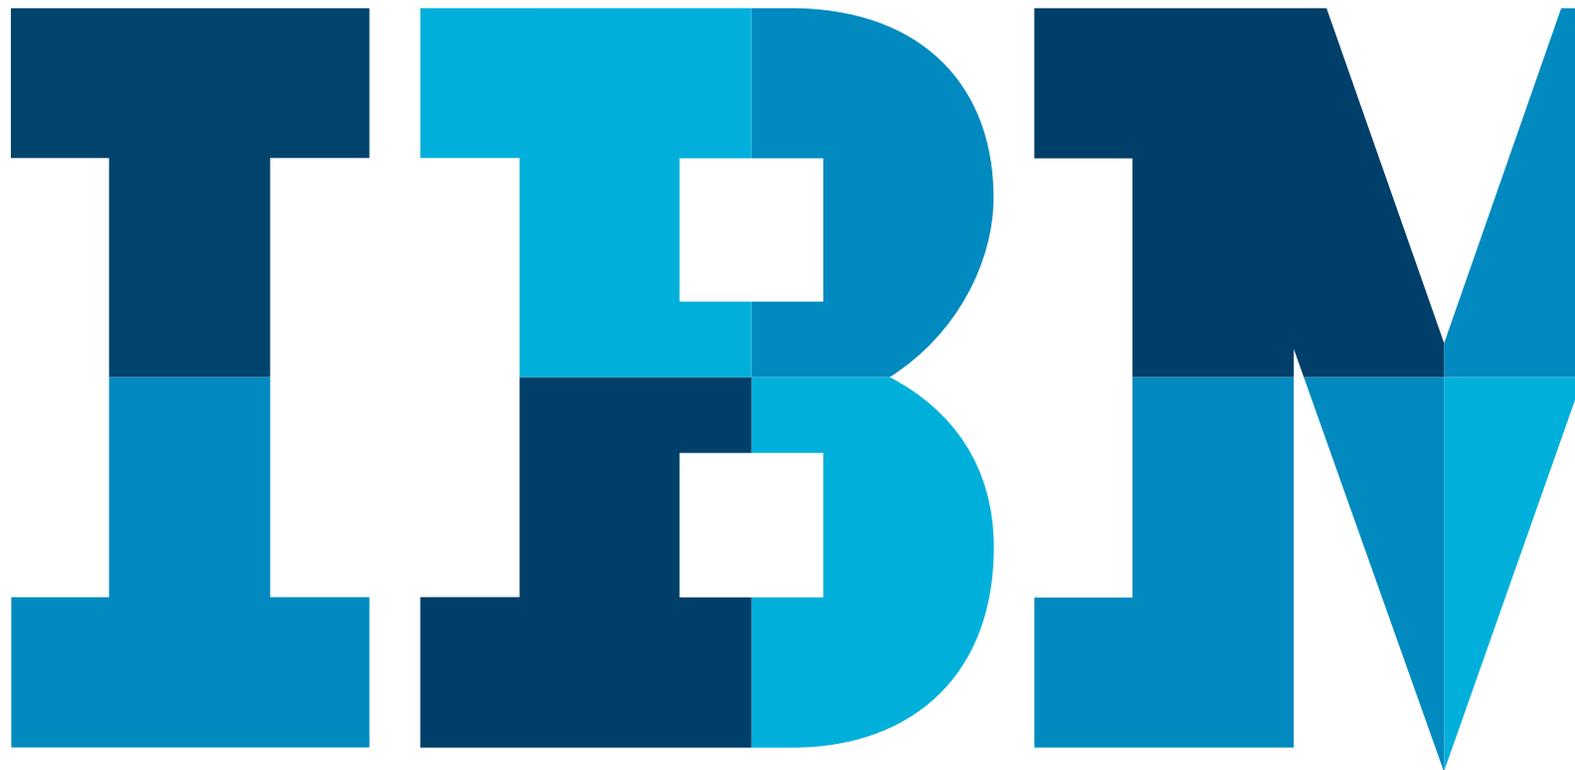


Oracle® Exadata® and IBM® Netezza® Data Warehouse Appliance compared

*by Phil Francisco, Vice President, Product Management and Product Marketing, IBM
and Mike Kearney, Senior Director, Product Marketing, IBM*



Contents

Introduction	2
Online Transaction Processing (OLTP) and data warehousing	3
Query performance	5
Simplicity of operation	10
Value	14
Conclusion	16

Introduction

IBM Netezza data warehouse appliances focus on technology designed to query and analyze big data. IBM Netezza data warehouse appliances are disrupting the market. Wishing to exploit data at lower costs of operation and ownership, many of our customers have moved their data warehouses from Oracle. Oracle has now brought Exadata to market, a machine which apparently does everything an IBM Netezza data warehouse appliance does, and also processes online transactions. This examination of Exadata and IBM Netezza as data warehouse platforms is written from an unashamedly IBM viewpoint, however to ensure credibility we have taken advice from **Philip Howard, Research Director of Bloor Research and Curt Monash, President, Monash Research.**

To innovate requires us to think and do things differently, solving a problem using new approaches. IBM Netezza data warehouse appliances deliver excellent **performance** for our customers' warehouse queries. IBM Netezza data warehouse appliances offer customers **simplicity**; anyone with basic knowledge of SQL and Linux has the skills needed to perform the few administrative tasks required to maintain consistent service levels through dynamically changing workloads. IBM Netezza data warehouse appliances' performance with simplicity reduces their costs of owning and running their data warehouses.

“Netezza was part of the inspiration for Exadata. Teradata was part of the inspiration for Exadata. We'd like to thank them for forcing our hand and forcing us to go into the hardware business.”

— Larry Ellison, January 2010

More important, our customers create new business value by deploying **analytic applications** which they previously considered beyond their reach.

“Netezza was part of the inspiration for Exadata. Teradata was part of the inspiration for Exadata,” acknowledged Larry Ellison on January 27, 2010. “We'd like to thank them for forcing our hand and forcing us go into the hardware business.”¹ While delivered with Larry Ellison's customary pizzazz, there is a serious point to his comment: only the best

¹ See http://oracle.com.edgesuite.net/ivt/4000/8104/9238/12652/lobby_external_flash_clean_480x360/default.htm

catch Oracle's attention. Exadata represents a strategic direction for Oracle; adapting their OLTP database management system, partnering it with a massively parallel storage system from Sun. Oracle launched Exadata V2 with the promise of extreme performance for processing both online transactions and analytic queries. Therefore, Oracle Exadata V2 is a general purpose platform for managing mixed workloads. Oracle Database was designed for OLTP. But data warehousing and analytics make very different demands of their software and hardware than OLTP. Quite simply, some workloads for data warehousing perform much better and are more cost effective on a system that is purpose-built for analytics. Exadata's data warehousing credentials demand scrutiny, particularly with respect to simplicity and value.

