Hybrid Storage for Data Warehousing

Colin White, BI Research
September 2011
Sponsored by Teradata and NetApp

Hybrid Storage for Data Warehousing

Ever since the advent of enterprise data warehousing some 25 years ago, organizations have faced the challenge of deploying enterprise data warehouse (EDW) solutions that provide good price/performance while also at the same time offering scalability for growing and new analytic workloads. This is especially the case today as companies experience dramatic leaps in data volumes and an increasing need to exploit the power of analytics for optimizing business operations in a constantly changing business environment.

This paper looks at new and evolving hardware and software options for building an EDW focusing specifically on the use of hybrid storage. As an example of hybrid storage in an EDW, the paper also reviews the implementation of solid-state drives (SSDs) in Teradata's Active Data Warehousing™ high performance hybrid storage platforms. The main objectives of the paper are to discuss the performance and cost tradeoffs of using SSDs, and to offer advice on their use in large-scale data warehousing considering bandwidth and high throughput requirements.

Before looking at the role of SSDs in detail, the paper first reviews data warehousing trends and discusses the EDW challenges and issues faced by IT today. It then moves on to look at how hybrid storage help solve those challenges.

The Performance Challenges of Enterprise Data Warehousing

Big data involves more than handling large data volumes

One of the biggest challenges facing organizations is how to handle rapidly increasing data volumes cost effectively. The term *big data* is often associated with this issue and there is a wealth of papers and web seminars covering this topic. It is important to realize, however, that in business intelligence and enterprise data warehousing, big data involves more than simply handling and managing large data volumes. It is also important to take into account the many different analytical workloads that exist in an EDW environment and the performance and service level requirements of those workloads.

Workload Complexity

Workload complexity and business agility are important considerations

In an EDW system, workloads may range in complexity from simple *ad hoc* reporting to the creation and running of sophisticated customer scoring models. An important consideration here is business *agility.* Companies using operational business intelligence (BI) analytics to enable faster decision making, for example, require the EDW system to handle not only complex mixed workloads, but also workloads that include time-sensitive applications such as fraud detection and call center optimization.[1]

When implementing time-sensitive analytical applications, it is necessary to balance the IT costs of supporting improved agility against the business benefits obtained. From a data perspective, one approach to achieving this balance is to place time-

[1] Teradata Corporation, the sponsor of this paper, calls such an agile EDW environment *Active Data Warehousing™*.

critical data on high-performance storage devices, and less used data on slower, lower-cost devices. The savings from using the lower-cost devices can help offset the costs of deploying the high-performance hardware. As we will see later, a data warehousing hardware configuration consisting of both SSDs and hard disk drives (HDDs) helps balance performance and hardware costs. The successful deployment of this hardware configuration, however, requires a good understanding of workload data usage patterns.

Data Usage Patterns

It is important to distinguish between frequently and infrequently used data – hot data vs. cold data

A data warehousing system stores and manages detailed, aggregated, current and historic data. Some of this data may be accessed on a regular basis, while other data may be used less frequently, or simply retained for regulatory reasons. The terms *hot data* and *cold data* are used to distinguish between frequently and infrequently accessed information. For good performance hot data should be readily available and have fast access times. Cold data, on the other hand, may be maintained on slower hardware devices to reduce costs. The problem here is that the temperature of data varies from workload to workload and even by time period. Month-end processing, for example, may increase the temperature of data. The solution is to migrate data between different types of data storage as its temperature changes – between an HDD and an SSD, for example.

