

Data Warehousing Appliances –Next Generation Approach to the Challenges of Big Data Environments

April 2007



Today's enterprises face an untenable problem with respect to their data warehousing and business intelligence infrastructure. On the one hand, companies must keep and store more data online longer to gain greater insight into their customer base and drive new levels of profitability, retention and loyalty. To that end, their data warehousing environments become strategic assets of the line of business.

However, on the other hand, current data warehousing technology is either too cost prohibitive, inflexible, or too operationally expensive to meet the ever increasing demands of the business. In short, IT is caught between a rock and a hard place – unable to cost-effectively meet escalating business demands with current approaches and technologies. As a result, data driven corporations are quickly realizing that current approaches to scaling and managing their data warehouses are insufficient in light of the new realities.

Enter the Data Warehouse Appliance (DWA). The DWA promises a self-tuning, self managing system that enables a DBA to modularly scale performance, while providing a low acquisition cost. As part of our ongoing research into data warehousing and big data environments, we have conducted a series of one-on-one interviews with Global 2000 IT executive management, storage administrators, and database administrators to understand the current challenges and requirements of large scale data warehousing. This technology brief summarizes the key findings of those interviews, reviews the current technology approaches to data warehousing (SMP, Clustering, and MPP) and posits how a new category of data storage technology –data warehousing appliances – is evolving to meet these new challenges.

Challenges of Big Data Environments

To meet the ever increasing requirements of the business, IT must keep a tremendous amount of current and historical information online. As a result, big data environments are becoming increasingly commonplace. Taneja Group classifies any data warehouse that is bigger than 5 TB as a big data environment. Clearly, the problems with

managing, tuning, troubleshooting these big data environments are exacerbated as the amount of the data grows. From our interactions with end users, we have found that there are three main challenges that big data environments enable for creating, scaling, and managing large data warehouse.

CHALLENGE #1: KEEP MORE DATA ONLINE LONGER

For users in charge of managing data warehousing environments, the amount of structured content is growing at an alarming rate. Many users with whom we have spoken cite that their data warehousing environments are doubling in size every year.

The primary driver for this growth in the database size is coming from the line of business. Business decision makers recognize the value for maintaining more historical data online for better analytics and decision making purposes. Traditionally, a corporation might keep three months of historical sales and transaction data online for detailed analysis. Increasingly, the business leaders are insisting that IT keep longer and richer histories so that marketing and sales professionals can conduct more complex scenario planning to determine the impact of different programs on the bottom line.

A secondary driver fueling the size of databases is due to a tightening regulatory and compliance environment. With recent regulations such as Sarbanes-Oxley (SOX), Healthcare Insurance Portability and Accountability Act (HIPAA), Gramm-Leach Bliley Act (GLBA), and Payment Card Industry (PCI), IT must now retain structured data online longer to meet eDiscovery and compliance requirements imposed by this legislation. As a result, this new regulatory environment exacerbates the already high organic growth of structured data.

As a result, DBA and storage administrators are being asked to greatly expand and manage much larger database and storage environments, yet still maintain or improve query and data analytic response times. Unfortunately, the overall performance and response time of current RDBMS systems is impacted as the database size increases. This fact is particularly true as the database grow in excess of 1TB. Techniques like database archiving allow IT to prune the size of database to improve performance, but don't necessarily allow that data to be kept online and fully queryable. To hit its performance SLAs, IT must leverage a variety of approaches to scale data warehouse, but each approach has its own drawbacks in terms of cost, scalability, and overall management. We will review the current approaches (SMP, Clustering, MPP) to scaling database performance in a latter section of this document.

CHALLENGE #2: THE EVER GROWING SERVER AND STORAGE PERFORMANCE GAP

Compounding this 100+% growth rate of structured data is the fact that there exist a huge technology asymmetry between the ability to process information and the ability to access information. Many have dubbed this gap or asymmetry the server and storage performance gap. Simply put, the access and retrieval times of spinning magnetic disks have not kept up with the advances in CPU processing capabilities. For example, CPU performance has dramatically improved, driven by Moore's law, multi-core processors, and threading technology to increase 2,000,000 times since 1987. In comparison,

T E C H N O L O G Y B R I E F

disk access time performance only improved by 11 times.