This eBook opens by reviewing differences between processing online transactions and processing queries and analyses in a data warehouse. It then discusses Exadata and the IBM Netezza data warehouse appliance from perspectives of their query performance, simplicity of operation and value.

All we ask of readers is that they do as our

customers and partners have done: put aside notions of how a database management system should work, be open to new ways of thinking and be prepared to do less, not more, to achieve a better result.

One caveat: The IBM Netezza data warehouse appliance team has no direct access to an Exadata machine. We are fortunate in the detailed feedback we receive from many organizations that have evaluated both technologies and selected IBM Netezza data warehouse appliances. Given Oracle's size and their focus on Exadata, publicly available information on Exadata is surprisingly scarce. The use cases quoted by Oracle provide little input to the discussion, which in itself is of concern to several industry followers, e.g., Information Week.² The information shared in this paper is made available in the spirit of openness. Any inaccuracies result from our mistakes, not an intent to mislead.

² See http://www.informationweek.com/news/business_intelligence/warehouses/showArticle.jhtml?articleID=225702836&cid=RSSfeed_IWK_News

Online Transaction Processing (OLTP) and data warehousing

OLTP systems execute many short transactions. Each transaction's scope is small, limited to one or a small number of records and is so predictable that often times data is cached. Although OLTP systems process large volumes of database queries, their focus is writing (UPDATE, INSERT and DELETE) to a current data set. These systems are typically specific to a business process or function, for example managing the current balance of a checking account. Their data is commonly structured in third normal form (3NF). Transaction types of OLTP systems are stable and their data requirements are well-understood, so secondary data structures such as indices can usefully locate records on disk, prior to their transfer to memory for processing.

In comparison, data warehouse systems are characterized by predominantly heavy database read (SELECT) operations against a current and historical data set. Whereas an OLTP operation accesses a small number of records, a data warehouse query might scan a table of billions of rows and join its records with those from multiple other tables. Furthermore, queries in a data warehouse are often so unpredictable in nature, it is difficult to exploit caching and indexing strategies. Choices for structuring data in the warehouse range from 3NF to dimensional models such as star and

snowflake schemas. Data within each system feeding a typical warehouse is structured to reflect the needs of a specific business process. Before data is loaded to the warehouse it is cleansed, de-duplicated and integrated.

This eBook divides data warehouses into either first- or second-generation. While this classification may not stand the deepest scrutiny, it reflects how many of our customers talk about their evolutionary path to generating greater and greater value from their data.

First-generation data warehouses are typically loaded overnight. They provide information to their business via a stable body of slowly evolving SQL-based reports and dashboards. As these simple warehouses somewhat resemble OLTP systems – their workload and data requirements are understood and stable – organizations often adopt the same database management products they use for OLTP. With the product comes the practice: database administrators analyze each report's data requirements and build indices to accelerate data retrieval. Creep of OLTP's technology and techniques appears a success, until data volumes in the warehouse outstrip those commonly managed in transactional systems.

In this century, corporations and public sector agencies accept growth rates for data of 30-50 percent per year as normal. Technologies and practices successful in the world of OLTP prove less and less applicable to data warehousing; the index as aid to data retrieval is a case in point. As the database system processes jobs to load data, it is also busy updating its multiple indices. With large data volumes this becomes a very slow process, causing load jobs to overrun their allotted processing window. Despite working long hours, the technical team misses service levels negotiated with the business. Productivity suffers as business units wait for reports and data to become available.

Technologies and practices successful in the world of OLTP prove less and less applicable to data warehousing...

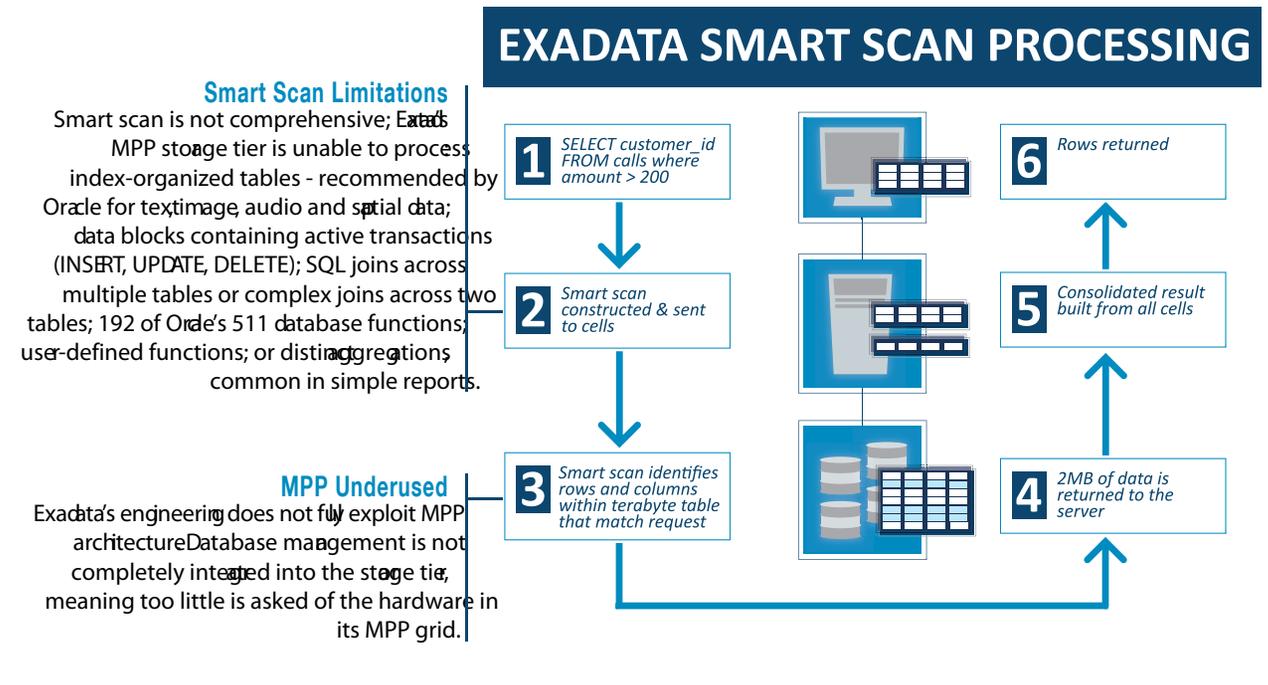
Organizations are redefining how they need and want to exploit their data; this eBook refers to this development as the second-generation data warehouse. These new warehouses, managing massive data sets with ease, serve as the corporate memory. When interrogated, they recall events recorded years previously; these distant memories increase the accuracy of predictive analytic applications. Constant trickle feeds are replacing overnight batch loads, reducing latency between the recording of an event and its analysis. Beyond the simple SQL used to populate reports and dashboards, the warehouse processes linear regressions, Naive Bayes and other mathematical algorithms of advanced analytics. Noticing a sudden spike in sales of a high-margin product at just five stores drives a retailer to understand what happened and why. This knowledge informs strategies to promote similar sales activity at all 150 store locations. The computing system underpinning the warehouse must be capable of managing these sudden surges in demand without disrupting regular reports and dashboards. The business users are demanding the freedom to exploit their data at the time and in the manner of their choosing. Their appetite for immediacy leaves no place for technologies whose performance depends on the tuning work of administrators.

