Improving System Performance Using Solid-State Disk

Solid-state disk (SSD) provides 100 times the I/O performance of hard drives (HDDs) but at 100 times the cost per gigabyte or more. Compared with other techniques such as in-memory databases and database tuning, SSD can be a cost-effective solution to improving the performance of I/O-bound applications. But in order to achieve the proper balance between cost and performance, it is necessary to first eliminate duplicate data produced by the applications, identify precisely which data and database elements are causing the I/O bottleneck, move only those data and database elements to the SSDs and then move them back to slower but less expensive HDDs as soon as such a move will not impact application performance.

This white paper explains how Cognizant used SSDs to achieve performance improvements of as much as 85% for two sampling applications accessing an Oracle 10g environment for a well-known U.S. market research firm. It describes how we evaluated the use of SSDs compared with several other performance-enhancing technologies and determined which data and database elements to move to SSDs to achieve the most cost-effective performance improvements. It also explains the differences between HDDs and SSDs, as well as among the several types of SSDs on the market.

This document will be useful for:

- Enterprise architects focusing on improving system performance.
- Management groups responsible for decisions on IT infrastructure budget.
- Performance engineering team members responsible for making hardware recommendations.

SSDs Explained

SSDs are storage devices that use memory (DDR or flash) as the primary storage media. While various types of SSD use different types of memory and vary in performance, all deliver far more performance than spinning magnetic disks because they use no moving parts to store and retrieve data.

Lab tests suggest that an enterprise-grade SLC (single-layer cell) flash drive can reach performance as high as 250,000 IOPS, compared with only 350 IOPS for a high-end conventional magnetic disk. The typical access time for a flash-based SSD is about 35 to 100 microseconds, compared with 5,000 to 10,000 microseconds for a rotating disk. An SSD can typically transfer data at 170 megabytes per second, compared with 80 megabytes per second for a rotating disk. Their lack of moving parts also means SSDs typically last longer, use less energy and generate less heat than magnetic media, which together can reduce maintenance and power costs over time.

Because SSDs appear as standard disk drives to servers and applications, they have other advantages over application- or database-level solutions. They pose no hardware or software compatibility issues for users or administrators and require no special training or tweaking of applications or databases.

- The prime — and significant — disadvantage of SSDs is cost. Traditional hard drives cost about 20 to 30 cents per gigabyte for 2.5-in. laptop drives and 10 to 20 cents per gigabyte for 3.5-in. desktop drives. Most consumer-grade SSDs from leading vendors now cost around $3 per gigabyte. An enterprise-grade SLC flash drive appropriate for use in this client's application would cost around $50 per gigabyte, while a high-end SAN with spinning hard drives costs around $5.50 per gigabyte.[1]

SSDs come in several different configurations. Low-end consumer SSDs are based on multi-layer cell (MLC) technology, which stores two bits of data in each memory cell and is half as fast (but half as expensive) as SLC technology, which stores one bit of data per cell. Enterprise-grade SSDs, which are focused on performance and data integrity, use an SLC design, along with the use of RAID algorithms, to ensure data protection in case of disk failure, error code correction, wear-leveling (to extend the life of the individual cells) and active-spare capacity to instantly migrate data from a failing drive to a healthy one. SSDs can also be built using NAND-based flash memory, or RAM memory, which provides even faster performance but does not retain data without a power source.

SSDs are available in several configurations:

- DDR solid-state disk provides the fastest performance and typically includes internal batteries to keep the unit powered on long enough after shutdown to write data to a backup disk within the unit.
- Cached-flash SSD combines the advantages of DDR SSD and all-flash SSD, balancing the speed of DDR SSD with the storage capacity and lower cost of flash SSD. A cached flash solid-state disk is as fast as a DDR solid-state disk for cache hits (when the data is read from the cache) and still 20 times faster than typical hard disk-based solutions if there is a cache miss that causes data to be read directly from the flash memory. However, this type of SSD is still relatively costly and not as fast as all-DDR SSDs.
- All-flash SSDs offer higher capacity at a lower price than either DDR or cached flash, because they do not require the additional batteries to flush the DDR cache during power outages, nor large amounts of expensive DDR memory. Instead, a small amount of DDR provides a buffer for writes and for metadata. The RamSan 620 used in our proof of concept is an all-flash SSD with sustained 250,000 random IOPS.