Figure 1. Tiered Storage

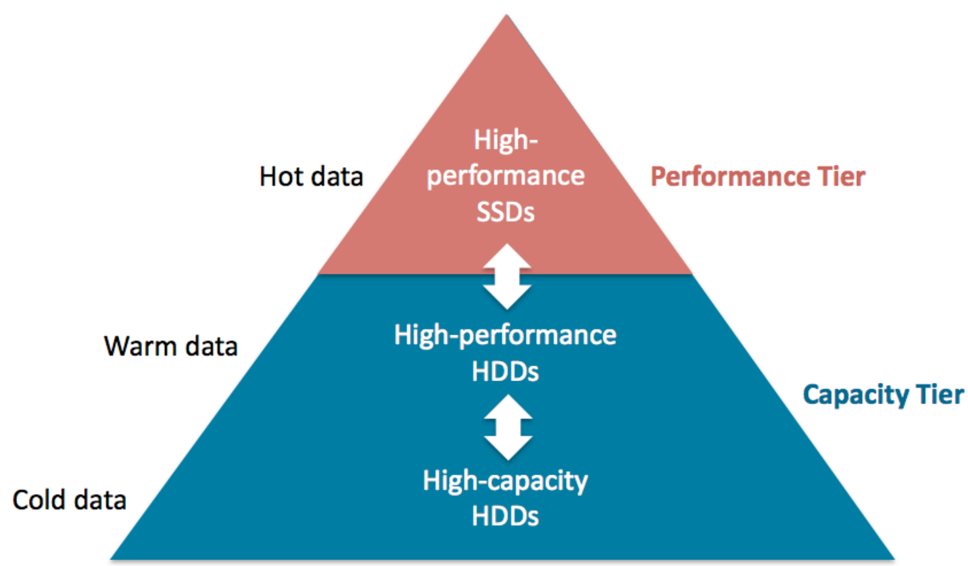

Data Migration Approaches

Need to migrate data between different speed devices as its temperature changes

Several vendors provide data migration between different hardware devices (tiered storage), but the level of automation and data granularity (block versus table level migration, for example) supported by products varies. Teradata includes Teradata Virtual Storage data migration technology with its active data warehouse systems. Teradata Virtual Storage allows organizations to use a mixture of SSDs and large and small capacity HDDs, which helps improve overall storage utilization. Teradata Virtual Storage keeps detailed statistics on data usage patterns and associates a temperature (hot, warm, cold) with each cylinder of data. It then moves or migrates the data between both SSD and HDD devices as the temperature of the data changes. This migration is done automatically by the software and is transparent to the DBA.

Capabilities such as Teradata Virtual Storage are valuable for balancing overall data performance against data storage costs for the majority of EDW analytical applications. The actual balance of hot to cold data usage determines the actual storage configuration required. Teradata experience shows that on average about 15-35 percent of EDW data is accessed 80 percent of the time. Moving the hotter data to faster data devices therefore improves the performance of most queries. The ability to migrate data at a low level of granularity is important for achieving optimal performance. This is why products that allow only table or table partition level migration are not always well suited to a large-scale EDW environment.

Automated approaches to data migration are usually more efficient

Automated and statistics-driven data storage migration is superior to manual approaches that require DBAs to determine which data is hot or cold and where it should be stored. There will inevitably be situations where some types of analytical processing will not benefit from an automated approach. In these cases, manual intervention may be necessary. An example would be a time-critical application that accesses a mixture of hot and cold data.

It is important to remember that an automated data storage migration feature works based on actual data usage, not on how time critical the data is to the application. Time-critical queries can of course be given priority on the use of available resources using the services of a workload manager. A well-designed workload manager coupled with enough HDD spindles to provide good input/output (I/O) performance will handle operational style analytical workloads when the data is not “hot enough” to be kept on SSD storage.

Balancing EDW Hardware Performance and Costs

Important to balance the EDW hardware configuration

An EDW system must support growing data volumes, different data usage patterns and increasing analytical workload complexity, while at the same time enabling business agility. The challenge for IT is how to build an EDW hardware configuration that balances these requirements against the costs of supporting them. One approach is to simply throw increasing amounts of low-cost commodity hardware at the problem. The increased floor space and power required by this approach, however, coupled with underutilized hardware resources and the additional administration costs involved, often offset the cost savings of using cheaper hardware. A better solution is to build a balanced and optimized EDW hardware configuration that maximizes the performance capabilities and capacity of the hardware.