Query performance

Query performance with Oracle Exadata

In acquiring Sun, Oracle has come to the conclusion the IBM Netezza data warehouse appliance team reached a decade earlier: data warehouse systems achieve highest efficiency when all parts, software and hardware, are optimized to their goal. Exadata is created from two sub-systems connected by a fast network: a smart storage system communicating via InfiniBand® with an Oracle Database 11g V2 with Real Application Clusters (RAC). A single rack system includes a storage tier of 14 storage servers, called Exadata cells, in a massively parallel processing (MPP) grid, paired with the Oracle RAC database running as a shared disk cluster of eight symmetric multi-processing nodes.

In acquiring Sun, Oracle has come to the conclusion the IBM Netezza data warehouse appliance team reached a decade earlier: data warehouse systems achieve highest efficiency when all parts, software and hardware, are optimized to their goal.



Oracle labels Exadata's storage tier as smart because it processes SQL projection, restriction and join filtering,³ before putting the resulting data set on the network for downstream processing by Oracle RAC. This technique is called smart scan. However, smart scan is not comprehensive; the storage tier does not process all restrictions. Oracle's online forum⁴ lists a number of operations including scans of index-organized tables or clustered tables as not benefitting from smart scan. Further to these, Christian Antognini, author of the book *Troubleshooting Oracle Performance*, writes a blog that suggests smart scan is not used with the `TIMESTAMP` datatype.⁵ Oracle recommends implementing fact tables in data warehouses as index-organized tables for efficient execution of star queries.⁶ Exadata's storage tier will not process restrictions on index-organized tables, but instead must pass all of the records downstream to the Oracle database. Exadata's approach of passing full records from storage to database tier is highly effective for OLTP as each transaction must only retrieve a small number of rows. However, a statistical analysis requiring a scan of a long (hundreds of millions or billions of rows), wide (hundreds of columns) fact table will generate a tidal wave of data to be inefficiently moved across the network. Exadata would achieve better performance and be more efficient if it processed all SQL predicates (`WHERE` clauses) in its MPP storage tier.

Exadata storage servers cannot communicate with one another; instead all communication is forced via the InfiniBand network to Oracle RAC and then back across the network to the storage tier. This architecture is beneficial to online transaction processing; where each transaction, with a scope of one or few records, can be satisfied by moving a small data set from storage to the database. Analytical queries, such as "find all shopping baskets sold last month in Washington State, Oregon and California containing product X with product Y and with a total value more than \$35," must retrieve much larger data sets, all of which must be moved from storage to database. This inefficient movement of big data adversely effects query performance.

Exadata's storage tier demonstrates other shortcomings. Exadata cells cannot process distinct aggregations, which are common even in simple reports; they are unable to process complex joins or analytical functions used in analytical applications. Unable to resolve these typical data warehousing queries in its storage tier, Exadata must push very large data sets across its internal network to Oracle RAC. This architectural flaw raises questions of Exadata's suitability for second-generation data warehouses which must run complex analytical queries.

Oracle positions its use of 40 Gb/sec switch InfiniBand as an advantage over IBM Netezza data warehouse appliance; in reality, Exadata needs this expensive network because of the system's imbalance and inefficiency. Exadata storage servers do too little work, so more data than necessary is put on the network to be moved downstream for processing by Oracle RAC, which is asked to do too much work.

At its database tier Exadata runs Oracle 11g V2 with Real Application Clusters as a clustered, shared disk architecture. Using this architecture for a data warehouse platform raises concern that contention for the shared resource imposes limits on the amount of data the database can process and the number of queries it can run concurrently. Time and customer experience will tell if this concern is justified.

³ A Technical Overview of the Sun Oracle Exadata Storage Server and Database Machine - An Oracle white paper, October 2009.

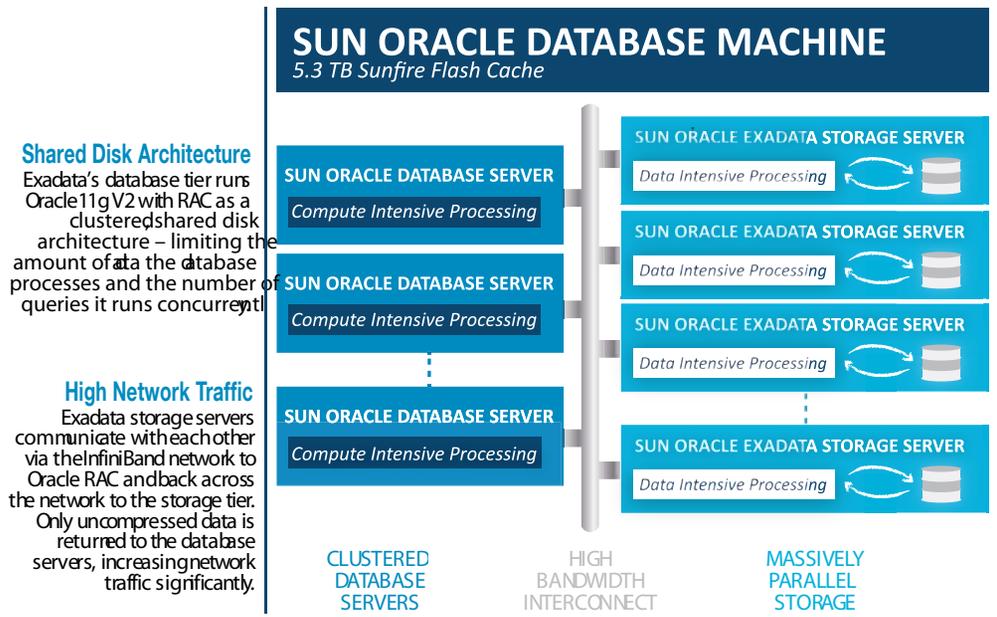
⁴ <http://forums.oracle.com/forums/thread.jspa?threadID=1036774&start=0>
The full list is: scans of index-organized tables or clustered tables; index range scans; access to a compressed index; access to a reverse key index; Secure Enterprise Search

