle a high number of complex queries. ♦

COMPANY WEBSITE: www.teradata.com

FOUNDED: 1979

SUMMARY: Teradata's database platforms are designed for scalability and affordability and are built for analyzing data and processing the increasing volumes and growing complexity of queries at the highest performance levels. The platform family offers customers options that take full advantage of all the power of Teradata anywhere in the enterprise—as an active data warehouse, enterprise data warehouse, entry-level data warehouse appliance, special-purpose data mart, or sandbox environment. Teradata platforms are recognized for speed, the amount of data that can be queried and the number and complexity of queries that can be run.

PRICING: The Data Mart Appliance 551 starts at \$56,000/terabyte; the Extreme Data Appliance 1550 starts at \$16,500/terabyte; the Data Warehouse Appliance 2550 starts at \$99,000/terabyte; and the Active Enterprise Data Warehouse 5555 starts at \$150,000/terabyte. <http://www.teradata.com/t/WorkArea/DownloadAsset.aspx?id=4682>

[INTRO](#)

[SELECTING THE RIGHT DATA WAREHOUSE PLATFORM FOR YOUR ORGANIZATION](#)

[INDEX](#)

[1010DATA](#)

[ASTER DATA SYSTEMS](#)

[IBM](#)

[ILLUMINATE SOLUTIONS](#)

[GREENPLUM](#)

[INFORMATICA](#)

[KALIDO](#)

[KOGNITIO](#)

[MICROSOFT CORP.](#)

[NETEZZA](#)

[ORACLE CORPORATION](#)

[PARACCEL](#)

[PERVASIVE](#)

[SAND TECHNOLOGY](#)

[SUN MICROSYSTEMS](#)

[SYBASE](#)

[TERADATA](#)

[VERTICA](#)

[ABOUT SEARCHDATA-MANAGEMENT.COM](#)

[METHODOLOGY](#)

VERTICA

Vertica Analytical Database

As businesses become more analytical to gain competitive advantage and comply with new regulations, data warehouses are pushed to answer more ad hoc questions from more people analyzing vastly larger volumes of data, often in real time. The Vertica Analytic Database is built specifically to meet this need. ● ◆

COMPANY WEBSITE: www.vertica.com

FOUNDED: N/A

SUMMARY: Vertica has these features and functions in the latest edition: Column orientation—query 50x-200x faster by eliminating costly disk IO; “scale-out” MPP architecture—scale limitlessly just by adding new servers to the grid; aggressive data compression—reduce storage costs by up to 90%; automatic high availability—run nonstop with automatic replication, failover and recovery; and deployment flexibility—deploy on Linux or VMware or in the Amazon Cloud to handle a wide variety of projects. ☉

PRICING: Vendor declined to provide pricing

VERTICA

Vertica Analytic Database for the Cloud

Vertica Analytic Database for the Cloud is an on-demand version of Vertica’s grid-enabled columnar database hosted on Amazon’s Elastic Compute Cloud. The pay-as-you-go offering enables companies to create large, high-performance analytic data marts without up-front data center costs and delays. ●

COMPANY WEBSITE: www.vertica.com

FOUNDED: N/A

SUMMARY: Vertica Analytic Database for the Cloud’s latest version features innovations designed to allow Vertica to manage terabytes of data faster and more reliably: “scale-out” grid architecture that handles changing workloads as elastically as the cloud; aggressive data compression to keep storage costs low; and an automatic K-Safety that provides replication, failover and recovery in the cloud. Vertica for the Cloud completely changes the economics of BI, making it possible to rapidly initiate a much broader spectrum of analytic projects and businesses, including ad hoc and short-lived business analytic projects and new analytic SaaS businesses. ☉

PRICING: Vendor declined to provide pricing.

[INTRO](#)

[SELECTING THE RIGHT DATA WAREHOUSE PLATFORM FOR YOUR ORGANIZATION](#)

[INDEX](#)

[1010DATA](#)

[ASTER DATA SYSTEMS](#)

[IBM](#)

[ILLUMINATE SOLUTIONS](#)

[GREENPLUM](#)

[INFORMATICA](#)

[KALIDO](#)

[KOGNITIO](#)

[MICROSOFT CORP.](#)

[NETEZZA](#)

[ORACLE CORPORATION](#)

[PARACCEL](#)

[PERVASIVE](#)

[SAND TECHNOLOGY](#)

[SUN MICROSYSTEMS](#)

[SYBASE](#)

[TERADATA](#)

[VERTICA](#)

[ABOUT SEARCHDATA-MANAGEMENT.COM](#)

[METHODOLOGY](#)

VERTICA

Vertica Virtualized Analytical Database

The Vertica Virtualized Analytic Database is the first high-performance analytic database to run as a virtual machine in private enterprise compute clouds. ●

COMPANY WEBSITE: www.vertica.com

FOUNDED: N/A

SUMMARY: The Vertica Virtualized edition combines a full-featured version of the Vertica Analytic Database with VMware's Hypervisor virtualization platform and an optimized self-contained Linux operating system. It is built as a simple-to-deploy, self-contained software package that runs on any VMware-supported hardware. Using VMware deployment tools, users can install any number of Vertica "virtual appliances" and then provide any one of the instances with a list of the IP addresses of the others, and Vertica will automatically unite them into a cluster. A virtualized Vertica cluster can use direct-attached storage (DAS) or network-attached storage (NAS). ○

PRICING: Vendor declined to provide pricing.