Although there are many hardware components in an EDW configuration, one of the more critical aspects of the configuration is the need to balance processor utilization and the I/O performance of the data storage subsystem. If the configuration is I/O constrained then the storage subsystem will need to be upgraded. On the other hand, if the configuration is processor bound then the number of server nodes will need to be increased. This assumes of course that the hardware provider allows server nodes and disk drives to be scaled independently of each other.

Unbalanced configurations waste resources and reduce price/performance

As hardware evolves and changes, the balance in the number of server nodes and disk drives required to support EDW workloads also changes. Organizations must ensure that the total processor capacity per node is matched with enough I/O throughput to keep the processors busy. If there is inadequate or too much I/O capacity, the system is not balanced, and hardware resources become underutilized, which means the configuration is not achieving maximum price/performance. Designers therefore need to balance the number of processors, I/O throughput, and the total cost per megabyte in order to get a well-performing, cost-effective system.

In parallel with processor performance improvements such as those provided by multi-core Intel Xeon processors, HDD vendors are supporting the storing of ever increasing amounts of data on a single drive. Unfortunately, the disk and I/O interface performance of these larger disk drives is not keeping up with improvements in processor performance. To alleviate this issue, vendors are adding SSDs to their offerings to help balance the hardware configuration and offset the performance deterioration caused by high-capacity disk drives.

Hard Disk Drive Performance

Regardless of the type of processing involved, the main consideration from an I/O performance perspective is the amount of time it takes to read and write data from/to the storage subsystem used to manage the data. If the storage subsystem is based on HDD technology, this performance is affected by the performance of the hard disk drives, disk interface and interconnect speed, and database and file system architecture.

Head positioning performance has the biggest effect on random data access

HDD performance can be broken down into two areas: head positioning performance (*seek time* and *rotational latency*) and data transfer performance (*data transfer rate*). Head positioning performance is the most important factor for random data access, while data transfer performance has a bigger influence on sequential data access.

HDD seek times can be reduced by using smaller disk platters and/or higher bit recording and track densities. Rotational delay can be improved by increasing the rotational speed of the drive, but heat and vibration limit the level to which rotational speed can be increased. Data transfer rate depends on rotational speed and bit recording and track density. It is also important to note that the outer tracks of a disk platter contain more data sectors, and thus provide a higher effective data transfer rate.

Performance is also dependent on the interfaces used to transfer data between the HDD and disk controller and from the disk controller to the rest of the system. Key interfaces include fibre channel (FC), serial attach SCSI (SAS) and serial ATA (SATA). FC and SAS provide better performance and are used for high performance data warehousing, but they are more expensive than SATA solutions.

The Impact of High-Capacity HDDs

The industry trend toward using high-capacity drives has the advantage of economical storage, but the downside of longer average data access times. Two 500GB disks drives, for example, will provide better performance for most applications than a single one-terabyte drive.

The higher seek times of high-capacity HDDs impacts EDW performance

For mixed analytical workloads and enterprise data warehousing, where data access is more random in nature, the use of high-capacity HDDs increases the time it takes to access large amounts of data compared with using a larger number of smaller disks. It is less of an issue for single business area projects with more focused analytical processing workloads. In these latter projects, sequential disk reads are prominent, and the use of high-capacity HDDs has less impact on performance.

To overcome performance issues with high-capacity disks, IT organizations are resorting to techniques such as *short stroking* where only a percentage (typically 10 percent to 25 percent) of a drive's capacity is used and data may even be kept to the outer tracks of the disk to maintain performance. Short-stroked systems emphasize performance over the cost per megabyte of storage.

Prior to the use of SSDs, the best solution for achieving good performance for frequently used (or hot) data was to use low-capacity and high-speed drives, over provision the HDD subsystem, and short stroke the data across all of the HDDs. However, this increases costs, wastes disk space, and leads to larger storage subsystems, which in turn, takes up more floor space and consumes more power.