⁵ See Christian Antognini's blog at <http://antognini.ch/2010/05/exadata-storage-server-and-the-query-optimizer-%E2%80%93-part-2/>

⁶ http://www.oracle.com/technology/products/oracle9i/datasheets/iots/iot_ds.html

Every disk in Exadata's storage tier is shared by all nodes in the grid running Oracle RAC. This communal storage creates the risk of a page being read by one node while it is being updated by another. To manage this, Oracle forces coordination between nodes. Each node checks the disk activity of its peers to prevent conflict. Oracle technicians refer to this activity as block pingging. Compute cycles consumed as each node checks disk activity of its peers, or that are lost as one node idly waits for another to complete an operation, are wasted. In an architecture specifically designed for data warehousing, these cycles would be employed processing queries, mining data and running analyses.

For all but simple queries Exadata must move large sets of data from its storage tier to its database tier, raising questions on its suitability as a platform for a modern data warehouse.



Shared Disk Architecture
Exadata's database tier runs Oracle 11g V2 with RAC as a clustered shared disk architecture - limiting the amount of data the database processes and the number of queries it runs concurrently.

High Network Traffic
Exadata storage servers communicate with each other via the InfiniBand network to Oracle RAC and back across the network to the storage tier. Only uncompressed data is returned to the database servers, increasing network traffic significantly.

Performance Bottleneck
Exadata's storage tier does not process restrictions on index-organized tables. All such records are loaded in the database server for processing.

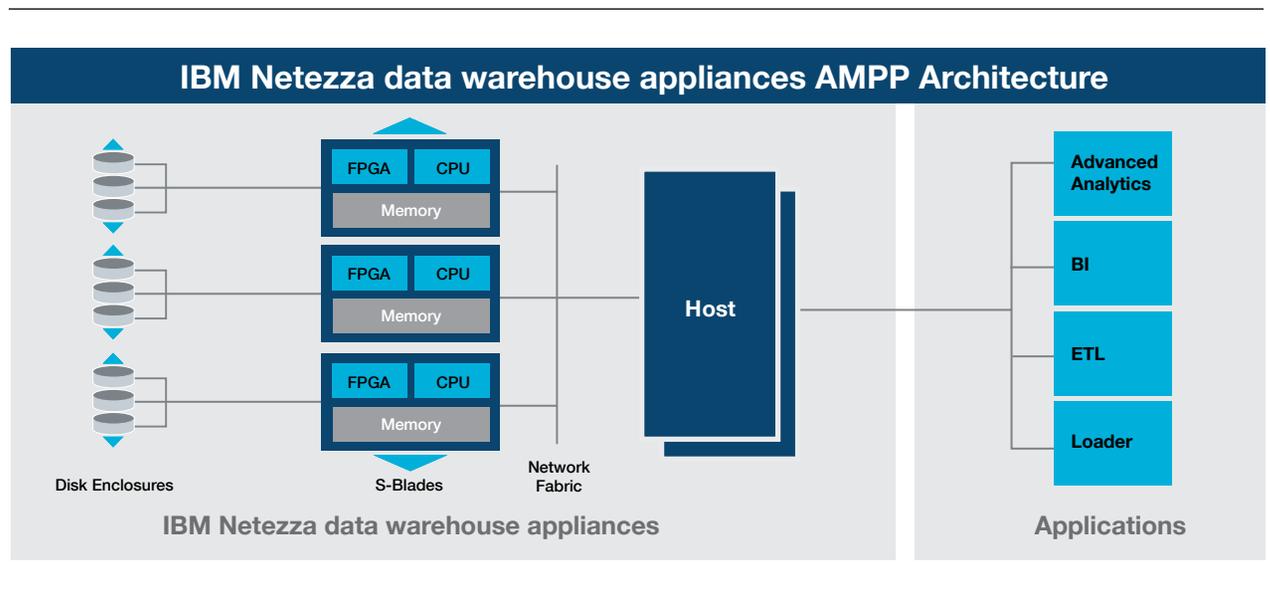
Management Overhead
Administrators design and define data distribution via partitions, files, tablespaces, and block/extent sizes.

Analytics Limitations
Exadata cells do not process distinct aggregations (common in simple reports), complex joins or analytical functions (used in analytical applications).

Marrying its existing database technology with a new “smart” storage tier, Exadata removes the disk throughput bottleneck Oracle suffers when partnered with conventional storage. Exadata may appear to present an interesting opportunity for CIOs looking to consolidate multiple OLTP systems to a single platform. However, upon closer examination, it becomes clear that the amount storage and storage software included in Exadata are overkill for the vast majority of OLTP environments. For all but simple queries Exadata must move large sets of data from its storage tier to its database tier raising questions on its suitability as a platform for a modern data warehouse.

Query performance with the IBM Netezza data warehouse appliance

The IBM Netezza data warehouse appliance is designed from the ground up as a data warehousing platform. IBM Netezza employs an Asymmetric Massively Parallel Processing (AMPP) architecture. A Symmetrical Multiprocessing host⁷ fronts a grid of Massively Parallel Processing nodes. The IBM Netezza data warehouse appliance exploits this MPP grid to process the heavy lifting of warehousing and analyzing data.



⁷ IBM Netezza data warehouse appliances have two SMP hosts for redundancy but only one is active at any one time.

A node in IBM Netezza data warehouse appliances grid is called an S-Blade (Snippet Blade), an independent server containing multi-core central processing units (CPUs). Each CPU is teamed with a multi-engine Field Programmable Gate Array (FPGA) and gigabytes of random access memory. Because the CPUs have their own memory, they remain focused exclusively on data analysis and are never distracted to track block ping-pong or cluster freeze activity at other nodes as occurs in shared memory database systems.