SSD Proof of Concept

As a part of an enterprise architecture review engagement, we performed a proof of concept of the use of SSD to improve the performance of a system of sampling applications for a well-known U.S. market research firm. A high latency in disk I/O was slowing the performance of the processes that selected panelist data from an Oracle database, eliminated unwanted data and balanced the data set to give it the desired characteristics to determine which panelists met the criteria for a follow-up survey.

The research firm had tried other methods of improving performance, such as tuning application queries, redesigning the database table structure, reindexing the database, as well as archiving some data to reduce data volume. Some of these methods compensated for poor disk performance to some extent but were time-consuming, did not address the disk performance issue directly and, therefore, had limited impact on system performance.

After evaluating various hardware and software-based technologies (see Appendix C), the research firm decided to implement a proof of concept to determine if SSD could be a cost-effective way to improve performance.

The first step was to closely examine the selection, elimination and standard balancing applications used in the sampling process, since these were the most I/O intensive. Further analysis identified the database tables used and the volume of data read and generated by each program, and found each creates four tables, all of which have a significantly large data volume. A preliminary look at the programs revealed that they created duplicate data in three out of those four tables. Modification of the programs led to significant simplification of the sampling process and consolidation of these three tables into one, achieving a data volume reduction of about 50% for each sampling process run.

We then worked with the client SMEs to identify the data and table space that could be moved to SSD. Candidates for movement were:

- The tables that are read during the sampling process.
- The indices, to speed the selects from the input tables, as well as updates/inserts on the output tables.
- The redo logs to improve the speed of database writes.
- The undo table space, in order to handle concurrent processes faster.
- The temporary table space to make the standard balancing program run faster, by improving the performance of the sorting logic in the balancing process.
- Index spaces and temporary table space to improve sequential read and sorting speed, respectively.

To ensure the more expensive SSD was used only when it would speed application performance, just the data read and written by these programs would be placed on SSD, with the output data moved frequently from the SSD to HDD.

Based on these calculations, it was determined that a 500 GB SSD was required. However, since 500 GB SSDs were not available, and to allow future reuse in production systems, CTS chose a 1 TB SSD.

Results

Replacing HDD with SSD reduced job completion times by 75% and 66%, respectively, for the elimination and selection processes, as shown in the tables below.

Elimination Module	
Environment	**Time (minutes)**
Non-SSD	2.15
SSD	0.54
The mean job completion time in the SSD environment is only 25% of that required in the non-SSD environment.	

Selection Module	
Environment	**Time (minutes)**
Non-SSD	7.40
SSD	2.50
The time taken in the SSD environment is almost a third compared with that in the non-SSD environment.	

Summary

While they cost about 100 times more per gigabyte than spinning disk, SSDs deliver about 100% more I/O performance. If used properly, they can be a cost-effective option for improving the performance of I/O-bound application processes, compared with other techniques such as in-memory databases and database tuning. Making the most efficient use of this expensive resource requires first eliminating duplicate data produced by the process; identifying precisely which data and database elements are causing the I/O bottleneck and moving only those to the SSDs; and moving data off the SSD onto slower, but less expensive HDDs as soon as such a move will not slow application performance.

Appendix A: Solid-State Disks

SSDs originated from two similar concepts developed in the 1950s: core memory and Card Capacitor Read Only Store (CCROS). In the 1970s and 1980s, SSDs were implemented in semiconductor memory for early supercomputers from IBM, Amdahl and Cray. However, the prohibitively high price of the built-to-order SSDs caused them to be used rarely. Flash memory (both NOR and NAND types) was invented by Dr. Fujio Masuoka while working for Toshiba. Most modern, high-performance, large SSDs are based on NAND flash drives.

Flash memory stores information in an array of memory cells made from floating-gate transistors – resembling a standard MOSFET with two gates – as well as a control gate (like any MOS transistor) and a floating gate. Storage of electrons in the insulated floating gate enables the flash memory to store data without a battery backup.

In an SLC flash, an array of transistors (also known as a cell) is used to store one bit of data. MLC flash memory can store more than one bit in a cell, but it is less dependable than the SLC flash. Data is stored directly on memory chips and accessed from them, resulting in storage speeds far greater than is even theoretically possible with conventional magnetic storage devices. To make full use of this speed, an SSD typically connects to servers or networks through multiple high-speed channels. Since SSDs use no moving parts, they last significantly longer, are more rugged and use significantly less power than hard disks.