Industry is moving toward using a mixture of SSDs and HDDs to provide a balanced hardware configuration

Until recently, organizations faced the choice between high-performance (and higher cost) and high-capacity (and lower cost) hard disk drives. With the cost of SSDs dropping rapidly, another option is to use SSDs for high performance and HDDs for high capacity. IDC predicts[2] that in the enterprise HDD market there will be a shift away from higher cost performance-optimized HDDs to lower cost capacity-optimized solutions and SSDs to complement HDDs in storage systems.

SSD Technology Overview

SSDs have no moving parts and use flash memory to store data

The key difference between a sold-state drive and other data storage technologies is that the storage medium used in an SSD is solid-state semiconductor, rather than the magnetic medium used in an HDD, or optical medium used in a CD or DVD. Another major difference is that an SSD has no moving parts, which improves reliability and reduces power consumption. The main components of an SSD are the drive controller and the memory to store the data. Most commercial SSDs today use NAND flash non-volatile memory.

Performance, power efficiency and drive reliability are key distinguishing features in SSD products. These latter factors are highly dependent on the quality of the SSD's controller. Some vendors in the storage industry have adopted the term *enterprise flash drive* (EFD) to identify drives with high performance and reliability, and to distinguish them from the SSDs used in devices such as notebook computers. Like an HDD, an SSD may support a variety of different host interfaces (SATA, SAS, USB, etc.). Some manufacturers also offer hybrid SSD and HDD solutions.

The Benefits of SSDs

Solid-state drives combine high and consistent performance with better reliability than mechanical HDDs. They also use less energy.

[2] Worldwide Hard Disk Drive 2010-2014 Forecast, IDC, April 2010.

SSDs eliminate heading positioning overheads, which results in significant performance improvements

Improved Performance: There are no moving mechanical parts in an SSD, and the positioning overheads of an HDD are therefore eliminated. In an SSD, data is accessed directly from flash memory and access time is in the order of microseconds compared to milliseconds with an HDD. Read performance is consistent because read times are the same regardless of where the data is stored on the SSD. Write performance is slower than read performance, but is still much faster than with an HDD. Most enterprise-level SSD controllers also support multiple parallel read/write operations.

The 300GB or larger SanDisk Pliant SSD drives used by Teradata, for example, can achieve a 525MB/sec read throughput rate and deliver between 47K and 65K I/Os per second (IOPS) depending on block size and read/write ratios.[3] This is considerably faster than enterprise-level HDDs, which typically provide data transfer rates in the order of 20-40MB/sec and throughput rates of 200-400 IOPS. That is in the order of 160 times more IOPS, thanks to SSD flash memory.

Benchmarks may be needed to determine SSD performance gains

Comparing the performance of SSDs and HDDs is of course a moving target given the rapid change of hardware technology. Also, comparing the raw IOPS rates of individual devices is not a good indication of how this will be reflected in actual application performance. Instead, SSD usage should be considered in terms of the overall system and the application workloads that run on that system. In most cases, an application benchmark will be needed to ascertain the real performance improvement of using SSD storage.

Lower Energy and Rack Space Requirements: An SSD has lower power requirements than an HDD, and given that multiple short-stroked HDDs can be replaced by a single SSD to supply the same IOPS, this results in considerable savings in both power needs and rack space requirements. SSDs employ industry-standard 2.5/3.5-inch form factors and are compatible with existing host interfaces, which eases SSD migration into an existing rack infrastructure.

Cost Considerations: Most of the debate concerning the use of SSDs focuses on their higher cost per megabyte compared with HDDs, and whether or not this cost can be offset by better performance. Vendors are forced into this debate because there is a lack of industry knowledge and customer awareness about the value and use of SSD technology. SSDs are a disruptive technology, and as can be seen in the discussion above, there are several factors that need to be considered when determining their cost savings.