An FPGA is a semiconductor chip equipped with a large number of internal gates programmable to implement almost any logical function, and particularly effective at managing streaming processing tasks. Outside of the IBM Netezza data warehouse appliance, FPGAs are used in such applications as digital signal processing, medical imaging and speech recognition. The IBM data warehouse appliance engineering team has built software machines within our appliances' FPGAs to accelerate processing of data before it reaches the CPU. Within each Exadata rack Oracle dedicates 14 eight-way storage servers to accomplish less than the IBM Netezza data warehouse appliance achieves with 48 FPGAs embedded within our blade

servers. Each FPGA – just a 1"x1" square of silicon – achieves its work with enormous efficiency, drawing little power and generating little heat.

IBM didn't take an old system with known shortcomings and balance it with a new smarter storage tier; IBM Netezza data warehouse appliances are designed as optimized platforms for data warehousing.

Inter-nodal communication across the IBM Netezza MPP grid occurs on a network fabric running a customized IP-based protocol fully utilizing total cross-sectional bandwidth and eliminating congestion even under sustained, bursty network traffic. The network is optimized to scale to more than a thousand nodes, while allowing each node to initiate large data transfers to every other node simultaneously. These transfers bring enormous efficiency to the processing tasks typical of data warehousing and advanced analytics. Just as SQL

statements benefit from processing within the IBM Netezza data warehouse appliance MPP architecture, so too do computationally complex algorithms at the heart of advanced analytics. Previous generations of technology physically separate application processing from database processing, introducing inefficiencies and constraints as large data sets are shuffled out of the warehouse to the analytic processing platforms and back again. The IBM Netezza data warehouse appliance brings the heavy computation of advanced analytics into its MPP grid, running the algorithms in each CPU physically close to the data, making data movement redundant and boosting performance. The algorithms benefit from running on the many nodes of the IBM Netezza data warehouse appliance MPP grid, freed from constraints imposed on less-scalable clustered systems.

Unlike the competition, IBM didn't take an old system with known shortcomings and balance it with a new smarter storage tier; IBM Netezza data warehouse appliances are designed as optimized platforms for data warehousing. IBM Netezza data warehouse appliances deliver performance generously, making life easy for programmers, administrators and users.

A customer of an IBM Netezza data warehouse appliance from the financial services industry used the Lean approach to analyze resource expenditure required to manage their Oracle data warehouse. They learned in building and maintaining indices, aggregates, materialized views and data marts that more than 90 percent of their IT team's work was either required waste or non-value added processing.

Simplicity of operation

Simplicity of operation with Oracle Exadata

Before the warehouse can run queries it must be loaded with data. Exadata's storage tier is an MPP grid. MPP systems achieve performance and scale when all nodes participate equally in the computational task at hand. Data must be evenly distributed, with the same amount of relevant data at each node for each query, to the extent possible. To evenly distribute data across Exadata's grid of storage servers requires administrators trained and experienced in designing, managing and maintaining complex partitions, files, tablespaces, indices, tables and block/extent sizes. "Even better might be a system that doesn't lean heavily on complex partitioning to achieve good performance."⁸

A customer of IBM from the financial services industry used the Lean⁹ approach to analyze resource expenditure required to manage their Oracle data warehouse. They learned in building and maintaining indices, aggregates, materialized views and data marts that more than 90 percent of their IT team's work was either required waste or non-value added processing. The cost of this waste translates to unnecessary hardware and software license costs, terabytes of wasted storage, elongated development and data load cycles, long periods of data unavailability, stale data, poorly performing loads and queries and excessive administrative costs.

Exadata does little to simplify managing an Oracle data warehouse. Administrators must manage multiple server layers, each with operating system images, firmware, file systems and software to be maintained. Oracle suggests that DBAs should expect to spend 26 percent less time managing 11g, the database version in Exadata, than they spend on older 10g deployments. If this is confirmed in practice and Exadata reduces by a quarter the time customers waste in valueless administration, Oracle has taken a step in the right direction. IBM Netezza's appliances are designed not to waste any of the customers' time. "The DBA team only backs up the environment and manages the high level security model for the appliance and that is it. They don't need to do anything else (for example, the concept of indexing is foreign to them when dealing with IBM)."¹⁰

⁸ Curt Monash at <http://www.dbms2.com/2009/09/21/notes-on-the-oracle-database-11g-release-2-white-paper/>

⁹ With roots in manufacturing, "Lean" is a practice using tools and techniques of Six Sigma to analyze wasteful expenditure of resources, and target activities adding no value to the product or service for elimination.

¹⁰ Customer using Oracle for OLTP and Netezza for data warehousing quoted from LinkedIn Exadata Vs Netezza forum at http://www.linkedin.com/groupAnswers?viewQuestionAndAnswers=&gid=2602952&discussionID=11385070&sid=1275353329699&trk=ug_qa_q&goback=ana_2602952_1275353329699_3_1

Not only do business users demand that their queries complete quickly, they also expect consistent performance; a report that completed in five seconds yesterday and three minutes today will likely create a ticket requiring a response from IT helpdesk staff. Warehouses are inevitably subject to the demands of varied, dynamic workloads. Data arriving from OLTP systems via batch jobs or trickle feeds are loaded, administrative tasks such as backup and restore and grooming run in the background out of view of the business and dashboards are constantly updating. At the same time, computational intensive applications – such as those predicting which claims or trades might be fraudulent or irregular – create sudden, heavy load on the warehouse infrastructure. Delivering consistent performance to the business makes two requirements of the warehouse: consistent query performance and effective workload management. This simplifies allocation of available computing power to all the jobs requiring service, usually based on priorities agreed with the business.

Oracle's philosophy of workload management is to offer administrators multiple tuning parameters. Oracle's parameters have a high degree of dependency on one another, and in Exadata some

must be set to the same value for every processor in their grid. This complexity forces administrators to experimentally change parameter settings, tuning their way around unexpected demands on the warehouse. Achieving and maintaining consistent performance for large communities of users, with different application and data requirements, through rising and falling loads, is a complex task requiring a high degree of Oracle experience and expertise of the warehouse administrators.

