

Tom Coffing Bio



Tom Coffing is one of the leading experts on Teradata, DATAlegro, Netezza, Neoview and Greenplum. Tom has written over 20 books on data warehousing, Teradata, and Netezza. Tom has taught over 1,000 Teradata classes and is considered one of the best technical speakers and writers in the industry. Tom founded Coffing Data Warehousing 15 years ago and is the Chief Executive Officer (CEO) and President.

Coffing Data Warehousing performs training, consulting, professional services, and has written almost 90% of the books on Teradata. Many people refer to Tom Coffing as Tera-Tom.

Tom Coffing's software development team has produced the Nexus, which is considered the Rosetta stone of data warehouse software. Not only is it considered the best query tool, but it has been designed to work with Teradata, IBM, Oracle, Netezza, SQL Server, DATAlegro, and Greenplum systems. The Nexus is also a DBA dream tool because it has "Point-and-Click" capabilities for compression, replicating data and DDL between different systems, comparing systems at the database, table, and data level and synchronizing the results. The Nexus is also used to build load scripts and to schedule queries and batch jobs. The Nexus has been tuned to make Teradata, Netezza, and DATAlegro co-existence a breeze.

Please feel free to contact Tom at Tom.Coffing@CoffingDW.com . For anyone wishing to download a free trial of the Nexus you can do so at the Coffing Data Warehousing website at www.CoffingDW.com.

Tom Coffing
CEO and President
Coffing Data Warehousing
Email: Tom.Coffing@CoffingDW.com
Website: www.CoffingDW.com
Direct: 937 855-4838
Cell: 513 300-0341

Teradata – Discoverers of the New World of BI

“Following the light of the sun, we left the Old World.”

– Christopher Columbus

When Christopher Columbus discovered America he couldn’t answer the simplest of questions. When he left Spain he didn’t know where he was going. When he got there he didn’t know where he was. And when he returned home he didn’t know where he’d been. Try explaining that to your boss! It definitely makes hope of future funding sail down the drain.


Teradata founders following the light of the sun worked right in their garage and founded a new computer world. In 1976, Teradata set their compass on building a database that could support a Terabyte, and they would do so following an unmapped course called Parallel Processing.

Teradata’s mission was to conquer the mainframe, which provided a centralized repository for running a company’s business by capturing their transactions. Analyzing data to make better decisions was uncharted water. Harboring feelings of resentment towards the high costs of mainframe processing, Teradata knew they could produce a faster and less expensive route to islands of data.

The Teradata concept was to take Intel chips (8080 chips) and get them to work together. Each chip would be connected to a disk. The rows of every table would be spread evenly among the disks, and then data could be retrieved in parallel.

The processing chips (called AMPs) and were each like a team of oarsman rowing a boat. If each did an equal amount of processing the system worked. Teradata continued to improve their craft knowing that there was no latitude for mistakes, because the system would either completely work together or fail.

24 years ago in 1984, Teradata’s ship came in with the release of the DBC 1012. Unlike Columbus, Teradata systems knew exactly where they were going. When they got there they knew where they were. And when they returned answer sets they knew they had just discovered Business Intelligence.



Teradata scales linearly like one Superhighway

Why Teradata Made Sense Then

“There are no speed limits on the road to excellence”

-Anonymous

A man driving down the highway got a call from his wife. She said, “The news just reported some idiot driving the wrong way down the highway. He said, “It’s not just one car, there’s hundreds of them!

Teradata built an industry with a vision of being 10 times faster and cheaper than a mainframe, but now Teradata prices have continued to head the wrong way. Teradata prices made sense 10 years ago because hardware was expensive and Teradata was the only avenue for parallel processing.

But today all **Data Warehouse vendors use parallel processing!** They all use commodity hardware, can scale to Petabytes of data, and each handles a mixed workload of queries. Teradata is facing oncoming competition and customers are changing gears to test drive newer systems in a buyer’s market.

Teradata still insists that all corporate data should reside on one single Enterprise Data Warehouse, but today’s data warehouse is far too complicated and diverse for everyone to share a single highway.

Ad hoc queries make tuning impossible, thousands of reports need to be produced each day; tactical queries compete against long-running queries. Data that is 3-months old and accessed regularly is parked next to older data that is rarely accessed. Mixes of 3rd normal form logical models access detail data which run head-on into dimensional data marts and users in completely different time zones need access to data during their rush hours. Some companies have so much traffic that they can’t even find time to complete their batch window, and when traffic suddenly stops nobody can find an exit strategy. With everything competing for CPU, memory and disk the pressure begins to take its toll.

The saddest part is that everything on Teradata Boulevard costs the same. Executives will accept slower driving if costs are low, but every mile costs as if the roads were paved with gold. Old cars cost as much as new ones and motorcycles cost the same as trucks. When you park in short term you don’t pay the same as parking in long term, so why do all users and data have to pay valet prices?

It is because Teradata has a monopoly. Trusting Teradata to reduce their prices is like expecting OPEC to drop oil prices. In time, gasoline prices will drop only when new forms of fuel are developed. We can’t currently break the OPEC monopoly, but the time has come to break the Teradata monopoly.

No country in the world uses only one giant highway. People travel through a distributed set of routes to reach a wide variety of destinations, and decide at what cost to travel. The data warehouse of the future will resemble a distributed system. Teradata customers in droves are making a U-turn, and this turn signals a whole new way of driving business value.


Netezza – Inventor of the Data Warehouse Appliance

“I cannot imagine any condition which would cause this ship to founder. Modern shipbuilding has gone beyond that.”

– E. I. Smith, Captain of the Titanic

Speeding through the treacherous waters of the fast moving data warehouse industry, Teradata hit an iceberg called Netezza, and a historical tragedy occurred. Netezza streamlined a design that provides faster performance at a much lower cost. Teradata was slow to adjust to the stern wakeup call and Netezza has been able to take a bow.

Netezza made everything so simple. They put a database, a server, and a disk module together and set sail for coordinates that maximize physical speeds for retrieving data while minimizing the shipping of unneeded data across the network channel.