SSDs cost more but provide better performance

There is no doubt that for many types of workloads SSDs offer significant benefits, and given that the cost of SSDs is dropping, these benefits can only improve. Measured in terms of cost for performance, SSDs are a good bargain even though their cost per megabyte is much higher than HDDs. This is why organizations need hybrid storage systems to balance costs and performance as disk technology changes.

[3] Source is the SanDisk Pliant Technology specification sheet for the LS300S SSD.

USING HYBRID STORAGE FOR DATA WAREHOUSING

A hybrid storage EDW configuration can use SSDs and a mixture of different HDD types

The best approach to data management in a hybrid SSD and HDD environment is to put hot data on the SSDs (eliminating the need for short stroking HDDs) and the remaining data on the HDDs. The HDD configuration can be optimized to suit the warm and cold data, and may consist of a mixture of low and high drive capacities and speeds as appropriate.

The challenge is to determine the workloads that can benefit from the use of SSDs, the SSD and HDD configuration required to support those workloads, and which data should be stored on SSDs rather than on HDDs.

In general, most high-volume OLTP workloads will benefit from the use of SSDs. These workloads are well defined, and it is easier to optimize the use of SSDs for OLTP than say for an EDW workload, which is more mixed and interactive in nature.

SSDs provide high performance for operational BI and random read/write EDW workloads

The high performance of SSDs is advantageous for EDW workloads that require high-volume and low-latency read/write operations or intensive random read operations. SSDs are, therefore, particularly valuable in operational BI environments where fast response times are required. Many Teradata active data warehouse customers, for example, will gain from the use of SSDs because they have certain performance critical operational BI workloads that involve a high percentage of random data access operations.

SSDs can also improve the performance of analytical workloads that sequentially scan data. Although the improvement in this case will generally be lower than for workloads where random data access dominates, the actual improvement will vary by workload. Sequential workloads that consume small amounts of processing power per row will see more improvement than those with high compute overheads per row. Sequential workloads that include random read operations (to access reference or dimension tables, for example), and mixed workloads with both short- and long-running queries that disrupt sequential table scans, will also gain from the use of SSDs.

Processor-constrained workloads will not benefit from the use of SSDs

Figure 2 shows a Teradata SSD query benchmark result for one of its customers. The client uses the benchmark to validate Teradata systems prior to purchasing upgrades. The 40 concurrent complex queries run in these tests use 7 terabytes of real customer data. Some 28 percent of the data was measured to be hot. The 40 percent performance improvement was achieved when 90 percent of the hot user data was migrated from the HDDs to the SSDs. The customer's workload was I/O constrained, which accounts for the performance benefits obtained by using SSDs. This will not always be the case since many BI queries are processor constrained.

Data loading is another critical data warehouse workload. Teradata claims internal SSD benchmarks reveal an average of 10 percent to 20 percent performance improvement for a variety of data load tasks.