In OLTP systems with a stable, well-understood population of transactions the business can be shielded from this complexity. Database administrators have ample opportunity during an application's development phase to analyze each operation's data requirements and have the time to design, test and tune the database. Data warehouses are different. An event in the outside world creates the need to analyze data in ways never before attempted. The immediate need for information leaves no time for administrators to analyze each query and optimize its data retrieval. A warehouse unable to process requests immediately, as they are formulated, denies the business opportunities for action.

“The way we did a proof of concept with them [IBM Netezza data warehouse appliance] was, they shipped us a box, we put it into our data center and plugged into our network. Within 24 hours, we were up and running. I’m not exaggerating, it was that easy.”

Simplicity of operation with an IBM Netezza data warehouse appliance

IBM Netezza's customers willingly confirm that our appliances are simple to install and use. “The way we did a proof of concept with them [IBM Netezza] was, they shipped us a box, we put it into our data center and plugged into our network,” he said. “Within 24 hours, we were up and running. I’m not exaggerating, it was that easy.”¹¹

This commentary is from the vice president of technology at a leading social networking company that is already using Oracle's database and RAC software.

¹¹ http://www.computerworld.com/s/article/9126930/EHarmony_finds_data_warehouse_match_with_Netezza?source=rss_news

There's something to be said for a simple approach

- NO cluster interconnect (GES and GCS) monitoring/tuning
 - NO RAC-specific knowledge/tuning (DBAs with RAC experience are less of a commodity)
 - NO dbspace/tablespace sizing and configuration
 - NO redo/physical log sizing and configuration
 - NO journaling/logical log sizing and configuration
 - NO page/block sizing and configuration for tables
 - NO extent sizing and configuration for tables
 - NO temp space allocation and monitoring
 - NO integration of OS kernel recommendations
 - NO maintenance of OS recommended patch levels
 - NO JAD sessions to configure host/network/storage
 - NO query (e.g. first_rows) and optimizer (e.g., optimizer_index_cost_adj) hints
 - NO statspack (statistics, cache hit, wait event monitoring)
 - NO memory tuning (SGA, block buffers, etc.)
 - NO index planning/creation/maintenance
 - Simple partitioning strategies: HASH or ROUND ROBIN
-

Reducing the time to get productive is a good start; IBM's philosophy is to bring simplicity to all phases of data warehousing. The first task facing a customer is loading their data. An IBM Netezza data warehouse appliance automates data distribution. Experience from proof-of-concept projects is that customers load their data to an IBM Netezza data warehouse appliance using automatic distribution, run their queries and compare results to their highly tuned Oracle environments. For all but the simplest queries, automatic distribution is good enough for an IBM Netezza data warehouse appliance to outperform Oracle. Customers may later analyze all their queries to identify those that can be accelerated by redistributing data on different keys. The IBM Netezza data warehouse appliance makes this task simple.

All queries submitted to the IBM Netezza data warehouse appliance are automatically processed in its massively parallel grid with no involvement of database administrators. Queries and analyses enter the IBM Netezza data warehouse appliance through the host machine where the optimizer, the compiler and the scheduler decompose them into many different pieces or snippets, and distribute these instructions to the MPP grid of processing nodes, or S-Blades, all of which then process their workload simultaneously against their locally-managed slice of data.

A Snippet arriving at each of the IBM Netezza data warehouse appliance S-Blades initiates reading of compressed data from disk into memory. The FPGA then reads the data from memory buffers and utilizing its Compress Engine decompresses it, instantly transforming each block from disk into the equivalent of 4-8 data blocks within the FPGA. The engineering behind the IBM Netezza data warehouse appliance accelerates the slowest component in any data warehouse – the disk. Next, within the FPGA data streams into the Project Engine which filters out columns based on parameters specified in the SELECT clause of the SQL query being processed. Only records fulfilling the SELECT clause are passed further downstream to the Restrict Engine where rows not needed to process the query are blocked from passing through gates, based on restrictions specified in the WHERE clause. The Visibility Engine maintains ACID (Atomicity, Consistency, Isolation and Durability) compliance at streaming speeds. All this work, the constant pruning of unneeded columns and rows, is achieved in an energy efficient FPGA measuring just one square inch. If the engineering behind the IBM Netezza data warehouse appliance doesn't need to move data, it doesn't.

The FPGA's pre-processing complete, it streams just the resulting trimmed down set of records back into S-Blade memory where the CPU performs higher-level database operations such as sorts, joins and

aggregations, doing this in parallel with all other CPUs within the MPP grid. The CPU may also apply complex algorithms embedded in the Snippet code for advanced analytics processing. The CPU finally assembles all the intermediate results from the entire data stream and produces a result for the Snippet, sent over the network fabric to other S-Blades or the host, as directed by the Snippet code. When data required by a JOIN is not collocated on a node, the IBM Netezza data warehouse appliance inter-nodal network fabric efficiently and simply re-distributes late in the processing cycle after the database has completed restrictions and projections. Some highly complex algorithms require communication among nodes to compute their answer. The engineering behind the IBM Netezza data warehouse appliance exploits a message passing interface to communicate interim results and to produce the final result.

And, as the original compressed data blocks are still in memory, they can be automatically reused in later queries requiring similar data via the IBM Netezza data warehouse appliance table cache – an automated mechanism requiring no DBA training or involvement.

Just three months after moving to an IBM Netezza data warehouse appliance, a customer relates that his team delivered more analytical applications than they could in the previous three years with Oracle.

Because the IBM Netezza data warehouse appliance applies full parallelism to all tasks, its workload management system plays a critical role in controlling how much of the appliance's computing resources are made available to each and every job. In the IBM appliance architecture, one software component controls all system resources: processors; disks; memory; network. This elegance is the foundation of the IBM Netezza data warehouse appliance Workload Management System. The IBM Netezza data warehouse appliance Workload Management System makes it simple for administrators to allocate computational resources to users and groups based on priorities agreed with the business and maintain consistent response times for multiple communities.

IBM Netezza data warehouse appliances eliminate the wasted work of database tuning. Equipped to make their own intelligent decisions, IBM Netezza appliances require no tuning and little system administration. The few administrative tasks necessary to maintain consistent performance through dynamic, changing workloads are within easy reach of anyone with Linux and SQL experience. All that is required of the administrator is to allocate the IBM Netezza data warehouse appliance resources to groups within the user community and hand control to the Workload Management System. Freed from constant cycles of database administration, technical staff engages with the business to investigate new, value-creating ways of exploiting data. Just three months after moving to an IBM Netezza data warehouse appliance, a customer relates that his team delivered more analytical applications than they could in the previous three years with Oracle. Processing analytical applications close to where data is managed, exploiting the same MPP platform as used for processing SQL, represents a real opportunity for organizations to increase dramatically the value they derive from data.

