

Overview

In today's increasingly complex and diverse telecommunications environment, the need for complete, accurate, and timely data analysis is essential to optimizing business processes for greater efficiency and profitability.

One of the major pain points facing telecom providers with a large number of subscribers, a high volume of data, and long-term storage needs is how to handle and analyze their business critical data in a profitable and efficient manner.

In the past, telecom providers had to rely on databases that were data segmented by application, department, and function; had highly relational Online Transaction Processing (OLTP) oriented data; were only able to process thousands to millions of records (rather than millions to billions of records); and had standardized data integrity rules across multiple application interfaces. All of which could be extremely time consuming, update centric, cumbersome, and costly.

Today, however, telecom providers are seeking greater agility and flexibility from their database solutions and TeleSciences has come up with the answer.

The TeleSciences Hybrid Database solution puts the power of data analysis back into the hands of telecom providers by providing them with access to a state-of-the-art database solution that not only utilizes data unified across application, department and function; it uses minimally relational flat/star oriented data, can process millions to billions of records, has high volume load/extract demands, and is capable of moving large volumes of data between existing datasets.

The Hybrid Database Solution¹

Utilizing a highly adaptive database solution is no longer just a competitive advantage – it is the answer to surviving in today's dynamic telecommunications industry.

¹ The name 'Hybrid Database' was chosen because it consists of both an expensive MPP as well as compressed file stores on an inexpensive Storage Area Network (SAN).

By using MPP (massively parallel processing)², proprietary algorithms, and domain-based business rules, the Hybrid Database solution allows a number of databases to be searched in parallel, thereby allowing business owners to gain credible, clear, and comprehensive insights into their data and focus on the effective use of that data. In addition, because the Hybrid Database solution is much more application and task specific than other available solutions, the process not only uses dynamic analysis to drill down into the data, it uses solution design to translate the information into usable data structures, which are capable of finding the root cause of the problem. In other words, the Hybrid Database solution makes it easier to answer ‘what is the problem?’ ‘Why did this problem occur?’ and ‘How can I fix it?’

Because the Hybrid Database is highly adaptive and was designed to work seamlessly with data warehouse appliances by market leaders Netezza and Teradata³, telecom providers can now run queries against billions of records in minutes (and often times seconds) and experience lightening fast data load and extract times of ¼ to ½ Terabyte (TB) per hour.

<Insert Gartner’s Magic Quadrant chart here. Insert chart comparing Netezza to other database solutions here. Insert chart illustrating Teradata’s node architecture. >

The Hybrid Database solution consists of a powerful and analytical database designed with ad-hoc reporting and analysis capabilities and a comprehensive, high performance file server integrated within core business analytics platform software. The TeleSciences solution is also fully integrated; thereby allowing for the seamless import and export of data between the high performance front end and the lower cost file storage server. And, universal access is available via any third party reporting, business intelligence or analytics tool.

And, since the data from each source is maintained separately and ONLY merged when a user request is made, the complexity of the database is dramatically reduced. This reduction in data complexity combined with state-of-the-art database techniques, such

²MPP is a type of computing that allows a number of databases to be searched in execute a single program.

³ Gartner has continuously rated Netezza and Teradata as the leaders in this market space.

as the temporal (time-related) placement of the data, results in allowing business owners to maintain many, many months of information and the nearly instantaneous delivery of its analytical reports.

How it Works

The Hybrid Database architecture is a highly efficient tool which allows large, complex datasets to be constantly searched, retrieved, analyzed, and stored in a timely manner. And, because searching and retrieving information from large datasets also means that output volumes can be extremely large, the Hybrid Database was designed to provide all of these functions in a rapid and cost effective fashion.

For example, telecom providers must have the ability to search and analyze CDR data efficiently and effectively. With the Hybrid Database, telecom providers now have a tool to perform this task swiftly and accurately because the CDR data is broken down by switch and loaded into the MPP at lightening speeds. However, because large telecom providers generate massive amounts of CDR data that would likely exceed the storage capability of the MPP, in the past providers were faced with only two options: purge the data or move the data. However, because the Hybrid Database was designed as part MPP and part compressed files stores on a SAN, once the storage capability of the MPP is exceeded, that data is automatically configured and compressed for 'long term' storage on the SAN; creating a much less expensive data storage option.

The Hybrid Database Architecture

The Hybrid Database has sophisticated query optimization routines, indexing, and partitioning to take advantage of high-performance parallel-processing and was specifically designed to address both the time and space complexities associated with using large datasets.

<Insert Hybrid DB Architecture: Conceptual View graphic>

As seen in the above graphic, Step 1 illustrates how an application makes a call to the Query Manager API to check the status of the data and request that the data for a particular timeframe be made available online (Level 1) for querying.

Once this happens, the request is passed on to an on-demand data loader (see Step 2). Now, the system can determine whether the requested data is available in long-term storage (Level 2) and whether or not a sufficient amount of space exists to bring that data online (Level 1). Once this determination is made, the data is loaded and made available online and the application is 'alerted' that the data load is complete and that the application can now perform the requested queries.

Once this step is complete, one of two things will happen: the application will make another API call to 'remove' the data that was brought online or the system will wait until scheduling attrition automatically 'removes' the data from online (Level 1) storage.