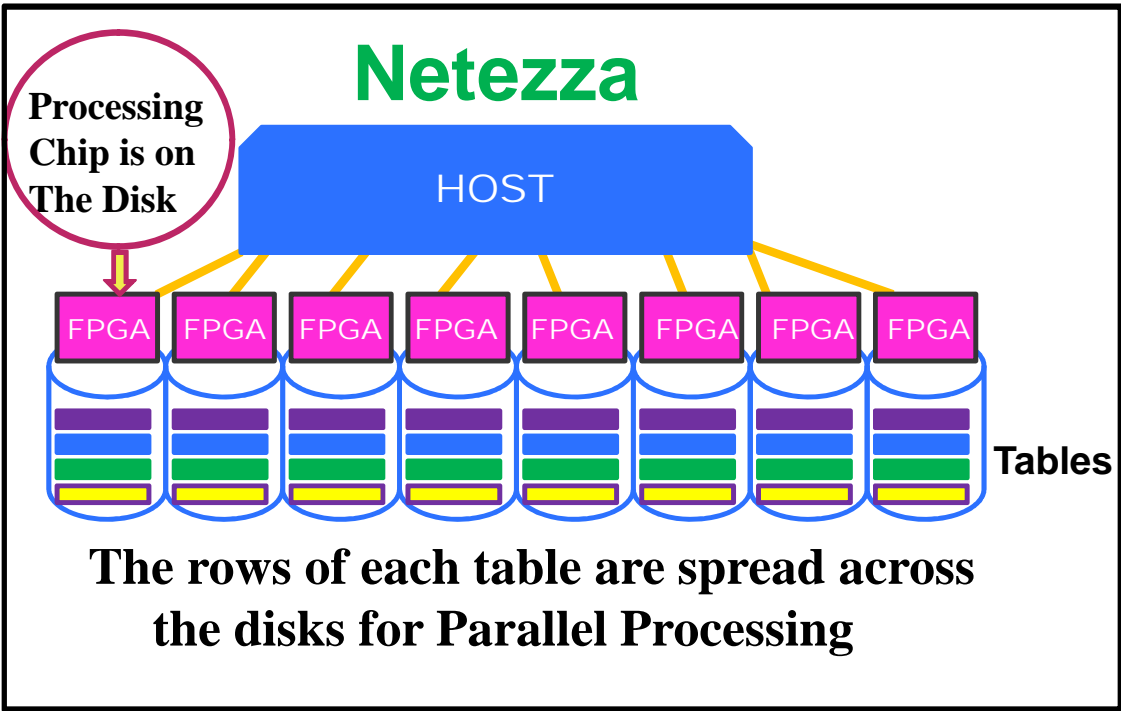


Behind the Netezza Engine

"The man on top of the mountain did not fall there."

-Longfellow

Parallel Processing is natural on Netezza because the rows of each table are spread across all the disks, just like Teradata, but the secret sauce is that Netezza put the processing chip right next to the disk. Data is analyzed at the disk level. No data moves off the disk into memory unless it is part of the answer set. This provides Netezza systems peak performance. The chip on top of the disk did not just fall there! Brilliant minds climbed the mountain, took a view at aggregation, and found a better way to "sum it".



Netezza Basic Terms

“Ambition is a dream with a V8 Engine.”

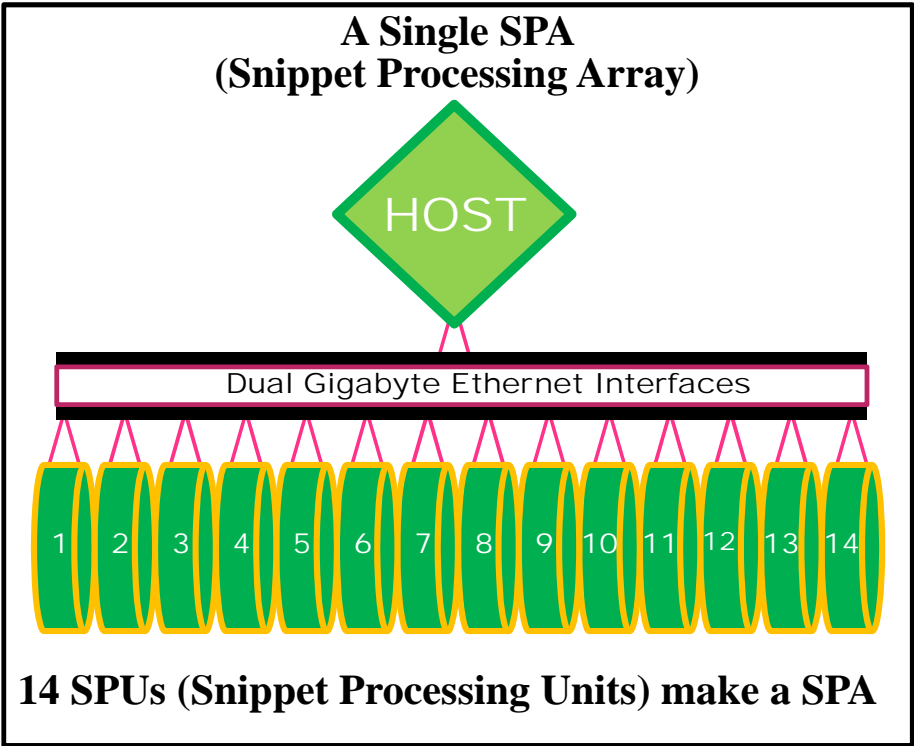
– Elvis Presley

Netezza has graced the landscape to become “the King” of all Appliances. They have the data warehouse industry all “shook up” because of their fast moves and great features. In customer tests the competition has looked like “nothing but a hound dog”, as Netezza has proved to be the faster “Retriever” of data, and their emphasis on green technology has given them fans that don’t need to scream.

For those who are familiar with Teradata systems, you know they call their parallel processors **AMPs** (Access Module Processors), but Netezza systems call their **SPUs** (Snippet Processing Units). Teradata places their AMPs inside a **Node**, while Netezza places their SPUs inside a **SPA** (Snippet Processing Array).

The Host is in charge of system administration tasks and interacts extensively with users. It performs much like the Parsing Engines (PE) Optimizer on Teradata.

Two Gigabyte Ethernet Interfaces connects the Host to the actual data, and is similar to Teradata’s BYNET.



