The Teradata Data Warehouse Appliance

Technical Note on Teradata Data Warehouse Appliance vs. Oracle Exadata

By Krish Krishnan

Sixth Sense Advisors Inc.

November 2013

Sponsored by:

TERADATA

Table of Contents

Background

This technical note is an Independent Analyst and data warehouse practitioner's viewpoint on the Teradata Data Warehouse Appliance vs. Oracle Exadata. The information contained in this note is derived from various sources of input, including the author's own implementation experiences, interactions with clients and other data warehouse professionals, discussions about the subject at various conferences, and public information available from the respective vendors.

The goal of this note is to provide a concise comparison of the data warehouse capabilities of the Teradata Data Warehouse Appliance and Oracle Exadata platforms.

Starting at the beginning, the primary goal of a data warehouse is to provide a consolidated enterprise version of the data, which is cleansed, integrated and available for query and reporting. The data warehouse has grown to be a key asset for every organization, and today the data warehouse and the business intelligence (BI) applications that are deployed on data warehouse platforms are critical to business decision making. As such, they now constitute a strategic part of the overall IT budget in most enterprises.

The biggest threat in the world of BI and data warehousing is a five-letter word – S-P-E-E-D. While BI applications often demand a quick response time and near real-time application availability, the underlying infrastructure of a traditional data warehouse becomes expensive as it scales up and out. C-level executives constantly question the cost and value of this asset – its total cost of ownership and the return on investment it supports.

Due to the continuous advances in technology, the cost equation for infrastructure to support high-performance data warehouses has come down considerably – by about 70% of what one might have spent a decade ago. The data warehouse appliance has emerged as the answer to a cheaper, faster and better data warehouse platform. There has been significant growth of data warehouse appliance vendors and equally faster consolidation of the vendor space over the past few years through mergers and acquisitions. The market has consolidated with key leaders Teradata Corporation, Oracle, and Netezza (now IBM) in the space. Nonetheless, ongoing confusion about the scalability and functionality of data warehouse appliances continues to exist, especially regarding Exadata and Teradata.

Several key factors driving data warehouse performance that continue to challenge IT departments are data loading, availability, data volume, storage performance, scalability, diverse and changing query demands against the data, and operational costs of the data warehouse system.

Loading Data

Loading data into a data warehouse is often time consuming. To summarize, the process of extracting data from various sources, processing them through data quality procedures, and then data profiling and loading them with or without transformations to a final destination may exceed available load windows. This can be especially challenging when low input volumes alternate with short, high-volume bursts of data. Finally the sheer volume of data in the data warehouse affects speed of adding additional data.

Data Availability

Timeliness of data availability has a profound impact on the need for a high-performance data warehouse environment. End-user requirements must be clearly documented for data to be pristine, integrated, and seamlessly available for downstream applications like reporting and analytics. Additionally, organizations often neglect or underestimate data growth projections, data demand projections, data retention cycles and associated query response SLAs. The demands on the data warehouse system to support larger data volumes, greater numbers of users, and a more complex and diverse workload can easily outpace the existing infrastructure.

Data Volumes

Data volumes in a data warehouse have exploded by gigabytes and even terabytes on a daily basis. Data sources have gone from traditional systems to social media, mobile devices, machine learning and much more. Demand for capturing and retaining more granular details has been increasing over the past four years, driving up data growth and storage levels. A few reasons for this data volume explosion include:

- Compliance requirements such as SAFE Harbor, SOX, HIPAA, GLBA, and PCI regulations have mandated that structured data must be retained online long enough to meet compliance requirements.
- Legal mandates: Recent lawsuits in pharmaceuticals, manufacturing and the tobacco industries require data retention for longer periods of time to assist in legal situations.
- Business users: Business leaders are seeing the value of information gathered from exercises in data mining and exploration as well as historical analysis.
- Predictive intelligence: This is a major reason for storing data longer, to support exploration warehousing and heuristics.
- Just-in-Case scenarios: While there is no clearly defined requirement for such hypothetical cases, the volume of data retained for such purposes continues to grow. This will eventually cause a highly adverse impact on the performance and response time problems for the data warehouse system as organizations keep their data for longer periods of time.