Value

Value with Exadata

As the waste analysis conducted by the financial services customer of both an IBM Netezza data warehouse appliance and Oracle highlights, using Oracle for data warehousing is labor intensive. Oracle suggests the latest version of their database management system might reduce this waste by 26 percent.¹² IBM Netezza data warehouse appliance customers attest that these low level, technically-demanding administration tasks are simply unnecessary; in this light it is indefensible that operating an Oracle database demands administrators spend the majority of their time on care and feeding of the underlying technology, while IBM Netezza data warehouse appliance customers spend that time creating value by exploiting their data.

Exadata's new storage tier adds another layer of complexity for administrators to tune and manage. Because Exadata is very new, and so few data warehouses using the technology are in production, projections on its cost of ownership are premature. However, customers should expect that achieving consistently high performance from Exadata will incur substantial costs in database design and administration.

While adding a new storage tier removes the disk throughput bottleneck to Oracle's database, Exadata's engineering is more adapted to massively

parallel processing than full exploitation of the architecture. Oracle's failure to integrate data management fully into Exadata's storage tier means too little is asked of hardware in its MPP grid. This inflates the cost of acquiring Exadata; customers pay for hardware that will never be fully exploited, and they pay for a new additional layer of storage software that has limited capabilities. These costs build over the lifetime of the warehouse. Customers pay for under-utilized space in their data centers which would return greater value if used to house a more efficient computer system.

While costs undermine value, a fundamental question is whether Exadata helps customers to create value. First-generation data warehouses play an important role in keeping an organization informed of the recent past, yet data unleashes greater potential through advanced analytics and other capabilities of second-generation warehouses discussed earlier in the paper. Oracle RAC teamed with traditional storage has had limited technical success in this area and has yet to be proved a success in this role to date. Exadata's storage tier is unable to process complex joins, distinct aggregations and analytical functions. It is difficult to envisage how two technologies, individually ill-equipped to analyze deeply very large data sets with high performance, will achieve this feat when connected by a fast network and housed in the same rack.

Value with IBM Netezza data warehouse appliances

The engineers of IBM Netezza data warehouse appliances integrate data management and analysis deep within massively parallel, shared-nothing grids. One result we plan from this innovation is simplicity for our customers, which translates directly to dramatically lower costs of owning and operating data warehouses than is possible with traditional database products, such as Oracle's.

Demands on data warehouses have moved beyond processing simple SQL; to fully exploit data requires the warehouse be capable of running predictive models, investigative graphs and other analytic applications. To illustrate, a financial services company – knowing the next most probable purchase by a family that recently purchased a mortgage and previously purchased investment products, loan products and a checking account but has never purchased insurance policies, is an investment product followed by another mortgage – can create targeted marketing campaigns of value to the customer and with a high chance of success.¹³

¹² <http://www.dbms2.com/2009/09/21/notes-on-the-oracle-database-11g-release-2-white-paper/>

¹³ See Dynamic Bayesian Networks for acquisition pattern analysis: a financial-services cross-sell application by Anita Prinzie, Marketing Group, Manchester Business School and Dirk Van den Poel, Department of Marketing, Ghent University

Evaluating the Systems		IBM Netezza	Oracle
Item		IBM Netezza 1000	Exadata v2 (SAS)
Performance and Architecture	MPP	<ul style="list-style-type: none"> · True MPP · Optimized for Data Warehousing and Analytics 	<ul style="list-style-type: none"> · Hybrid – parallel storage nodes and SMP clustered head node · A generalized architecture
	Hardware Architecture	<ul style="list-style-type: none"> · Full processing S-Blades (1 CPU core + 1 FPGA core / 1 disk drive) · SMP host node used primarily for user/applications interface · Independent blade-to-blade redistribution 	<ul style="list-style-type: none"> · Intelligent storage (1 CPU core / 1.5 disk drives) · SMP Cluster nodes running Oracle 11g RAC · InfiniBand (Exadata nodes to SMP cluster) · Head node engagement in all data redistributions
	Data Streaming	<ul style="list-style-type: none"> · FPGA performance assist on S-Blade – decompression, predicate filtering, row-level security enforcement · >95 percent of work done on S-Blades 	<ul style="list-style-type: none"> · Exadata nodes primarily used for decompression and predicate filtering · Most DW and Analytics work done in SMP head node
	In-Database Analytics	<ul style="list-style-type: none"> · Fully engaged MPP platform for analytics · User-defined functions, aggregates and tables · Language support: C/C++, Java, Python, R, Fortran · Paradigm support: SQL, Matrix, Grid, Hadoop · Built-in set of >50 key analytics (fully parallelized) · Integrated Development Env.: Eclipse and R GUI w/ wizards 	<ul style="list-style-type: none"> · Analytics processing limited to head node cluster only · User-defined functions and aggregates · Language support: C/C++, Java · Paradigm support: SQL, Matrix (minor) · Basic analytics functions
	Scale	<ul style="list-style-type: none"> · Linear performance and data size scalability · Full-featured, enterprise-class workload management and other features 	<ul style="list-style-type: none"> · Non-linear performance and data size scaling – performance and i/o bottleneck at/to head node cluster
Simplicity	Appliance System Mgmt and Integration	<ul style="list-style-type: none"> · No tuning, no indexing, no partitions · Balanced system developed to deliver best price-performance 	<ul style="list-style-type: none"> · Heavily tuned performance dependency · Performance depends on physical database design skills, including indices and partitions

This analysis, beyond SQL's capabilities, requires a technique called Dynamic Bayesian Networks. However, the analysis uses the same data processed by SQL to create reports and dashboards suggesting an expanding role for the warehouse.

IBM Netezza data warehouse appliances are designed from the ground up for processing both SQL and the applications of advanced analytics. IBM Netezza data warehouse appliances free customers from proprietary languages. Customers can port existing applications to IBM Netezza data warehouse appliances or choose to develop new analytic applications in the language of their choice, including R, C/C++, Java, Python, and Fortran. Customers can leverage a built-in library of parallelized, in-database algorithms, including data preparation, data mining, predictive analytics, geospatial, and matrix algebra. Additionally customers can choose to work with Hadoop / MapReduce as, for example, a highly scalable ingestion mechanism to preprocess enormous data sets generated by public facing web applications and web logs before they are loaded into IBM Netezza data warehouse appliances for on-demand analysis.