The Teradata Architecture

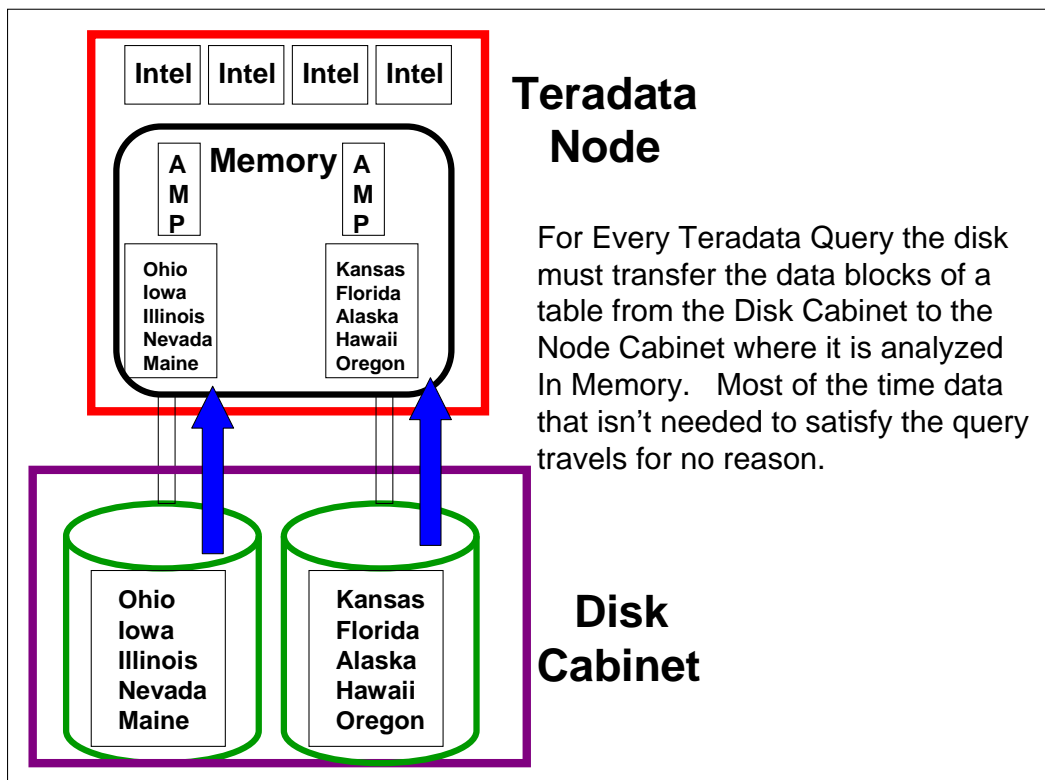
“Kites rise highest against the wind – not with it.”

– Sir Winston Churchill

Teradata was the true pioneer of parallel processing, but Netezza saw the “blitz” of data traveling through Teradata systems and decided that the disk battle would be their finest hour. The biggest difference between Teradata and Netezza is that Teradata moves a lot of data that won’t satisfy the query and Netezza does not.

In the picture below, if a Teradata user asks for the rows of Ohio and Oregon to be returned, then each AMP sends a request to a **separate disk cabinet** to find the State table. The disk finds the data block and sends it across the network to the Node’s memory. Every state in the block travels from the disk cabinet to the node cabinet where it is analyzed to see if it matches Ohio or Oregon. It is a waste of time and travel for 90% of the data.

This concept is like moving the contents of your entire refrigerator into the dining room just to eat a small snack.



The Netezza Secret Sauce

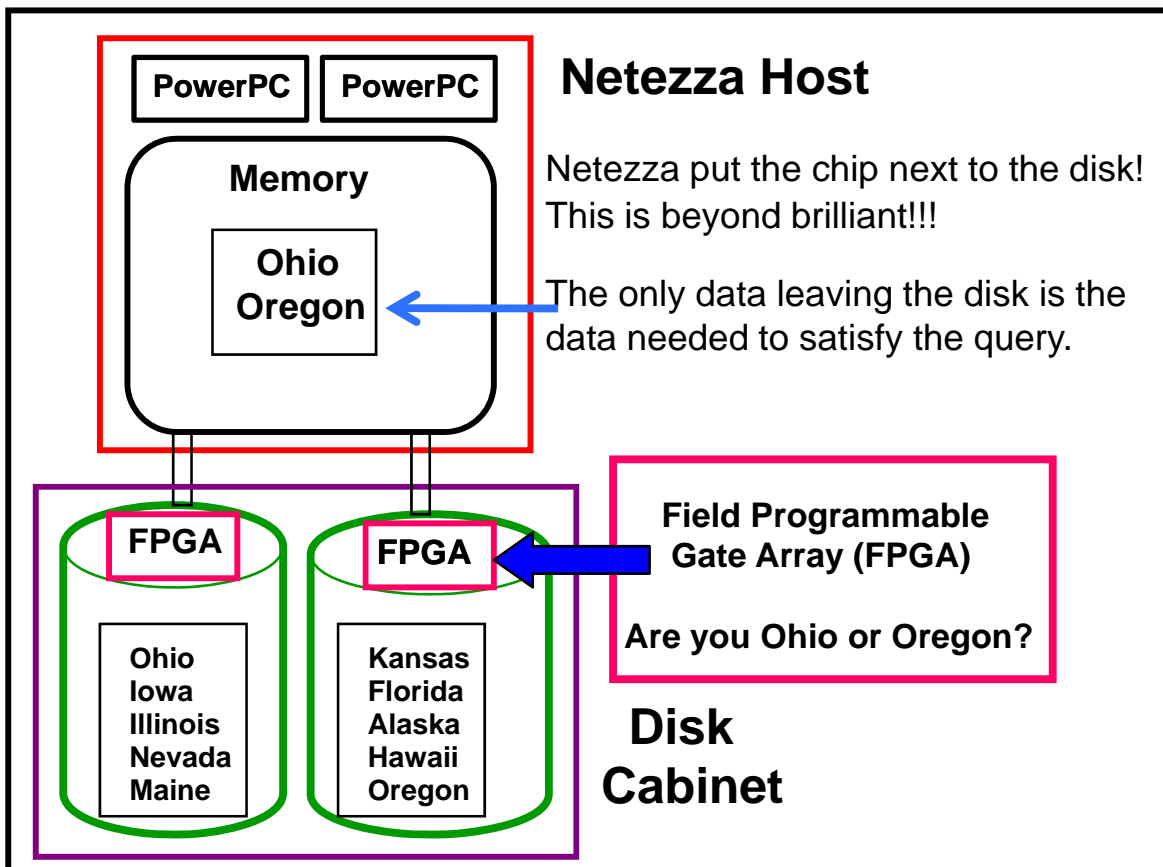
“A bird does not sing because it has the answers, it sings because it has a song.”