Storage Performance

Disk and storage systems have consistently increased available storage volumes over the years while keeping costs relatively stable, but disk performance has not kept pace with CPU performance. Architecturally, sharing storage across all areas of a data warehouse infrastructure constrains availability and performance of the I/O subsystem. ETL and BI queries produce large amounts of traffic consuming a lot of space and network bandwidth. If multiple queries are accessing this shared storage architecture, even the best-in-class hardware and disk are not going to enable faster query processing or a lightning response time for producing a result set. On top of this, if we add mixed query workloads on the storage architecture, we are going to start seeing extended access times due to I/O bandwidth limitations, resulting in poor query performance. Even continued strides in improving the overall storage performance won't make it more optimal for data warehousing. Faster is better is not sufficient in this space.

Operational Costs

In many organizations, the cost of operating and maintaining a data warehouse has been among the largest IT budget items. The increased granularity of data available and the need to keep historical data longer, is causing a two-way explosion, resulting in an unmanageable amount of information being processed by the data warehouse. In addition to the exploding data volumes, the warehouse also needs to support a growing volume and diversity of work against the data, including related activities like data mining, predictive and heuristics analysis or operational tasks. A direct by-product of these workloads is the heavy demand it puts on IT organizations administration resources (DBA, System Administrator, Network Administrator roles). The overall cost of running and maintaining the data warehouse is daunting and has frequently overwhelmed IT organizations.

These are just a few of the various pain points, issues, and challenges that are faced in managing and maintaining a "high-performance" data warehouse. Users continue to ask, "Is there a remedy in sight? Can we have our cake and eat it too? Are the types of responses or looks that we have been seeing on the faces of the people that have been running these data warehouses justified?"

The Data Warehouse Appliance Selection

In a nutshell, a self-contained integrated solution stack of hardware, operating system, RDBMS software and storage, optimized for data warehouse workloads is what the appliance is all about. The appliance has a self-managing, self-tuning, plug-and-play architecture that can be scaled out in a modular and cost-effective manner. The predominant architecture of the data warehouse appliance is commonly referred to as massively parallel processing (MPP) architecture. This is the key architecture that all leading vendors claim to provide, but do not be fooled by what everybody claims, as all appliances are not equal. What differentiates the scalability and performance of the appliance? Many factors need to be considered. This document will examine two platforms: Teradata and Exadata.

Teradata or Exadata? The biggest question that lingers in the minds of business and IT leaders is which platform should an organization build their data warehouse and business intelligence initiatives on? The consistent theme that the user community has expressed is the need to maintain performance by lowering response times, while not exacerbating implementation and ongoing maintenance costs.

Under the Hood

Let's examine the appliances under the hood, though it is comparing apples and oranges, the reader will benefit from this comparison.

Oracle Exadata

Oracle Exadata is a packaged solution offering from Oracle, configured with bundled hardware, storage and database, which is touted to be optimized for handling scalable data warehouse-type workloads in query and analysis. Oracle has repeatedly claimed that the Exadata X3 system has significantly improved performance and scalability system-wide, resulting in scalability up to petabytes and beyond.

There are actually two Oracle Exadata Database Machine products: Exadata X3-2 and Exadata X3-8. The X3-8 is a full-rack only system, primarily intended for large OLTP and consolidation environments. Both platforms consist of the latest release of the Oracle Database, Oracle RAC (Real Application Clusters) Database server grid, an InfiniBand interconnect, the Oracle Enterprise Linux operating system, and the Exadata Storage Server Grid using either high-performance (600GB) or high capacity (3TB) disk storage.

When a query executes, the data stored in the Exadata Storage Server grid and the storage servers act as a sort of pre-processor for accessing data from disk in an optimized fashion, using what Oracle calls Smart Scan before passing the results to the database itself. This can significantly reduce the amount of data that the database has to process and can potentially benefit users in a non-intrusive data warehouse environment. To improve performance in OLTP environments, Oracle Exadata also includes flash storage for caching hot data across all configurations. Let's examine the performance accelerators that Oracle claims to have developed for Exadata.

Hardware Improvements – Oracle in my opinion has thrown a lot of Iron to try to increase data warehouse performance.

Look at these numbers – to achieve desired performance levels this is the hardware platform that one will purchase. But what dismays many users is that the Oracle RDBMS is still the primary gateway to Exadata. The hardware constraints that exist due to Oracle base performance issues still remain.

The net result – if you are spending millions of dollars to buy the Exadata system (excluding support costs), what you will get is a lot of hardware that attempts to mask deficiencies that exist in the database, unsuccessfully in many cases. To configure and optimize the database, customers are still faced with the complexities of Oracle.

This brings up the next important set of features that Oracle claims differentiates Exadata.

Exadata Intelligent Storage

Exadata has multiple layers of storage designed to improve I/O performance.