[INTRO](#)

[SELECTING THE RIGHT DATA WAREHOUSE PLATFORM FOR YOUR ORGANIZATION](#)

[INDEX](#)

[1010DATA](#)

[ASTER DATA SYSTEMS](#)

[IBM](#)

[ILLUMINATE SOLUTIONS](#)

[GREENPLUM](#)

[INFORMATICA](#)

[KALIDO](#)

[KOGNITIO](#)

[MICROSOFT CORP.](#)

[NETEZZA](#)

[ORACLE CORPORATION](#)

[PARACCEL](#)

[PERVASIVE](#)

[SAND TECHNOLOGY](#)

[SUN MICROSYSTEMS](#)

[SYBASE](#)

[TERADATA](#)

[VERTICA](#)

[ABOUT SEARCHDATA-MANAGEMENT.COM](#)

[METHODOLOGY](#)

SearchDataManagement.com is a guide for information management and business intelligence professionals. With its combination of original articles, news, learning guides, expert advice, white papers, webcasts and customized research, SearchDataManagement.com offers a rich collection of insights on how to efficiently manage the data supply chain. The site also offers tips on vendors and product selection, as well as expert advice.

Visit SearchDataManagement.com for:

- ▶ **Independent content:** Award-winning, vendor-independent news and analysis.
- ▶ **Expert advice:** Our Ask the Experts section features advice from some of the leading authorities in the data management domain.
- ▶ **Learning guides:** Search our comprehensive library for useful how-to guides on any data management topic.
- ▶ **Vendor-produced content:** We hand-select white papers and webcasts to address the most relevant data management trends, issues and solutions.

**BECOME
A MEMBER
TODAY!**

Guide methodology: TechTarget has not evaluated the products listed and/or described in this Directory and does not assume any liability arising out of the purchase or use of any product described herein, neither does it convey any license or rights in or to any of the evaluated or listed products. TechTarget has prepared this Directory from sources deemed reliable (including vendors, research reports and certain publicly available information). TechTarget has used good-faith efforts to indicate when content has been provided by a vendor and, in some cases, has removed what was deemed to be overt marketing language.

TechTarget is not responsible for any errors or omissions in this Directory or for interpretations thereof, and expressly disclaims all warranties as to the accuracy, completeness or adequacy of all content contained herein. This disclaimer of warranty is in lieu of all warranties whether expressed, implied or statutory, including implied warranties of merchantability or fitness for a particular purpose. Information in this Directory is current as of Q3 2009. For more recent information, please check the vendors' websites. The opinions expressed herein are subject to change without notice.

©2009 TechTarget, Inc. All Rights Reserved. Reproduction and distribution of this publication in any form without prior written permission is forbidden. TechTarget and the TechTarget logo are registered trademarks of TechTarget, Inc.; all other trademarks are the property of their respective companies.

To compile this guide, our editorial team initially consulted research reports by major analyst firms covering the business intelligence market and contacted vendors about products reviewed by those firms. Editors also conducted additional Internet research and solicited feedback from our expert contacts. A notice about the project was posted on SearchDataManagement.com and listed regularly in our email newsletters.

Vendors were invited to submit listings via a form on the website. For vendors that did not submit listings, our editorial team compiled listings by excerpting information from the vendor's website. All entries, whether they were vendor-submitted or compiled by our team, were edited for length and clarity and to remove overt marketing language. In order to best assist our readers in assessing products, our editorial team attempted to obtain basic pricing information for all products in this directory—requesting information from vendors multiple times via email. Vendors that did not respond, or refused to provide any pricing information, have this statement on their listings: "Vendor declined to provide pricing."

Collection of data for this directory took place during Q2 and Q3 2009. As with any directory of this kind, products and vendors may change substantially at any time. Though every effort was made to make this directory as complete and accurate as possible, there may be changes, errors, omissions or vendors in this market not included in this guide. Nothing in this guide should be construed as endorsements, professional suggestions or advice. This directory should be used simply as a resource. We strongly urge you to supplement this with your own research and to contact vendors for the most up-to-date information about their companies or products. It is our intent to update this directory annually, but that is subject to change.