-Anonymous

Netezza doesn't sing because it has a song! It sings because it has the answers – and can fly! By putting the chip next to the disk it gave their system wings. Data never moves into memory unless it is part of the answer set. Does a bird leave the nest for no apparent reason? So why should data? The secret chip is called a Field Programmable Gate Array (FPGA).

In the diagram below when a user requests to see the data from the State Table and only wants to see Ohio and Oregon, then the words Ohio and Oregon are written on the FPGA chip. No data gets by the FPGA chip unless it is from the states of Ohio or Oregon.

This is like taking only the food you are going to eat and moving it to the dining room.



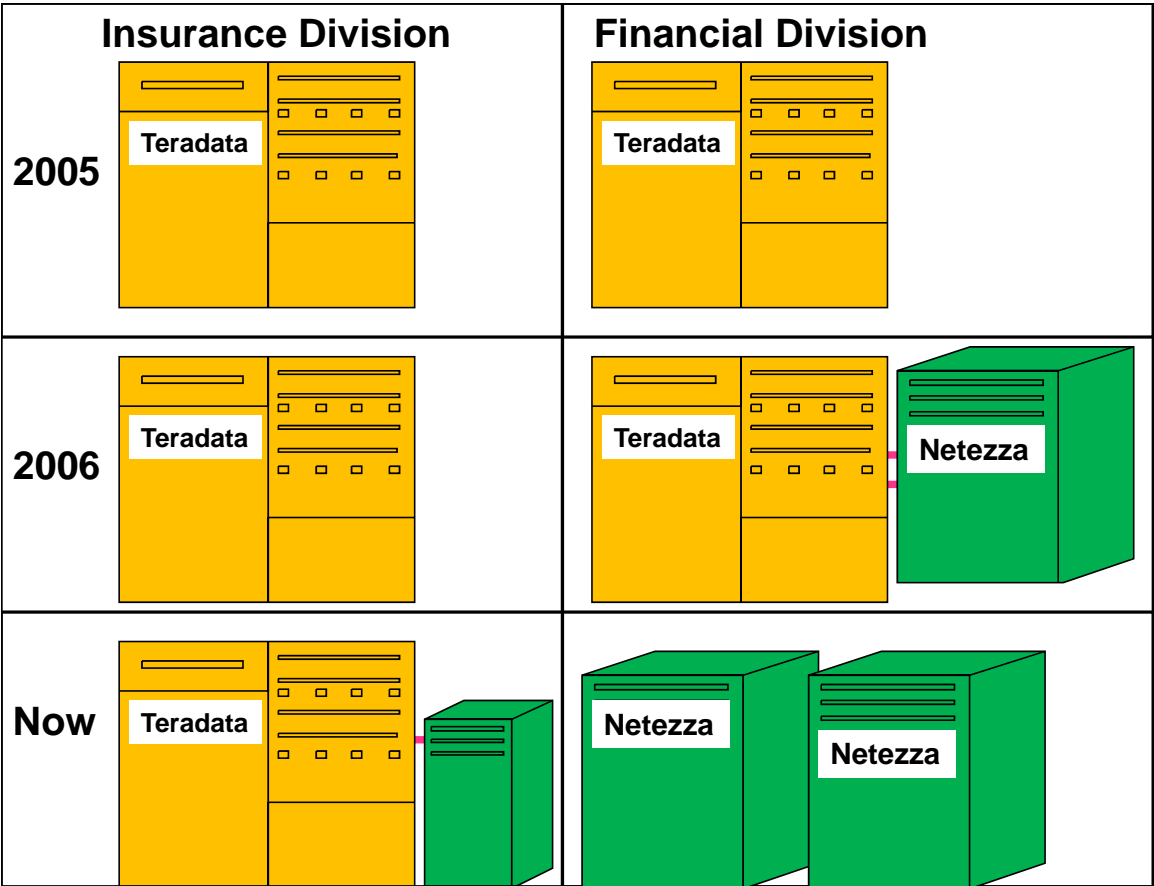
Case Study: A Major Insurance and Financial Company

“I don’t know who my grandfather was. I am more interested in who his grandson will become.”

– Abraham Lincoln, 16th president of the United States

A case study about a major insurance and financial company proved that a house divided truly can stand. A major insurer drew the Mason-Dixon lines of business between their Insurance and Financial divisions. Each division had separate Teradata warehouses in 2005, but a civil war emerged in the financial warehouse when a Netezza sales rep was given a direct order.

The Financial Division only started with a small Netezza system, but after Netezza proved itself battle ready, the division completely replaced Teradata with a Netezza system. Now, The Insurance Division is contemplating a small Netezza box in order to perform a better and faster UNION!



Teradata and Netezza Compete Head-to-Head

“When you are courting a nice girl an hour seems like a second. When you sit on a red-hot cinder a second seems like an hour. That’s relativity.”

– Albert Einstein

It doesn’t take an Einstein to understand that the theory of relativity is that Netezza systems are relatively faster and relatively less expensive.

$\Sigma=MC_2$ has now become **E**xpectations = **M**inimal **C**ost with **2** times the performance.

In a recent statement to Teradata employees, senior management provided some 2007 “Good News” and “Bad News”. The good news is that when Teradata competed against IBM they won 87% of the time! The bad news is that when Teradata competed against Netezza (in a warehouse environment) they lost 50% of the time. Some believe that Teradata qualified their losses with the words “when competing in a warehouse environment” to hedge the fact that they really lost close to 80% of the time.

Netezza has sold over 150 systems to some of the largest customers in the world and almost 20% of their systems have been sold to customers who currently have Teradata.

I had a customer call me who was looking to purchase a Teradata warehouse. They asked if I would do their Teradata training. I said “yes”. They said that they were definitely buying Teradata, but they were doing a Proof of Concept (POC) against Netezza to negotiate a better price. They recently called and asked me if I could do their Netezza training. I said, “I thought you were set on buying Teradata”. They said they were surprised, “but after running the POC that Netezza was most likely to win”. They were surprised at how fast Netezza ran a mixed workload of queries and the costs were phenomenal.

2007 Results for head-to-head competition

When **Teradata** competed against **IBM** they won 85% of the time

When **Teradata** competed against **Netezza** they lost over half

Why Teradata Customers are moving to Netezza

“When you go into court you are putting your fate into the hands of twelve people who weren’t smart enough to get out of jury duty.”