- **Disk** – Exadata Cells or Exadata Storage provide a classic shared disk architecture that consists of arrays of disks that are managed by Oracle automatic storage management (ASM). Data is striped by Oracle's Stripe and Mirror Everything (SAME) storage allocation policy, which distributes data across multiple Exadata storage cells. As a consequence of SAME

Oracle Database Server Grid

- 8 compute servers
- 128 Intel Cores
- 2 TB DRAM

InfiniBand Network

- 40Gb/sec unified server and storage network
- Fault Tolerant

Exadata Storage Server Grid

- 14 storage servers
- 100TB raw SAS disk storage
  or
  504TB raw
  SAS disk storage
- 22.4 TB of flash storage

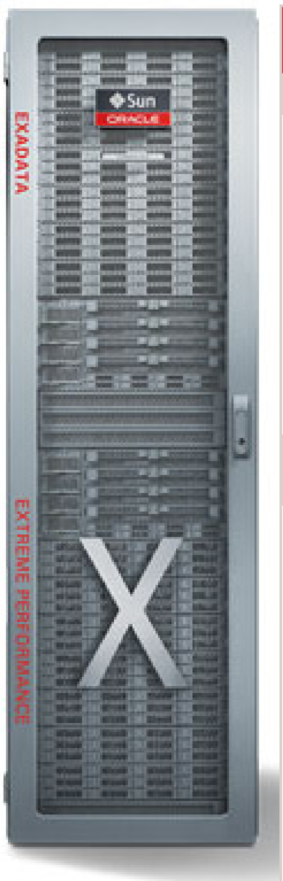

Exadata Database Machine X3-2 Full Rack	Exadata Database Machine X3-8 Hardware
8 x Database Servers	**2 x Database Servers**
Each with: – 2 x Eight-Core Intel® Xeon® E5-2690 Processors – 256GB Memory – Disk Controller HBA with 512MB Battery Backed Write Cache – 4 x 300GB 10,000 RPM SAS Disks	Each with: – 8 x Ten-Core Intel® Xeon® E7-8870 Processors (2.40 GHz) – 2 TB Memory – Disk Controller HBA with 512MB Battery Backed Write Cache – 8 x 300GB 10,000 RPM SAS Disks
14 x Exadata Storage Servers X3-2	**14 x Exadata Storage Servers X3-2**
With: – 12 x 600GB 15,000 RPM High Performance SAS disks or 12 x 3TB Performance SAS disks or 12 x 2TB 7,200 RPM High Capacity SAS disks Includes: – 168 CPU cores for SQL processing – 22.4TB Exadata Smart Flash Cache	With: – 12 x 600GB 15,000 RPM High Performance SAS disks or 12 x 3TB 7,200 RPM High Capacity SAS disks Includes: – 168 CPU cores for SQL processing – 22.4TB Exadata Smart Flash Cache

Figure 2. Adoption levels reveal data warehouse success.

policy, every parallel query will request data from all Exadata cells. Because of Oracle's shared disk architecture, concurrent queries or updates produce potential contention on each disk and will cause I/O bandwidth issues. The storage architecture is tightly coupled with ASM and configured to filter data as it reads it. This built-in storage filter is what Oracle calls Smart Scan or intelligent storage.

X3 can be configured with either 600 GB drives or 3TB drives. Oracle has named these high performance and high capacity respectively. The naming indicates the intended usage of these platforms. Oracle tries to sell the 3TB drives based on high storage and price per Terabyte. However, these fatter 3TB drives spin slower at 7.2K and will reduce overall system performance and concurrency. Customers will see that Exadata will struggle as the 3TB drives begin to fill up. When that happens, we would expect to see a decent performance loss since each drive has to do much more work. Long queries will tie up a drive and cause queuing to occur which will cause concurrency and performance issues.

- **Smart Scan** – One of the features of the Oracle Exadata Storage Server is Smart Scan. The Exadata cell can perform some filtering operations on subsets of the data. Thus, in many queries, predicates can be "pushed down" to the Exadata level for preliminary evaluation, eliminating the retrieval of unnecessary data. While this is true on performance, it does not mean that you can see 100x types of increase in performance, you will see smaller datasets transported across the network.
- **Flash Cache** – Oracle has implemented a database cache directly in the Exadata Storage Server. The Exadata Smart Flash Cache allows frequently accessed data to be kept in very fast flash storage while most of the data is kept in very cost-effective disk storage. Oracle claims that flash cache is smart because it knows when to avoid trying to cache data that will never be reused or will not fit in the cache. In real-life situations of data warehousing and business intelligence, one cannot pin data in flash cache. Because of the dynamic nature of many DW environments, knowing what questions are going to be asked and what data would need to be pinned would be difficult or nearly impossible. Oracle has a whitepaper (**Exadata Smart Flash Cache and the Oracle Exadata Database Machine**) that notes that OLTP receives the bulk of the benefits delivered by Flash Cache.