In today's data warehousing systems, disk drives represent the only mechanical device in the data path. Due to the physical limitations of how fast a drive can spin and seek data, each disk can only deliver a finite amount of I/Os per second. However, disk I/O rates are extremely slow compared to the rate of modern day CPUs to process data and the network rates to transmit it. As a result, the disk becomes a critical bottleneck that must be addressed in any large scale data warehouse implementation.

IT faces huge challenges in coaxing significant I/O throughput and response times of the underlying storage system in order to meet the insatiable requirements of a large RDBMS engine. Clearly, the overall throughput and response time of the underlying storage infrastructure directly affect what end user see as query response time. Storage performance tuning and layout is a critical issue in database design and scaling multi-terabyte data warehouse implementations.

CHALLENGE #3: HIGH OPEX COSTS OF TUNING & MANAGING BIG DATA WAREHOUSES

As the database size grows, the administrative overhead of managing a database grows exponentially along two dimensions – database management and tuning and storage management and tuning. For the DBA, a complex data warehouse requires a heavy investment in designing and profiling queries for optimum performance,

balancing data load times against the need to index as many columns for maximum query flexibility and performance, and partitioning data to avoid hot spot and bottlenecks. In particular, partitioning of data in a data warehouse magnifies the administrative overhead because it means that every database management operation must be performed separately on each partition. It is not uncommon to have a single DBA dedicated to managing a single table or set of tables in a very large data warehouse.

For the storage administrator, large databases result in potentially laborious, repetitive provisioning of LUNs, careful balancing of I/O performance and data protection requirements on the underlying physical storage media, evenly striping and re-balancing data across multiple disks and storage systems to achieve maximum I/O performance and access times. All of these operations are time intensive and many of these tasks defy automation.

Furthermore, the type of tuning and management required to maintain a large scale data warehouse requires some of the most highly skilled professionals. As can be imagined, as the database grows in scale, the amount a business must spend in terms of OPEX to maintain and grow it, increases dramatically. The OPEX cost of administering a large scale data warehouse does not grow linearly or in proportion to the database size, instead the OPEX costs scale exponentially as the size of the database grows. OPEX costs can be the number one inhibitor to further growing a very large data warehouse.

Traditional Approaches to Data Warehousing

Due to the tremendous growth in data warehouse size over the past decade and the need to improve query response times, three separate architectural approaches have evolved to address the scalability challenges of big data environments. In this section, we briefly summarize each of the three main architectural approaches and their advantages and disadvantages as it relates to cost, scalability, and overall ease of data management.

SYMMETRIC MULTI-PROCESSING

SMP, also commonly referred to as scale up, addresses data warehousing scalability by adding more CPU power in a tightly connected system. SMP systems, also known as a "shared everything" system, are configured such that all of the CPUs share a single memory space and I/O bus. This physical hardware architecture greatly simplifies the database design because all commercial databases (Oracle, DB2, SQL Server) are designed to be multi-threaded. As a result, a single database instance and its various processes or threads can be scheduled across all available CPUs and thus take advantage of the added cores or CPUs in an SMP design.

SMP systems are by far the most common deployment model for OLTP databases and small data warehouses or data marts. However, as the size of the database grows, the SMP approach has two distinct disadvantages. First, SMP systems can be quite expensive and they require IT to

provision the system for the peak expected load of the system during the equipments life. A typical 64 CPU large SMP UNIX system can easily cost \$1 to \$2M dollar for the physical hardware, not including the actual storage capacity or database software licenses required for the data warehouse. Second, the SMP model does not support a "pay as you grow" model for scaling the data warehouse. In short, the IT team must do a forklift upgrade of the server and add a bigger, faster server with more CPUs to add processing power to a redlining data warehouse. A forklift upgrade is disruptive to the business operations and can be very labor intensive.

CLUSTERED

Clustered databases fall into the camp of scale out systems; systems that can be modularly expanded and scaled by adding additional server capacity. Traditionally, clustering has been a technology for ensuring high availability. However, Oracle with its Oracle Real Application Clusters architecture has pioneered a shared disk form of database clustering that some users have used for data warehousing.