- Norm Crosby, Comedian

Teradata has always made a case for one Enterprise Data Warehouse in order to have one version of the truth, the whole truth, and nothing but the truth. Although there are benefits it is hard to defend the fact that every Teradata query costs the same. Would you pay a lawyer top dollar for a traffic ticket? Would you pay a lawyer double if you had evidence they were going to be outperformed? Object to the motion that theorizes that because Netezza is half the price they are half as good.

Many companies are making their next trial a “Proof of Concept” on Netezza. When the verdicts delivered and Netezza is found to perform 10-100 X faster on certain applications, and costs half, management will be pleading for a system.

In the past, companies brought in Netezza to drop their Teradata costs when negotiating a Teradata upgrade. Now, companies are bringing in Netezza because of lower prices and better performance.

A company with Teradata and Netezza in a co-existence strategy can save costs, improve performance, relieve the burden on their Teradata system, and allow for two equal sides to compete for the business. “Due Process” should really mean “Do process” data where it makes the most business sense for your client, on your server of choice.

In the past, Teradata customers testified as witnesses to Teradata’s superiority mainly because Teradata was the only parallel processing game in town. Teradata acted as judge, jury, and executioner of all queries, but now there is real parallel processing competition. Don’t put your fate in the hands of any one vendor or sentence your employees to heavy fines for each count, sum or join they run. The jury is still out on which system best suits your company long-term, so consider getting them both.

Teradata and Netezza combined		
Faster Processing	Versatility	Less burden on one system
Cost Savings	Competition	Easier to meet Batch Window

Why Netezza is Ready for the Big Leagues

“It’s déjà vu all over again!”

-Yogi Berra, Hall of Fame Catcher

For the past seven years Netezza has continued to fine tune its game. Netezza found their first base of success with customers whose applications used Full Table Scans and aggregation. Placing the chip next to the disk gave them a major advantage. It’s like a baseball Pitcher who gets to throw much closer to the plate while the competition throws from a further distance. Netezza has an unfair advantage.

Teradata continued to balk at Netezza, placing them at the bottom of their competitor list, but Netezza moved to second when they improved their ability to perform Tactical queries. While Teradata was able to perform these short sub-second queries with indexing, Netezza was able to do it with Zone Maps. With a new ability to handle a mixed workload, Netezza became a deep threat to Teradata.

Because of the FPGA chip next to the disk and the use of Zone Maps, Netezza has provided a delivery of systems that need no warm-up or tuning. Just load and go and you are guaranteed the system will fit like a glove inside almost any environment.

Netezza has always had a game plan to provide low cost, highly effective data warehouse Business Intelligence to customers. Their low cost, mixed workload capabilities and no tuning required approach has been like a called third strike against Teradata.

Netezza systems are fast now, but on pace to become six times faster over the course of the next year. This is because Netezza is utilizing 700,000 gates of the 1.4 million FPGA chip gates. Netezza is working to improve their ability to use more FPGA gates and has the potential to improve their already fantastic statistics. Perhaps this is why Netezza systems are always “outstanding in their field”.



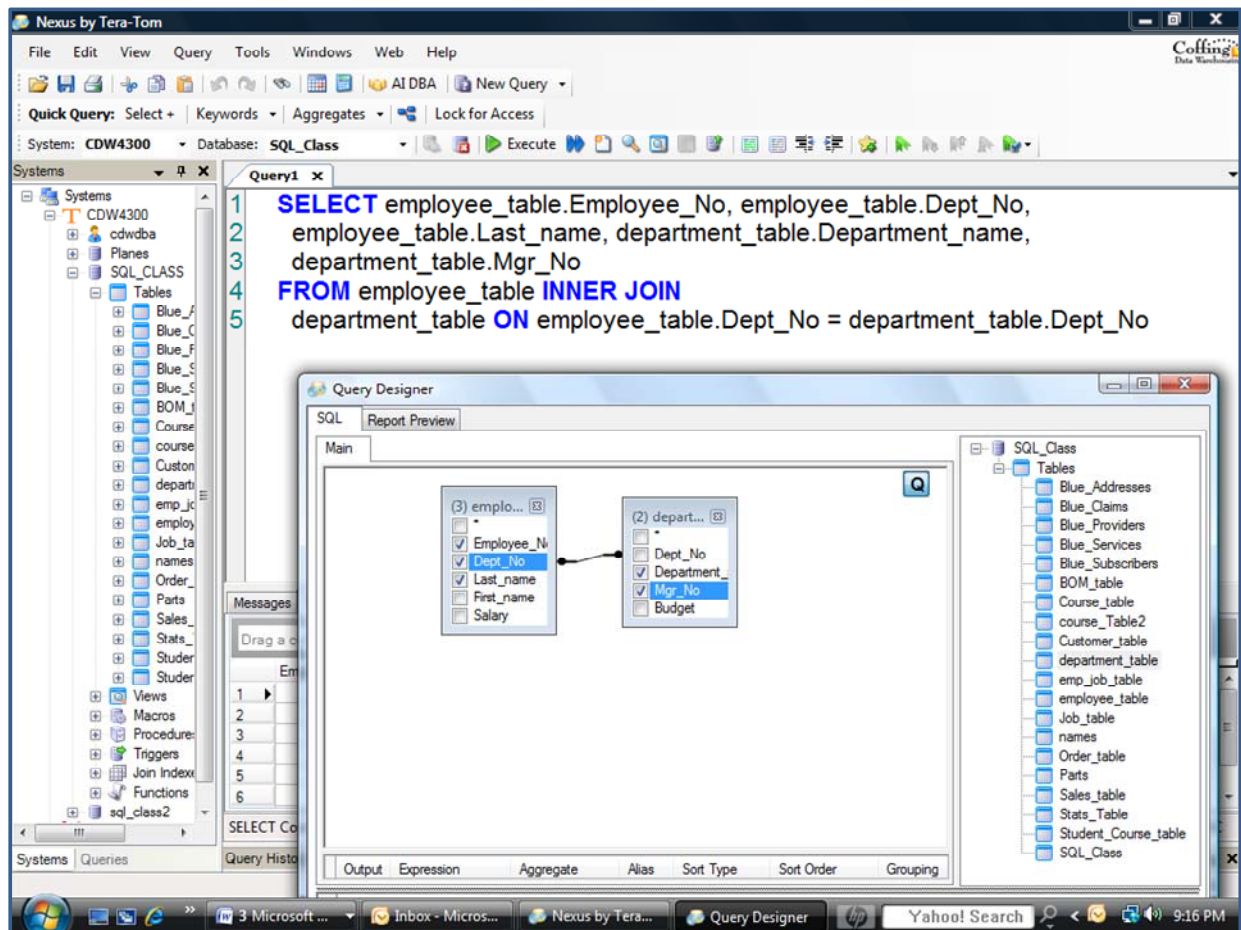
Nexus – a Co-existence Tool for Teradata and Netezza

Some companies worry about a multi-vendor environment because of the tools. Worry no longer! The Nexus is the tool of choice to bring both Netezza and Teradata together.