Oracle has claimed that if data is cached in the flash, the query will be optimized on the inbound side to the storage layer, and this provides scalability. In certain cases, Oracle suggests that you can pin multi-terabyte tables in the Flash Cache; while this idea is great, is it practically feasible? What Oracle does not mention when using flash as disk is the requirements to RAID protect or make the data permanently stored in Flash redundant (mirrored). In any real world production environment, this availability factor would have to be analyzed and set up for what Oracle suggests to be feasible. This would directly impact the amount of flash available and as well as its intended usage. Oracle strongly recommends using the flash as a cache and not as disk due to the high support and complexity requirements. The other area that is not mentioned is that data most move from the disk and flash storage tier up to the database (RAC) nodes when needed by a query. Large amounts of data being moved from the storage tier and between the RAC nodes will cause the Infiniband network to saturate. This limits the amount of data ingest from the storage to the RAC nodes to around 25GBs which is about the same speed as the high performing disks. From a data warehouse/BI perspective, Flash is useful only for reference data and master data that are read only in nature. It is best used for OLTP and in using the same with data warehouse workloads, no significant I/O boost is observed, as data is read only. The feature is great, but its usage is not visible in terms of data warehousing system performance.

Hybrid Columnar Compression

Hybrid Columnar Compression (HCC) is an Exadata-only feature that was introduced in Oracle 11gR2. HCC increases compression rates and thus decreases overall data sizes, which helps with I/O performance and system footprint. HCC is a multi-step process and works by taking rows in adjacent data blocks, called a compression group, vertically partitioning them by column, compressing each column partition, and storing the resulting compressed column partitions side-by-side in one or more data blocks as needed. The goal is to eliminate duplicate values in blocks and decrease data sizes.

- HCC works well on Exadata, as Exadata Smart Scan (the ability to do column and predicate filtering at the Exadata layer) can operate on compressed data. After initial filtering in the Exadata layer, functions including joins, sorts, group by, and analytic aggregation requires decompression and reassemble column values into rows.
- HCC, by Oracle's own recommendation, works best on static data, which means operational BI cannot truly use that feature when deployed on Exadata. (Operational BI - Competitive pressures are forcing companies to react faster to changing business conditions and customer sentiments. As a result businesses are using BI to help drive and optimize business operations on a daily or even intra-daily basis. This is referred to as operational business intelligence. This type of activity requires hosting OLTP and EDW types of data and workload on a singular platform, and that by itself means change in everything from data models to analytics).
- HCC can only be applied at bulk load time, which assumes more of a data reload approach to ETL/CDC.
- Only bulk SQL inserts are eligible for compression using HCC.
- SQL updates to HCC-compressed data cause the updated rows to revert to row organization, and migrate to uncompressed traditional rows, which results in decreased scan rate and OLTP compressed traditional row organization which increases storage requirements.

With so many caveats, HCC's real value is currently limited. Hence the HCC feature needs future enhancements to be generally applicable. While HCC can and will benefit some static data warehouses, more active data warehouses, especially ones that require mixed operational and analytical workloads, may see little benefit from HCC. In the field, user testing of VLDB data sets, the HCC feature has not scored well with users.

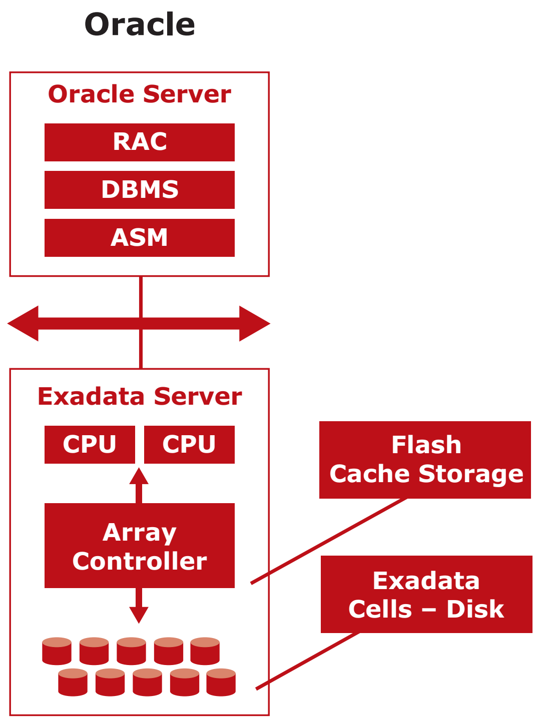

Figure 2. Exadata Internal Snapshot (source: Oracle).