Figure 2.
Teradata SSD query benchmark

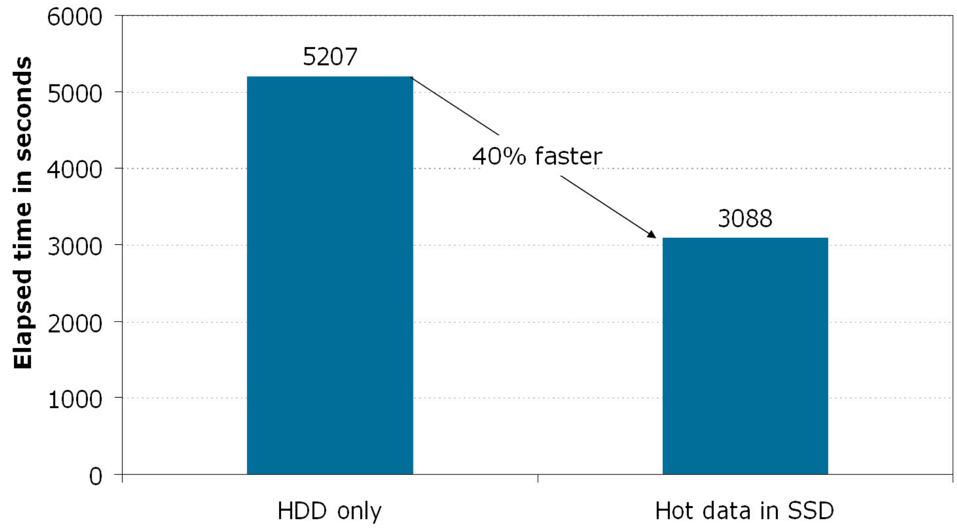

Automated versus Manual Approaches to Storage Optimization

Many RDBMS vendors are delivering optimized solutions for enterprise data warehousing, and these solutions often include the option to add SSDs to the configuration. These systems provide optimized hardware and software configurations that ease the task of balancing the hardware for optimal performance.

The various vendor offerings differ in the techniques and tools provided to help with the management of hot and cold data across the SSD and HDD configuration. These capabilities range from a purely manual approach where the DBA has to do the work, to a fully automated approach where the DBMS software makes the determination. Teradata chose to use its Teradata Virtual Storage technology to provide a fully automated approach when using SSDs in its Active Data Warehousing™ systems.

Manual data migration between SSDs and HDDs requires significant DBA skills and resources

The benefit of a manual approach to storage optimization is the level of control it provides. The downside is the skills and significant DBA resources required to do this tuning. Manually managing hot and cold data is a multi-step process. First, the DBA must identify the hot and cold data, a task that may be difficult depending on the tools provided by the vendor. Next, the DBA must close the system down and move the appropriate data onto SSDs and HDDs. In some cases, this may require the DBA to redesign database schemas to optimize partitions for SSD usage. Also, since BI users evolve and change their queries over time, the DBA must monitor these changes and revise the configuration as appropriate to maximize performance. To reduce this manual effort, some suppliers offer rules-driven storage migration from HDD to SSD. This again usually involves taking the system offline to do the migration.

Data granularity is a key factor in the effectiveness of data migration technology

An issue with both the rules-driven and manual approaches is that they usually support the migration of data only at the database, tablespace, table, or partition level. This level of granularity is insufficient for an EDW environment involving complex workloads. Moving partitions or large chunks of data onto SSDs will usually result in a mixture of both hot and cold data on the SSD, which reduces the effectiveness of the SSD. Migration at the data block or cylinder level is preferred, assuming the RDBMS can keep the data blocks full and hot data rows together.

In general, automated approaches to handling hot data are superior for enterprise data warehousing. Of course, the test here is the intelligence of the automated software. In addition, there will also be outliers in the workload that will not benefit from automation or the performance SSDs offer.

SUMMARY

SSDs improve the performance of a wide range of workloads, but it is important to set expectations because not all workloads benefit equally

Solid-state drives provide needed relief in storage performance and offer significant benefits for many types of workloads. Data loading, operational BI, and many complex queries will achieve faster response times and better throughput when using SSDs. Whereas I/O constrained systems will benefit from SSD use, processor-constrained workloads will not. It is therefore important to set user and management expectations appropriately.

Another benefit of using a combination of SSDs and HDDs lies in the flexibility to configure I/O performance using fewer drives compared with earlier I/O constrained HDD systems. To take advantage of this benefit, however, customers should look to the intelligence of data migration software to ensure the best use of the SSDs.

In summary, hybrid storage provides significant performance benefits for large-scale enterprise data warehousing, and many organizations can benefit from its use.

I would like to thank the following companies for their support in the publication of this paper.

Sponsored by

TERADATA®

NetApp®
Go further, faster

About BI Research

BI Research is a research and consulting company whose goal is to help companies understand and exploit new developments in business intelligence, data management, and collaborative computing.

EB-6436 > 0911