The largest Teradata customers in Australia, the Middle East, USSR, Europe, and the United States have purchased the Nexus. Some of the largest Netezza customers in Canada, Europe and the United States have also purchased the Nexus. Almost all customers who have a co-existence strategy with Teradata and Netezza have purchased the Nexus.

It is hard to imagine, but the Nexus can query Teradata, Netezza, Oracle, DB2, SQL Server and every other appliance simultaneously. The Nexus even has a visual Query and Join Builder that anyone can master. Plus, Teradata DBAs love Nexus because of its ability to do “Point-and-Click” Compression, Replication, Synchronization, Scripting, Metadata, and Scheduling. Why have ten tools for ten different vendors when you can have the best tool for all.

If you currently have a Teradata system and would like to try Netezza then I have great news. The Nexus will take Teradata tables (DDL) and automatically convert them and load to Netezza tables. Nexus will also handle Teradata and Netezza DBA functions. Download a free Nexus trial at www.CoffingDW.com.



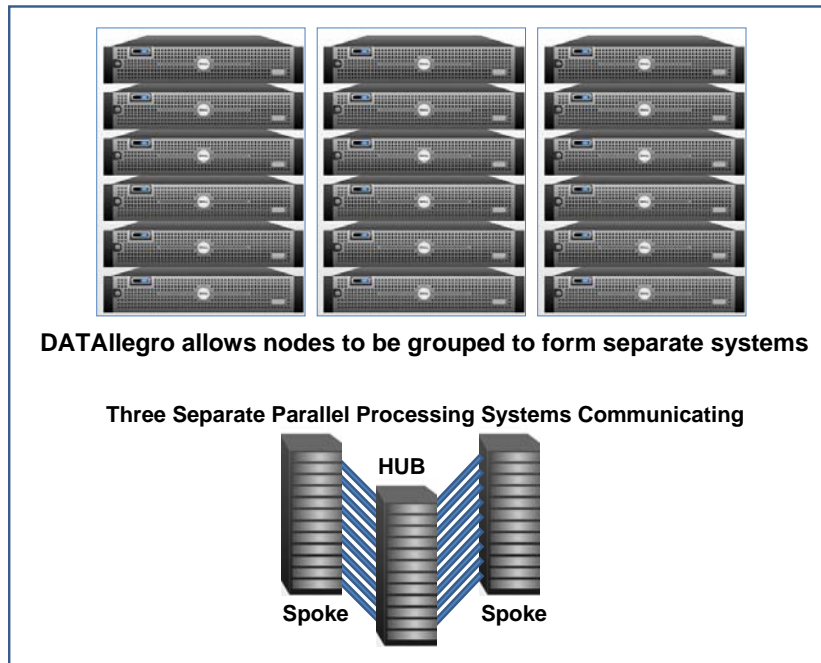
DATALlegro – A Distributed Processing Architecture

“An invasion of Armies can be resisted, but not an idea whose time has come.”

– Victor Hugo

Not paying an Army and a Leggy for a data warehouse is an idea whose time has come. Utilizing Dell Servers, Cisco InfiniBand Routers, EMC Disks, and a proven open source database in Ingres, DATALlegro has brought world-class commodity hardware and software together, and the invasion has begun.

Stuart Frost, CEO and Founder of DATALlegro, has built an architecture that is going to be a major force to be reckoned with. DATALlegro systems are uniformly physically connected, but act logically as a set of distributed systems. In general, this type of operation provides a true “Hub and Spoke” architecture that combines speed, power and flexibility with a price point that deserves four stars.



DATALlegro - The Best of Both Worlds

“A man who views the world at 50 the same as he did at 20 has wasted 30 years of his life.”

– Muhammad Ali

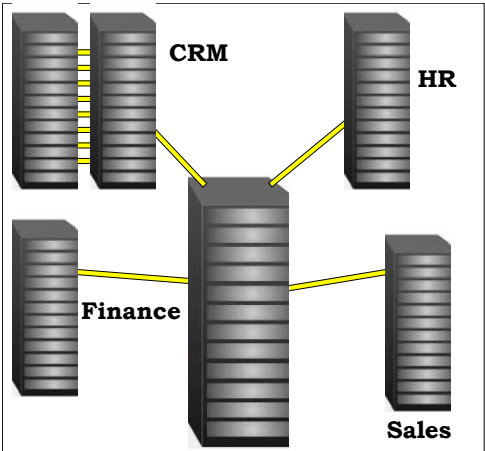
Nobody “spoke” better than Muhammad Ali, and nobody has better “Spokes” than DATALlegro. If playing with Lego blocks as a kid rings a bell, then designing your data warehouse using Dell Server blocks has to be ranked as the greatest of concepts. Dell Servers are laced together to form grids that allow the creation of a Hub and Spokes that are guaranteed to fit like a glove in any environment.

Consider the power of querying detail data combined with the speed of implementing new data marts. Beating the competition to the punch is what data warehouses are all about. DATALlegro has undercut the competition by promoting their distributed system, which offers much better potential than an expensive centralized platform.

As no champion boxer can expect to win without a great jab and power punch, no data warehouse can expect to compete without access to detail data and data marts. 80% of queries can be satisfied by a data mart, but 80% of the return on investment comes from querying and data mining detail data. A distributed system versus a centralized system is an easy decision.

Data marts can be placed in logical spokes, or you can even consider having detail data for a region, time zone or continent on spokes that make logistical sense. You might even take all product data and place it separately from data that doesn’t have a relation. This can all be done without duplicating data.

Muhammad Ali and DATALlegro are both known for delivering great combinations.