Mixed Workload Capability

In the author's definition and measurement perspective, mixed workload processing is a favorite area of in-depth analysis by users today. It is the capability of a data warehouse to process many large and small queries of different types (BI, Analytical, ETL, ELT, multi-dimensional, Ad hoc, and OLAP) in the same system. This means one can run operational BI and analytical BI queries at the same time and get satisfactory performance that meets defined SLAs. Oracle definitely has more ground to make up in this area, which is traditionally a very strong point for Teradata. While the Oracle 11g database has improved in workload management, it certainly cannot be optimized to run mixed workloads across the same data warehouse.

To get good performance, there are several layers to constantly tune. First you will start with tuning the database, then the flash cache, the smart scans and result cache, the Exadata storage layer, and then manage system resources as they map to SLAs. You have to keep doing this over and over as data volumes and workloads change. Small, simple data warehouses will not exhibit any issues, but large data warehouses will struggle and probably fail. This is where Exadata is still Oracle 11g.

Management

Oracle claims that Exadata is self-managing. With a complex multi-tiered architecture, how is this possible? If you have to tune the RDBMS, RAC, ASM, and Exadata Storage layers, what is self-managing here? The sum of all parts is less than the whole. This is an area that DBAs will still have to work on, because the architecture is less integrated than advertised.

Scalability

Linear scalability refers to the capability of a system to scale in users and query volumes when new hardware is added to existing infrastructure. The architecture of Exadata is not linearly scalable. You can buy a quarter, half or full rack of Exadata system, but just adding more hardware does not mean that data volume and query workload will see linearly increasing scalability. Exadata's shared memory and disk RAC architecture, as well as its need for cache coherency are key limitations in its ability to scale. The RDBMS, ASM and Exadata storage layers have to be rebalanced and tuned to perform optimally.

Consolidation

Oracle claims that you can run multiple applications on one single database and Exadata platform. Oracle's workload management capability is difficult to set up and manage. Since Oracle RAC is a shared disk and memory architecture, workloads cannot be dynamically de-allocated once running in the system. As a result, an Oracle DBA must restrict workloads by RAC nodes, I/O and CPU. This complex relationship makes it difficult to integrate more than a few applications together on a single database and future restricts the overall functionality gained on Exadata. In the real world as customers add applications we see a degradation in performance that can lead to missed SLA's even when a few applications begin to run simultaneously on a single Exadata system.

Supportability

Problem with patching and testing cycles – I have also talked to customers about the high level of patching seen on the Exadata platform. Two types of patches are regularly seen: database and storage patches. Typically Oracle releases one or two database patches per month and a quarterly storage image patch which requires a system restart. In both cases, DBAs need to test these patches to ensure that no performance degradation is seen. These tasks are time consuming to the DBA, which affects the ability to deliver new applications to the business in a timely fashion.

Teradata Data Warehouse Appliance

Let's compare Teradata Data Warehouse Appliance on the same parameters.

Hardware Improvements

Teradata has a new series of platforms for the data warehouse appliance market.

Like Oracle Exadata, there is a lot of processing and disk power behind the Teradata appliances. The only difference here is the underlying shared nothing architecture and database are tested over 30 years, and like real quality wine, each version of the new features to the database has added a good deal of features to the data warehouse appliance family, technically the entry point to Teradata solutions.

Family Member	DATA MART EDITION	DATA MART APPLIANCE	EXTREME DATA APPLIANCE	DATA WAREHOUSE APPLIANCE	ACTIVE ENTERPRISE DATA WAREHOUSE
Scalability	Up to 6TB	Up to 8TB	Up to 234PB	Up to 21PB	Up to 61PB
Workloads	Test/ Development or Smaller Data Marts	Test/ Development or Smaller Data Marts	Analytical Archive, Deep Dive Analytics	Strategic Intelligence, Decision Support System, Fast Scan	Strategic and Operational Intelligence, Real-Time Update, Active Workloads

Table 1. Teradata's Workload-specific Platform Family.