Because the Hybrid Database uses an indexing-based approach with 'tags' to store the data (i.e., a compressed file store structure provides the main indexing and a metadata structure provides alternate indices), once the data is stored on the SAN, the Hybrid Database 'looks' at the data and then collects only the data that fits the criteria by a configurable granularity. For example: CDR data would include first month, day, and switch. In this example, the CDR data can be retrieved from the directory by month, billing date, or switch and placed into the MPP. Once the data is in the MPP and the analysis takes place, the data is then automatically loaded back onto the SAN.

In instances where a telecom provider needs the same data used for separate analysis as well as real time insight into that data, the Hybrid Database offers two approaches. The first approach is to create 'two' smaller data records consisting of ONLY the elements required for the life of that data. The second, and much more proficient way, is to create an indexing scheme that creates a list of the files based on the data elements and allows the records to only be stored one time. For example: If a user wants reference data records retrieved by area code, the database index would be queried for all of the required area codes. Once that happens, the MPP is loaded with the files that match the index criteria and the user can now extract the desired information and perform the necessary analysis. (Other types of reference data that gets loaded includes: Local Exchange Routing Guide (LERG)⁴, NECCA, subscriber tables, band lists, phone types, circuit inventory, and switch and truck configuration data.)

⁴ The LERG contains information about the current network configuration and scheduled changes within the local exchange providers' networks. The LERG informs telecom companies which end

Timed Queries

Telecom providers often use queries that contain computations over large datasets and require capabilities that can move that data in an efficient and cost effective manner. With the ongoing need for retrieving and analyzing massive amounts of data, it became apparent that a more sophisticated method for handling complex queries and large data sets was needed. The Hybrid Database was specifically designed to suit these high-end applications.

<Insert Application Database Requirements graphic>

To execute timed queries for large datasets, the data must reside on the MPP. If a particular dataset is not in the MPP, it must be moved from the SAN into the MPP. Once that process is complete, the user who made the query is notified via email that the requested dataset is ready. For example: If a telecom provider has a regular monthly report showing six months of accounts in arrears, the analysis in this situation must compare billing against the previous six months of usage. In this case, the user can schedule the data to automatically load into the MPP from the SAN in advance of the query, thereby saving time and money. Once the analysis process is completed, the data can be loaded back from the MPP into the SAN, thereby freeing up valuable MPP space.

The Framework

The framework of TeleSciences Hybrid Database solution is BIAS, a Business Integration Appliance Services tool. This tool is a simple, lightweight, and effective task management system that can run, monitor, and execute computer related tasks. In addition, BIAS' user friendly interface provides a visual mechanism for the creation of simple or complex tasks, a dashboard monitoring capability with a left hand side navigational menu, as well as a centralized information logging facility. And, through the simple design of its user interface, BIAS allows business owners to define and control user access and permissions, creating non-invasive management functionality with a high level of security.

office or tandem office the NNX resides in and how calls should be routed and rated so that they can properly terminate to the appropriate telephone number at the proper rate.

<Insert chart illustrating how BIAS works with the Hybrid Database Solution>

BIAS also uses smart grids, giving business owners the capability to seamlessly manage tasks across disparate systems along a network. And, because it is designed to utilize existing processing power, BIAS takes full advantage of centralized monitoring and control.

Some of the advantages of BIAS include:

- a complete solution in one package,
- lower cost,
- open business rules engine,
- simple hardware and footprint,
- scalable open architecture, and
- an unique design that utilizes the value of MPP databases.

<Insert a chart illustrating asymmetric massively parallel processing here >

Why Your Business Needs the Hybrid Database Solution

Due to its unique architecture for data storage and processing, the TeleSciences Hybrid Database solution utilizes powerful analytical tools and delivers exponentially higher data volume processing – all at a fraction of the cost of other more traditional high volume data retention solutions.

By integrating the most advanced data collection and management tools from industry leader TeleSciences and utilizing the value of MPP, the Hybrid Database solution opens a never before offered window into enterprise performance. It also easily and intelligently support a businesses more traditional database models.

The Hybrid Database solution uses highly reliable datasets which are complete, accurate, and timely. And, its statistical and quantitative analysis features allow for explanatory and fact-based decision making while providing business owners with the ability to view the data, business processes, and results. In addition while the solution

presents business owners with answers – not just more questions – it also provides the flexibility to support a variety of very diverse requirements (application vs. application).

Based upon real-time performance data, the Hybrid Database solution allows business owners to make informed decisions efficiently and effectively. And, coupled with its scalability, lightning fast processing, and data collection and prep from any and all network elements using the most advanced mediation and analysis tools available, the Hybrid Database solution is indeed the solution that Telecom providers have been waiting for.

Leading companies throughout the world count on TeleSciences' Hybrid Database solution to analyze data, reduce costs, assure revenue, optimize facility performance, and maximize margins. For more information about the Hybrid Database solution, call ??????? or visit the TeleSciences website at www.telesciences.com.

Meeting the Requirements for Telecom Leaders Worldwide

TeleSciences provides a large range of telecom solutions to fixed, mobile, and IP network operators and equipment vendors. The company's deep expertise at providing telecom providers with fully integrated solutions that seamlessly collect data from any and all network elements with scalable business analytical tools is unparalleled.

Leading telecom providers are using TeleSciences solutions to bring enterprise-class analytical capabilities to their services-based systems. Here's a list of some of TeleSciences' customers: (List your customers here).

TeleSciences was founded in 1967 and located in Moorestown, New Jersey, USA.