Architectural Difference between Teradata and DATAlegro

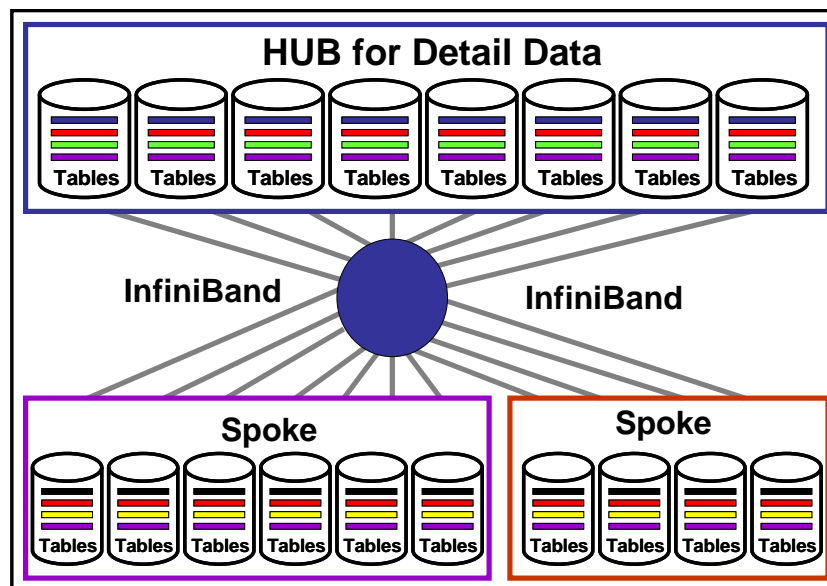
*“I saw the angel in the marble and carved until
I set him free.”*

--Michelangelo

DATAlegro has painted a portrait of an entirely new data warehouse landscape. This striking pose was captured by using a distributed model. DATAlegro began sculpting their masterpiece in 2003, and versions I and II taught them many lessons, but the silicon ceiling was raised when they produced a work of art with DATAlegro V3. Rome wasn't built in a day.

The brilliance behind DATAlegro is that parallel processing is used in every part of the system. They parallel process detail data in a hub, and simultaneously parallel process data marts in spokes. Distributed parallel processing communicates like one big system, but spokes in a sense are their own separate system. It's like Michelangelo painting different parts of the Sistine Chapel with great detail and colors, but tying the entire masterpiece together as one entity.

Teradata systems have been very successful, but always act as a centralized single system. The needs of the enterprise are too diverse to realistically allow a single system to process everything. Why process newer data the same way as older data? The Sistine Chapel wasn't painted with a roller. Allowing users to access their portion of a distributed system as if it were their own separate system will allow users to put the business in the best light.



Multi-Temperature Data Warehousing

“If I have seen farther than others, it is because I was standing on the shoulders of giants”


– Isaac Newton

The cost of storage and queries always costs the same price on Teradata. Understanding the gravity of the situation, DATAlegro built a system where the **price per query** falls quickly, and drops even more in their Multi-Temperature data warehouse.

A Multi-Temperature data warehouse allows customers to decide at what price point they want to pay for storing and querying certain tables and data.

Most companies keep data for years, but the latest 3-months is accessed much more heavily. It makes sense to query the most recent 3-months of data on your fastest and most expensive system, but be able to pay much lower costs on data rarely accessed. DATAlegro provides both higher and lower performance at higher and lower costs for different spokes in the warehouse.

Companies can keep paying the same price for every Teradata query or use DATAlegro to experience an equal but opposite reaction in price.



Combining your Teradata System with DATAlegro

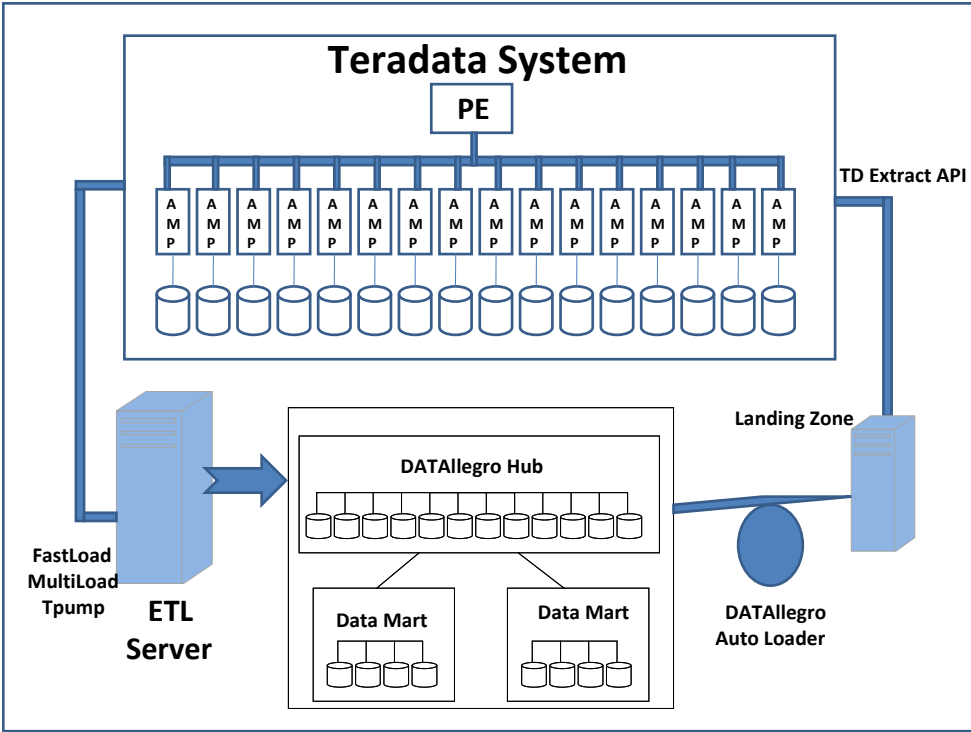
“First you imitate, and then you innovate.”

- Miles Davis – Jazz Musician

Although nobody expects Teradata customers to replace their current investment immediately, many companies like the sound of a Teradata and DATAlegro duet. Customers have taken note of the fact that DATAlegro reduces costs immediately, but has the potential to completely take the lead.

Allegro is a musical term meaning “A direction to play lively and fast”. DATAlegro is a data warehouse term that strikes a chord for fast speeds on large amounts of data. DATAlegro accomplishes this with commodity hardware and incredible networking between the Hub and Spokes. DATAlegro is getting the InfiniBand back together again, as this Cisco network standard links Intel, Dell, and EMC hardware for a price performance that is breaking records.

The Hub and Spoke architecture allows companies to orchestrate where they need to place their power and investment, so adding DATAlegro to your current Teradata environment will relieve the stress on your Teradata system, while setting a tone for real competition. Having additional options and cutting costs is the reason that everyone is getting so jazzed about the future of data warehousing.