Model	DATA MART APPLIANCE	EXTREME DATA APPLIANCE	DATA WAREHOUSE APPLIANCE	ACTIVE ENTERPRISE DATA WAREHOUSE
Nodes	Efficient single rack system including:	Integrated system of node and storage including:	Integrated system of node and storage including:	Integrated system of node and storage including:
Description	Single node – with Dual Intel Six- or Eight-Core Xeon® Processors	MPP nodes with Dual Intel Eight-Core Xeon Processors	Eight MPP nodes per cabinet with Dual Intel Twelve-Core Xeon Processors	MPP nodes with Dual Intel Six- or Eight-Core Xeon Processors
Operating System	SUSE Linux	SUSE Linux	SUSE Linux	SUSE Linux
Storage	Up to 48 300-, or 450-, or 600GB enterprise-class HDD drives and up to 8 400GB SSD	168 2TB or 3TB SAS drives per cabinet	288 300-, or 600-, or 900GB SAS drives per cabinet	Up to 174 300-, or 450-, or 600GB enterprise-class HDD drives per node and up to 20 400GB SSD per node
Total Capacity – User Data	Up to 8TB	Up to 234PB	Up to 21PB	Up to 61PB
Scalability	Single cabinet design	Scales up to 2,048 nodes	Scales up to 2,048 nodes	Scales up to 2,048 nodes
Availability	Includes RAID1 disk mirroring	Includes RAID6 data protection, global hot spare drives, and optional Hot Standby Nodes	Includes RAID1 disk mirroring with node failover and recovery	Includes RAID1, system availability with node failover and recovery; performance continuity with Hot Standby Nodes
Memory	Up to 256GB	Up to 256GB per node	Up to 512GB per node	Up to 512GB per node
Interconnect	N/A	Teradata BYNET® over 1Gb Ethernet or BYNET® V5	Teradata BYNET over 1Gb Ethernet or BYNET V5	Teradata BYNET V5
Workload	Data mart; test and development	Analytical Archive, Deep Dive Analytics	Integrated data warehouse; analytical sandbox; traditional decision support	Integrated data warehouse; active data warehouse
System Management	Cabinet level	Fully integrated across cabinets and system	Fully integrated across cabinets and system	Fully integrated across cabinets and system

Table 2. Teradata Platform Family (Source: Internet).

Let us look at the existing hardware architecture of Teradata.

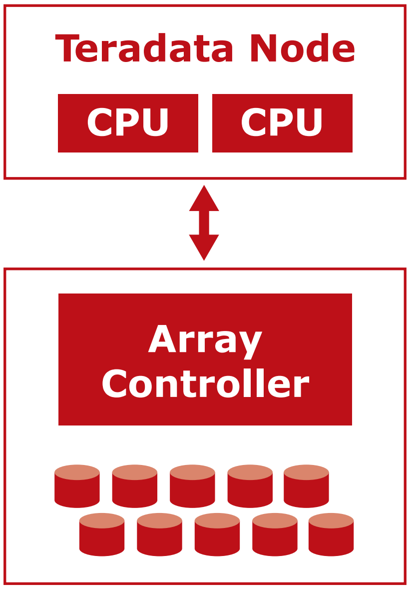

Figure 3. Teradata architecture (Source: Internet).

A typical node consists of multi-core processors, memory, and a disk subsystem that is virtually divided up into parallel shared nothing units called Access Module Processors (AMPs). Each AMP has its own little piece of CPU, memory and I/O as well as disk and the data that it owns. This type of architecture that is built from the ground up is why larger Teradata appliances move ahead of the pack. While industry experts have presented different opinions on MPP architectures, this architecture has been around for 25+ years and still continues to dominate the EDW world with proven customer scalability into the Petabytes.

At this juncture, let's deep dive into the important set of features that Teradata has emphasized as differentiators:

Teradata's Intelligent Memory (TIM) vs. Exadata Flash Cache

The first thing you need to realize is that Flash is not memory as Oracle likes to confuse customers. Teradata's is using database main memory DRAM that is 25X faster than flash so how could flash that is accessed at the block level be considered memory? Teradata is offering a solution that provides the right placement of data for the temperature of the data at the time it is needed. I think the WHOLE platform has to be examined – it is not fair to compare one feature in isolation to another in isolation. Oracle's Flash is in the storage tier, far from the database, while Teradata Intelligent Memory (TIM) is close to the database. In large memory systems, TIM greatly reduces physical I/O for the hottest data; reading from flash is still I/O and therefore much slower than memory access. Both are caches, flash cache is LRU-based, TIM is frequency-of-access-based & does a better job of keeping active, hot data in cache.