Conclusion

IBM Netezza data warehouse appliances emerged as a principal alternative to Oracle for data warehousing. Moving data warehouses and marts from Oracle to an IBM Netezza data warehouse creates new opportunity, not risk. A majority of IBM customers have already walked this path, many of them by partnering with system integration companies with strong track records for successful migrations. Exadata is an evolution of Oracle's OLTP platform, and is positioned as a general purpose platform for both OLTP and analytics. Oracle's database management system is designed for OLTP where data volumes are relatively modest compared to data warehouses. The database activity of an OLTP system can be assessed before it is put into production; administrators have the time to design, test and optimize each transaction's data retrieval. Data warehouses must immediately process whatever query the business needs to ask of their data; technologies requiring administrator mediation are ill-suited to the task. Conscripting this technology into a role other than transaction processing places enormous stress on people and processes harnessed to manage and operate a data warehouse.

“This [IBM Netezza data warehouse appliance] is the first database product with a long term product roadmap that aligns perfectly with our own roadmap. We call this our on-demand database.”

— Chief Data Officer, Large Equities Exchange Group

Oracle advises customers that Exadata is architecturally similar to IBM Netezza data warehouse appliances but better because IBM Netezza data warehouse appliances do not support every data type or SQL standard, and that it doesn't support data mining or high concurrency. Customers of IBM Netezza data warehouse appliances disagree: “This [IBM Netezza data warehouse appliance] is the first database product with a long term product roadmap that aligns perfectly with our own roadmap. We call this our on-demand database,”¹⁵ said the chief data officer of a large equities exchange group.

¹⁵ www.netezza.com/customers/nyse-uronext-video.aspx

Given their different workload characteristics, few customers attempt to run OLTP and data warehouse systems on the same infrastructure; to do so demands constant tuning and optimizing.

Technicians are placed in a difficult situation: either accept compromised performance for both OLTP and data warehousing, or ceaselessly reconfigure the database in a vain attempt to satisfy conflicting demands of the different workloads.

As mentioned earlier in this eBook, organizations evolving into second-generation data warehouses are running their OLTP and warehouse systems on different platforms, each specifically configured to the needs of their workloads.

The only data warehouse that really matters is your data warehouse – your applications running on your data in your data center. An on-site proof-of-concept (PoC) creates the opportunity for an IT department to thoroughly investigate a technology, learning how they can use IBM Netezza data warehouse appliances to help their business peers extract greater value from data. Making the most of this opportunity requires the PoC to be managed with the same discipline afforded other projects.

Curt Monash offers sage advice in his blog “Best practices for analytic DBMS POCs,”¹⁶ including involving an independent consultant to steer the project to a successful outcome. For qualified organizations wanting to understand how their warehouse performs on an IBM Netezza data warehouse appliance, at no cost and with no risk, IBM offers its TestDrive. To book one, go to <http://www.ibm.com/software/data/netezza/>.

IBM Netezza data warehouse appliances:
to use it is to enjoy it.

¹⁶ www.dbms2.com/2010/06/14/best-practices-analytic-database-poc/ffimore-2297

Pass it along

Please share this document with your friends and colleagues.

Give us feedback

What did you think about the ideas and arguments in this eBook? Let us know what you liked, disliked or might want to discuss further.

- Chat with the IBM Netezza Data warehouse Community:
www.enzeecomunity.com/groups/

Contact us

- Visit IBM Netezza’s blogs:
<http://thinking.netezza.com>
- Visit the IBM website: <http://www.ibm.com/software/data/netezza/>
- Visit the IBM Netezza Data Warehouse Community website: www.enzeecomunity.com

About the authors

Phil Francisco, Vice President, Product Management and Product Marketing, IBM

Phil Francisco brings over 20 years of experience in technology development and global technology marketing. As Vice President of Product Management and Product Marketing at IBM, he fosters new business and product strategies, directs the product portfolio and drives product marketing programs. Prior to IBM, Francisco was the Vice President of Marketing at PhotonEx, a leading developer of 40 Gb/s optical transport systems for core telecommunications network providers. Before PhotonEx Francisco served as Vice President of Product Marketing for Lucent Technologies' Optical Networking Group, where he worked with some of the world's largest telecommunications carriers in planning and implementing optical network solutions. Mr. Francisco holds a patent in advanced optical network architectures. He received B.S. in Electrical Engineering and B.S. in Computer Science degrees magna cum laude, from the Moore School of Electrical Engineering at the University of Pennsylvania. He earned his Master's degree in Electrical Engineering from Stanford University and completed the Advanced Management Program at the Fuqua School of Business at Duke University. Read Phil's blog: http://thinking.netezza.com/blogs_by/phil-francisco.

Mike Kearney, Senior Director, Product Marketing, IBM

Mike has worked in information technology for more than twenty-five years. At IBM he specializes in communicating the value of Netezza's appliance performance, value and simplicity to organizations looking to derive greater return from their data. Mike has previously worked at Vignette, BMC Software and Oracle Corporation and at companies in the financial services, telecommunications and energy industries. Mike has a B.Sc. from London University and an M.Sc from Coventry University.

About Netezza, an IBM Company

Netezza, an IBM Company is a global leader in data warehouse, analytic and monitoring appliances that dramatically simplify high-performance analytics across an extended enterprise. Netezza's technology enables organizations to process enormous amounts of captured data at exceptional speed, providing a significant competitive and operational advantage in today's data-intensive industries, including digital media, energy, financial services, government, health and life sciences, retail and telecommunications.

About IBM Data Warehousing and Analytics Solutions

IBM provides the broadest and most comprehensive portfolio of data warehousing, information management and business analytic software, hardware and solutions to help customers maximize the value of their information assets and discover new insights to make better and faster decisions and optimize their business outcomes.

For more information

To learn more about the IBM Data Warehousing and Analytics Solutions, please contact your IBM sales representative or visit <http://www.ibm.com/software/data/netezza/>

© Copyright IBM Corporation 2011

IBM Corporation
Software Group
Route 100
Somers, NY 10589 U.S.A.

Produced in the United States of America
August 2011
All Rights Reserved

IBM, the IBM logo, ibm.com and Netezza are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

Microsoft and SQL Server are trademarks of Microsoft Corporation in the United States, other countries or both.

Oracle and Exadata are registered trademarks of Oracle and/or its affiliates

Other company, product and service names may be trademarks or service marks of others.



Please Recycle