Three Areas to Consider DATAlegro Immediately

“A Journey of a thousand miles begins with a single step”

- Chinese Proverb

A journey of a single step can deliver a Teradata/DATAlegro co-existence strategy. This journey isn't long, but the views are great. Consider multiple areas that can have impact quickly. Start with a project that will provide high visibility to management and can impact the business financially. This makes the short trip well worth the effort.

A first step is to identify if there is a department or business unit that has a budget, but cannot be serviced by IT? In less than 12 weeks you can identify the requirements, and users can be ready for business. DATAlegro is easy to setup so 12 weeks is generally the estimated time for the journey to reach its first destination.

The second step is to examine migrating historical data off of Teradata? Why pay high prices for data you rarely access? This data costs the same on Teradata, but Multi-Temperature implementations on DATAlegro allow for low price points for storing and querying data that is rarely accessed. Migrating older data from Teradata will also free up valuable space, and make travel on your Teradata system uncluttered. Plus, it will save you an expensive Teradata upgrade.

The third step is to take a strong look at intensive batch processes such as aggregations and data marts. Placing this on a DATAlegro system is a great idea. With every user currently competing for space and time on your Teradata system, users will find it refreshing to query a DATAlegro with minimal traffic.

Understanding what data should be processed on Teradata and what data should be processed on DATAlegro is the key to process improvements and time to market. It is the path of least resistance.

The beauty of a DATAlegro system is that it can act as a small appliance, but can also compete directly with the largest of the traditional data warehouse players.

Converting Teradata to DATAlegro

“Every sunrise is a second chance.”

– Unknown

There is a saying, “Every morning in Africa a Gazelle wakes up and knows if it is to survive it must outrun the fastest lion. Every morning in Africa a lion wakes up and knows if it is to survive, it must out run the slowest gazelle. Either way when the sun comes up in Africa you hit the ground running”.

There are many reasons to convert a centralized system to a distributed solution, and most of these revolve around speed and survival. Is there only one watering hole in Africa? Africa works best when the rainy season comes alive because water is distributed in tributaries everywhere.

A company who has one vendor providing a centralized solution will be hunted down by competitors in the future. You can't outrun the costs or afford to lie in wait as your vendor brings out new technology.

When a centralized solutions resource begin to dry up, the system becomes slow, and every user fights to merely get a drink of water, until management pours out a large amount of money. A distributed system allows users to get their fill of answers.

DATAlegro makes converting from Teradata faster than a “Cheetah” with utilities to move data and DDL directly. DATAlegro can take data off of Teradata in binary form and load it faster than it was originally loaded into Teradata. Data hits the hub and is distributed to the spokes like rainfall filling up tributaries. These utilities are included with the implementation package supplied by DATAlegro.

Break the Monopoly Game with a Proof of Concept (POC)

The biggest reason that a single centralized solution is so high is because customers have allowed them to gain a monopoly over their data warehouse. The game of monopoly is played with the concept of controlling all the money and property, and the same holds true for some data warehouse vendors.

Here is how the game of Monopoly is played in data warehousing.

The data warehouse salesperson starts the game by rolling into your office, and tells you that their new hardware is twice as fast, and that you need an upgrade.

Your next move is to try and balance yourself as they reveal the price fully knowing that your users are demanding more speed and that you don't have any options but to pay the community chest.

There next move is to agree to drop the cost a little, but only if you agree to purchase all of your consulting, education, tools and professional services as part of the contract. These services then come in the form of pre-paid credits. They convince you that they are the only game in town and that you will fail if you take a CHANCE with another company.

Game over – they win! They now own all your properties by controlling the hardware, software, education, consulting, maintenance and professional services. They now set new rules.

There is no "Get out of Jail Free" card! Teradata always passes GO and collect \$200 to \$500 (an hour) for every role. And you must pay triple for every utility.

Why pay millions for a system and then pay an enormous luxury tax?

The brilliance behind their Monopoly game is that upper management negotiates the contract and middle management gets the services. One manager said it best as, "The services are free to me because it doesn't come from my budget". Just like Monopoly nobody is playing with real money anymore, and the vendor has just become your banker.

Don't play Monopoly. Ask your vendor to play Trivial Pursuit.

First Question: What are the advantages of a Distributed system over a Centralized solution?

Answer: Speed, price, upgrade capabilities, multi-temperature, flexibility, logistics and competition.

Then do a Proof of Concept (POC) with DATAlegro. What have you got to lose? By utilizing a competitive platform you are in a position to negotiate with power. With help about a POC please contact me or download my POC Whitepaper for instructions. It takes less than a week!

Then tell your Teradata sales rep you want to play a new game. "The price is right"!

Summary

“Write a wise saying and your name will live forever.”

- Anonymous

Teradata was designed back in the 1980's and deserves credit for being the market leader. Their greatest contribution has always been parallel processing. For almost three decades they have provided a centralized solution to many of the largest companies in the world.

Today's data warehouse is much more complicated than ever before. Geographic locations in different time zones need to access the warehouse at their peak times. Ad hoc queries make it difficult for tuning, and enormous pressure is placed on the warehouse to produce thousands of reports. Couple this with the fact that some departments need to access data that is 3-months old while others need to access data that is 3-years old. There is also a mixed variety of logical models such as 3rd Normal Form and Dimensional models residing on the same platform, and a more diverse mix of queries, users, and applications than ever before. All are competing for CPU, memory, and disk. It just doesn't make logistical sense to try and force it all onto one centralized platform.

DATALlegro is a great solution to utilize in conjunction with your Teradata system. DATALlegro utilizes parallel processing, fast load and backup utilities, and a Hub and Spoke architecture. Utilizing Dell Servers, Cisco Routers, EMC Disks, and a proven open source database in Ingres, DATALlegro has brought world-class commodity hardware and software together, delivering incredible speed and performance, while dropping costs dramatically.

Netezza is the king of the Appliances and they have sold to over 150 customers worldwide. Their patented streaming architecture places the chip next to the disk so only the data needed to satisfy a query is transferred to the host for processing. Their ability to load and go without any tuning makes them perfect for any data mart, but the scalability of Netezza has allowed them to put massive data warehouses in to sites like New York Stock Exchange, Catalina Marketing, and Amazon.com.

A co-existence strategy between Teradata, Netezza and DATALlegro is allowing companies to have the best of both worlds. Please don't hesitate to contact me if you need additional information or help with a POC.

Sincerely,

Tom