In a shared disk clustered system, multiple SMP servers communicate over a shared, high speed interconnect and share access to a single disk resource. Shared disk databases make heavy use of a lock manager to arbitrate and synchronize access to the underlying shared disk resource. Furthermore, shared disk databases, such as Oracle RAC, also implement a shared cache across all the physical nodes in the cluster, allowing nodes to share the contents of their

T E C H N O L O G Y B R I E F

volatile buffer caches through the cluster interconnect. Data blocks are shipped directly from one node to another using interconnect messaging, eliminating the need for extra disk I/Os. However, in big data environments, end user report that the interconnect can become saturated with too many data block transfers, resulting in scalability limitations.

Nonetheless, clustered database systems, such as Oracle RAC, offer the promise of delivering near linear scalability by adding low cost modular servers and storage to grow processing and storage IO throughput capabilities on demand. Moreover, the cluster appears as a single database to the application and can be managed just like a single instance database running on an SMP system. Therefore, shared disk databases are often seen as low cost alternatives to buying big iron SMP servers.

However, shared disk clustered databases are not without their drawbacks. An un-partitioned clustered database system faces diminishing returns as it is scaled out. For example, a shared disk database will incur significant lock and memory synchronization contention, along with interconnect contention in a large data warehousing environment. The risk in this environment is that the cluster will spend more of its time coordinating access and less of its time actually processing queries. As a result, it is common to see many share disk architectures be partitioned so as to mitigate the contention on a particular resource. This reduces contention and improves overall scalability, but it increases the OPEX administration of the database since now

each partition must be treated as a separate object within the database, greatly increasing database management overhead.

MASSIVELY PARALLEL PROCESSING

Massive Parallel Processing (MPP) systems have evolved to allow administrators to scale processing and database size on demand, but not be subjected to the diminishing returns as the system is scaled out. MPP, also commonly referred to as shared nothing systems, are designed around the concept that data warehousing workloads and queries can be cleanly divided into separate independently executable and parallelized operations across a federated system. MPP systems are clusters of two or more SMP server nodes, where each node has its own operating system, memory and exclusive access to a partitioned data set. Queries sent to the data warehouse are de-constructed into parallel queries that are executed by individual nodes. The results of these parallelized queries are rolled up and summarized after each node completes its processing.

Traditionally, MPP systems, such as Teradata and IBM DB2 EEE, have offered near unlimited scalability, but has come at the cost high overall management and OPEX cost to maintain a MPP system of significant scale. Typically, the setup for a MPP system is more complicated, requiring thought about how to partition a common database among processors and how to assign work among the processors. Moreover, once partitioned, administrators must engage in ongoing tuning to ensure that data is redistributed and partitioned optimally across all nodes in

MPP system. This re-partitioning can be quite challenging if the data warehouse is growing quite rapidly and users are performing new and different types of queries on the data set periodically. As a result, the implementation and ongoing management costs of a MPP databases can run in the tens of millions of dollars.

Data Warehousing Appliance – A Concept Whose Time Has Come

As the needs for deep, real time analytics have increased, the limitations and costs to the three architectural approaches have become evident. Given the continued 100+% growth in structured data projected, IT needs a system and architecture that delivers near limitless scalability and performance, while minimizing the laborious processes of provisioning storage, tuning the database, and adding additional processing power. In short, the user community wants a massively parallel data warehousing system that is easy to own and operate.

Enter the Data Warehousing Appliance (DWA). The origins of the term Data Warehousing Appliance (DWA) can be traced back to 2002 or 2003 when Foster Hinshaw, the founder of Netezza and now founder and CEO of Dataupia, coined the term. If Bill Inmon is the founder of the data warehouse, Foster Hinshaw should be considered the father of the DWA.

Essentially, DWAs simplify the deployment and operation of an MPP system so that corporations can cost-effectively scale their data warehouse. DWAs are generally defined to be a turnkey, fully integrated stack of

hardware, OS, and RDBMS software that is purpose built and optimized for data warehousing workloads. Ultimately, the vision of a DWA is to provide a self-managing, self-tuning, plug-and-play data warehousing system that can be scaled out in a modular and cost effective manner.