Because flash cards are block devices, they are accessed using processes similar to disk. DRAM memory is accessed differently – memory is not a block device. As a consequence, there are latencies in flash card access that are not present with DRAM. In addition, the flash cards reside in the Exadata storage layer remote from the database nodes' CPUs'. While IB rates are fast, accessing flash based data takes significantly longer than accessing DRAM. When large amounts of data are brought back from the flash we have a bottleneck in Oracle's use of Infiniband. No such limitations exists with Teradata's use of In Memory or when using their SSD drives solutions.

Compression

Multi-Value Compression (MVC), also called field compression, allows compression to be independently optimized for the data domain of each column. The granularity of compression is the individual field of a row. Field compression offers superior performance compared to row-level or block-level compression. The Teradata Database operates directly on the compressed fields– there is no need to uncompress a compressed row or field for query processing.

In the Teradata Database, up to 255 distinct values in each column can be compressed. If the column is nullable, then NULLs are also compressed. The best candidates for compression are the most frequently occurring values in each column.

Teradata also has block level compression. This allows data to be compressed at the block level within the storage sub-system which provides higher levels of compression. Teradata has dedicated compression cards that provide automatic block level compression of data. The compression ratios are around 4-6x in real-world implementations. If additional levels are desired, Teradata Database offers columnar compression which can be used in block-level compression. Teradata Database's columnar compression provides a similar ratio to Oracle's HCC compression. Unlike Oracle's compression methods, Teradata Database's methods can be used together to provide higher levels of data compression.

Recent versions of Teradata now offer true Hybrid Columnar storage and compression as well as Algorithmic compression. A combination or all of these can be applied to individual data use cases to reach extremely high levels of compression, without all of the limitations found in Oracle's single use of Exadata's Hybrid Columnar Compression.

Mixed Workload Capability

A key customer requirement today is mixed workload processing. This is the capability of a data warehouse to process large and small queries in the same database. This means one can run operational BI and analytical BI queries at the same time and get consistent performance per defined business SLAs. Teradata Data Warehouse Appliance's workload manager is the best-of-breed in managing mixed workloads in an appliance environment.

Teradata Data Warehouse Appliance has a single integrated workload management framework, where a set of rules is defined to assist in the monitoring and management of the varying mixed workload of an active data warehouse through Teradata's Web-based Viewpoint management portal.

Management

Teradata Database is self-managing and self-tuning from the database perspective. This has been an architectural feature since early days and now with the best-in-class workload management and data distribution algorithms, it remains unique and is at the head of the class. Functions such as memory management, storage allocation and management, reorgs, parallelism, are all automatic with no DBA setup or intervention needed. Teradata Database's ease of management enables lower TCO.

Scalability

Linear scalability refers to the capability of a system to scale and perform consistently when new hardware is added or new user/query volumes are added to existing infrastructure, and in this aspect, most appliances are scalable. But scalability is not just limited to data growth in transactional volume, but refers to the growth in dimensional data, number of concurrent users, complexity and volumes of queries, and workload mix, collectively referred to as multi-dimensional scalability. The architecture of Teradata presents a system that has been architected to handle the requirements of multi-dimensional scalability. In most situations, as users start measuring scalable (sustained) performance against growing workloads in terms of volume and complexity, the Teradata platform architecture definitely provides a more robust scale-up and scale-out model compared to Exadata and other leading appliances.

Consolidation

Teradata Data Warehouse Appliance allows for the dynamic control of workload resources. Running queries can be moved to a lower priority if needed reducing the amount of memory, CPU, and I/O needed on the system. This dynamic ability allows for the full utilization of all system resources and allows for a far greater number of applications to run within a single system and database.

Supportability

Teradata customers frequently talk about the fact that they only require one DBA to support a Teradata system no matter how large. Teradata also provides for on-site technical support for system patches and upgrades. This lowers overall costs and reduces the amount of time needed to take care of a Teradata system.

In summary the key difference that Teradata has established as an appliance is scalability, sustainability and maintainability in terms of performance, data growth, query complexity and system administration. These architectural traits have been designed and built into the appliance rather than rely upon a hybrid mix of software and hardware optimization. Let me draw a parallel here again. A bread toaster has a simple operation – to toast bread. If that is the fundamental premise of an appliance, at the bottom line, Teradata has been able to fulfill the needs of a data warehouse appliance. Read on to the next section to see what users have experienced in the field.

Criteria	Teradata Warehouse Appliance	Oracle Exadata
Support for Data Warehousing	Mature	Mature
OLAP Capability	Mature	Mature
Scalability	Mature	Mature
Multi User Management	Mature	Evolving
Maintenance Effort	Mature	Evolving
Training & Resource Availability	Evolving	Evolving
Customer Footprint	Mature	Evolving
Performance	Mature	Mature
BI Enablement	Mature	Mature
Partitioning	Mature	Mature
Workload Management	Mature	Evolving
MPP Architecture	Mature	Mature

Mature   Evolving

Figure 4.