Although this may seem like nirvana, the vendor community has been particularly active in defining and innovating around this vision of the DWA. The grandfather and original data warehouse appliance is Netezza. However, since Netezza's market entry, several other firms, such as DataAllegro, Dataupia, and Kognitio have entered the market with variations and improvements on the original concept of a DWA.

End User Requirements for Data Warehousing Appliances

Our research and one-on-one interviews with large scale data warehousing end users confirmed the needs for a new approach to database management and scalability. From these interviews, four main requirements emerged that point to continued interest and adoption of DWAs.

1. **Requirement: Pay As You Grow Scalability.** End users want a solution that allows them to buy exactly what they want in terms of performance and capacity today, yet enable them to modularly expand the system on demand without disruption as the business needs dictate. Implicit in this requirement is the need for the DWA to be massively parallel and to offer almost limitless scalability so that the end user has

T E C H N O L O G Y B R I E F

investment protection built into his initial deployment.

- Requirement: Leverage Existing RDBMS Investments.** End users have deep investments in RDBMS technologies, such as Oracle, DB2, and SQL Server, and their corresponding tools. Users overwhelmingly indicated to us that they prefer to leverage their existing investments in infrastructure, tools, and training and are reluctant to throw away these investments in order to get the scalability and management that they crave in MPP DWAs. Presently, most DWAs, with the exception of Dataupia, force end users to give up their RDBMS engine of choice in order to attain the benefits of DWA. Clearly, as the DWA mature, this criterion will be a point of differentiation between the different DWA vendors.
- Requirement: Low Cost of Acquisition.** End users are looking to limit the initial CAPEX spend and grow their data warehousing environments over time. As a result, most end users indicated a preference for industry standard hardware, such as x86 processors and SAS/SATA drive technology, so that they could tap into the volume economics of those technologies and ride them down the cost curve. In short, users realize that the sheer scale and growth of these data warehousing environments mandate a need to use the most cost-effective processing and storage technologies available. This is another area where DWA vendors will differentiate themselves. Several vendors

are utilizing industry standard designs coupled with novel software technologies, while others have elected to developed custom ASICs and FPGAs in order to deliver high performance.

- Requirement: Simplicity of Operations.** Clearly, clustered shared disk databases and traditional shared nothing MPP architectures are proving too complex and cost prohibitive to scale and tune for larger data warehousing environments. As a result, end users are keenly interested in deploying systems that simplify the installation, provisioning of capacity, and tuning of the database system. To that end, the appliance model holds great promise. End users are looking for a system that will automatically redistribute data across partitions to optimize performance, provision additional capacity at a click of a button, and require no additional tuning, indexing, and partitioning to scale effectively. Ultimately, end users realize that as database size doubles per annum, they must find solutions that give them economies of scale in terms of data management. Otherwise, IT will be faced with a cost-prohibitive proposition of adding more human capital in order to scale, tune, and manage their strategic data warehouse.

Taneja Group Opinion

We believe that the current data warehousing management and scaling paradigm is fundamentally flawed. The continuing data tsunami, need to keep more structured data online longer, and demands to do real-time

T E C H N O L O G Y B R I E F

business analytics will force a change in how structured data in large data warehouses is managed over the next three to five years. A people intensive approach to tuning, scaling, and managing ever larger databases is untenable going forward. Business as usual simply will not cut it.

The answer to this looming data management crisis is DWA. DWA market is still in its infancy. Nonetheless, it offers the potential to allow IT to meet the demands of the line of business by delivering real-time analytics on ever growing data set composed of current and historical data. Based on our research and interviews, the need for a better

solution is readily apparent. For user planning a “big data” data warehouse, we recommend that you give careful consideration to what your data management and scalability requirements will look like 3-5 years out. It is our belief that data warehousing appliance will play a pivotal role in addressing these long term scalability and data management requirements in the future. If you are considering such an implementation, we strongly urge you to evaluate this next generation of data warehousing appliances. We don't think you will be disappointed.

NOTICE: The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.