From the Field

Here is a real-world perspective on the appliances and a selection criteria that my customers have applied to select a data warehouse appliance. Figure 4 shows the capability comparison of data warehouse appliances by the author (the chart is derived from the author's experience in the past years with real-life situations.)

A true scorecard from multiple consulting engagements is presented here, the scores are cumulative sets assigned by actual field users.

Based on the real-world feature comparison and additional data gathered from different engagements across the world, we can confidently conclude that the Teradata platform definitely is purpose built and with all the features and functions to address data warehouse requirements. In multiple field tests side by side, Teradata Database has proven scalability and performance on a sustained basis, and this is proof enough for executives to make decisions. Oracle Exadata and its family of products have been created to handle complexities in the OLTP space. To claim that OLTP and EDW, from a workload and complexity perspective, to be the same is not appropriate. I urge readers of this paper to carefully consider all the points of why you need a data warehouse appliance and how to select a vendor for the same.

Another trend that I am beginning to see is the evolution of Exadata as a consolidation play by Oracle. From this perspective, we cannot simply consolidate, the EDW or data marts with existing OLTP databases and claim either performance and cost benefits.

In discussing the scores with different CxOs and senior executives, the pricing model of Oracle Exadata database machine is not less expensive compared with Teradata Data Warehouse Appliance. The overall feature/functionality and price/terabyte/performance proves Teradata Data Warehouse Appliance to be lower cost in actual comparison. This is a key point that every organization can test during a POC/POV.

Another view point, from procurement and financial teams, is the acquisition costs of Oracle Exadata, which even with best efforts, leaves a huge hole in the budget. This has made the Teradata Data Warehouse Appliance even more appealing from a non-technical perspective. This is very interesting to hear from non-Teradata and recent (less than one-year old) Teradata customers.

The views expressed by the author is an independent expert's perspective and is not influenced by any individual or corporation.

	Teradata Data Warehouse Appliance	Oracle Exadata
Performance		
Data Load Performance	3	2
Aggregations	4	3
Single User Sequential Dashboard	5	3
Concurrency Testing	5	2
Mixed Work Load	4	2
Tuning		
In DB Aggregation	4	2.5
Indexing requirements	4	2
Ease of Tuning when issues occur	4	2
Interfaces		
Third Party BI Tool Compatibility	4	3
Failover/Fault Tolerance	4	2
Compression	4	2.5
Scalability of Machine	5	3
Financial		
Initial Purchase Costs	3	2
Future Purchase Costs	2	3
Service Startup Costs (Installation and Training)	3	3
Maintenance & Support Costs	3	3
Infrastructure		
Floor Space	2	2
Networking	3	3
Power	2	3.5
Database Backup	3.5	4
Application Monitoring	3	3
Operational Effort	2	2
DA/DBA Tools	2	2

5 – Best in Class 4 – Excellent 3 – Good 2 – Fair 1 – Poor

Figure 5. Comparison Scorecard

About the Author

Krish Krishnan is a worldwide-recognized expert in the strategy, architecture, and implementation of high-performance data warehousing solutions and big data. He is a visionary data warehouse thought leader and is ranked as one of the top data warehouse consultants in the world. As an independent analyst, Mr. Krishnan regularly speaks at leading industry conferences and user groups. He has written prolifically in trade publications and eBooks, contributing over 250 articles, viewpoints, and case studies on big data, business intelligence, data warehousing, data warehouse appliances, and high-performance architectures. He co-authored *Building the Unstructured Data Warehouse* with Bill Inmon in 2011, and his first independent writing project, *Data Warehousing in the Age of Big Data*, in May 2013.

With over 21 years of professional experience, Mr. Krishnan has solved complex solution architecture problems for global Fortune 1000 clients, and has designed and tuned some of the world's largest data warehouses and business intelligence platforms. Mr. Krishnan is currently promoting the next generation of data warehousing, focusing on big data, semantic technologies, crowdsourcing, analytics, and platform engineering.

Mr. Krishnan is the president of Sixth Sense Advisors Inc., a Chicago based company providing independent analyst services and management strategy, innovation and advisory consulting, in big data, data warehousing, and business intelligence. He serves as a technology advisor to several companies, and is actively sought after by investors to assess startup companies in data management and associated emerging technology areas.