The differences in speed and price between consumer SSDs and EFDs (Enterprise Flash Drive) are shown below:

SSD vs. EFD Comparison

	SSD	EFD
Vendor	OCZ	High-end Texas Memory System 1 TB flash-based SSD
How it Looks		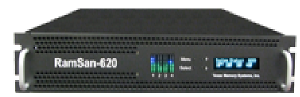
Technical Specification	• Size: 256 GB, 512 GB, 1 TB SLC flash storage • Random sustained throughput: › 10,000 IOPS › 600 MB/sec • MTBF: 1,200,000 hours • Power consumption: N/A	• Size: 1-5 TB SLC flash storage • Random sustained throughput: › 250,000 IOPS › 8 GB/sec • MTBF: Uses fault-tolerant flash (FTF). The flash used in the RamSan-620 is enterprise-grade SLC flash. • Power consumption: 230 watts
List Price	$4,000 for 1 TB	$61,450 for 1 TB

Figure 1

In addition to their much higher I/O throughput, some EFDs have other capabilities that increase their usefulness. The RamSan 620, for example, includes dual-port 4G-bit Fibre Channel controllers, the option of up to four DDR InfiniBand ports and the ability to designate one flash card within the drive as an active spare that works hand-in-hand with the chip-level RAID. If one of the cards experiences a failure that degrades its RAID protection, the system immediately migrates its data to the hot-spare to return to a fully redundant state.

The RamSan-620 also includes wear-levelling technology that maximizes system life by spreading out writes, and the ability to upgrade a single 2U chassis from 1 to 5 TB of flash storage and with as many as eight Fibre Channel ports or four InfiniBand ports in the future. Multiple units can be added to meet any capacity or performance requirement.

Appendix B: Detailed Performance Analysis

The mean timing data presented below is calculated from the data gathered from the database logs, assuming a linear fit of around 50% completion time. The following tables show the time taken for 50% of the jobs to be completed in the two environments. These mean timings show the impact of SSD on performance of selection and elimination jobs.

Elimination Module	
Environment	**Time (minutes)**
Non-SSD	2.15
SSD	0.54
The mean job completion time in the SSD environment is only 25% of that required in the non-SSD environment.	

Selection Module	
Environment	**Time (minutes)**
Non-SSD	7.40
SSD	2.50
The time taken in the SSD environment is almost a third compared with that in the non-SSD environment.	

Performance Analysis

The following sections display performance data gathered from the balancing and elimination applications within the client's sampling process. The process's start and end times, spanning several weeks, were taken from database log tables.

Performance data was analyzed to find the percentage of processes that completed in a fixed interval of time (2 minutes, 4 minutes, etc.). Finally, the data was plotted against time for performance comparison. In the section below, performance data for an environment without SSD is compared with one with SSD.

Result for Elimination Module

As-is Performance Data

The following table shows the cumulative percentage of elimination processes completed as a function of time in minutes. This data is collected from the as-is system before SSD deployment, showing 99% of processes were completed within 80 minutes of start time.

Cumulative % without SSD	48.64%	66.51%	72.47%	76.97%	80.11%	82.86%	86.84%	88.45%		99.80%
Time (minutes)	2	4	6	8	10	12	14	16		80

Performance Data with SSD

The following table shows the cumulative percentage of elimination processes completed as a function of time in minutes. This data is generated from the system after SSD deployment.

Ninety-nine percent of the processes are completed within 12 minutes of start time. This constitutes a processing time reduction of 85%.

Cumulative % with SSD	81.22%	89.14%	93.25%	95.57%	97.57%	99.26%	99.26%	100.00%
Time (minutes)	2	4	6	8	10	12	14	16

Performance Comparison

The percentage comparison of job completion in the two environments is represented below:

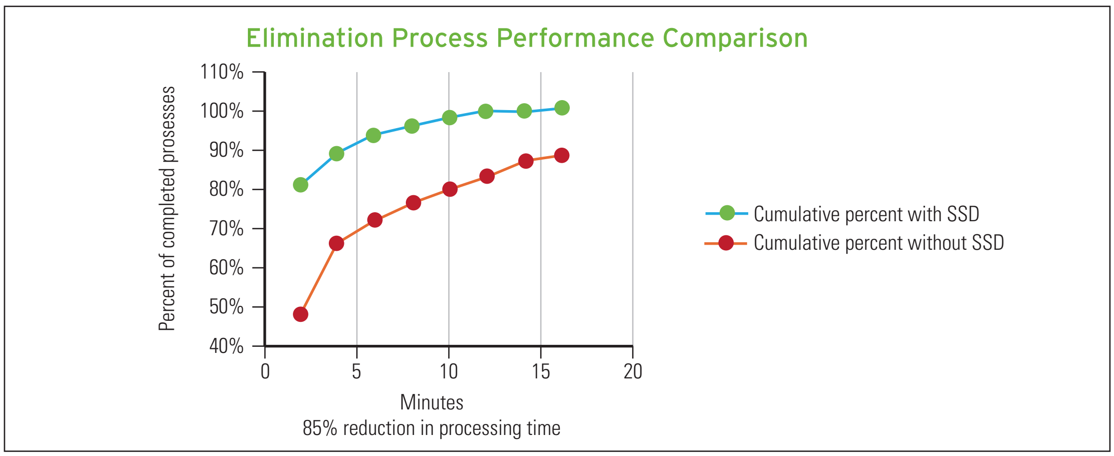

In the above graph, each point represents the percentage of total jobs completed within the specified time duration. The data is taken from the tables presented above. In general, we see that in both environments, more jobs are completed in the first half of the time than in the second half. In other words, in the SSD environment, while a 100% completion took 16 minutes, about 95% completion was done in eight minutes.

In the non-SSD environment, we perceive the same trend. This indicates that some jobs in either environment always tend to linger. These are most likely to be complicated jobs running large queries, or they may be running at a time when the system is busy processing other jobs.

Result for Selection Module

As-is Performance Data

The following table shows the cumulative percentage of selection processes completed as a function of time in minutes. This data is collected from the as-is system before SSD deployment, and shows that 99% processes are completed within 72 minutes of start.

Percent completion without SSD	0.00%	2.68%	33.13%	57.29%	70.49%	77.02%		99.32%
Time (Minutes)	2	4	6	8	10	12		72

Performance Data with SSD

The following table shows the cumulative percentage of selection processes completed as a function of time in minutes. This data is generated from the system after SSD deployment, and shows that 99% of the processes are completed within 12 minutes of start. This constitutes a processing time reduction of 83%.

Percent completion with SSD	37.25%	88.21%	96.67%	97.93%	98.44%	99.14%
Time (Minutes)	2	4	6	8	10	12

Performance Comparison

The percentage comparison of job completion in the two environments in the selection module is represented below.

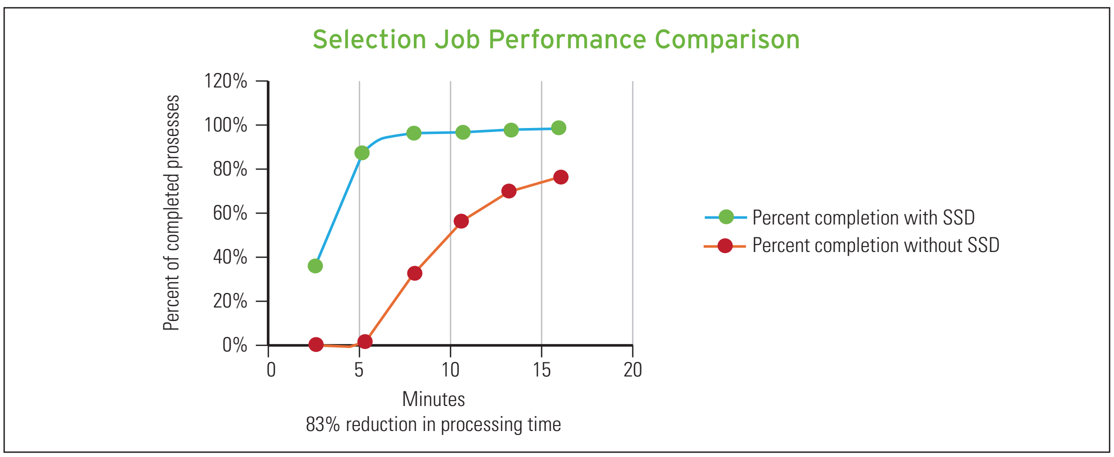

In the above graph, each point represents the percentage of total selection jobs completed within the specified time duration. The data is taken from the tables presented above. In general, we see that in both environments, more jobs are completed in the first half of the time than in the second half. In other words, in the SSD environment, while a 99% completion took 12 minutes, more than 96% completion was done in six minutes.

As in the case of elimination, we see that some jobs in either environment (SSD and non-SSD) always tend to linger. They are most likely the complicated jobs running large queries, or they may be running when the system is busy processing other jobs.

Result Analysis

In the following section, we present an analysis of the result observed above. The mean timing data is calculated from the data presented above, assuming a linear fit of around 50% completion time.

Elimination Module

The following table shows the time taken for 50% of elimination jobs to be completed in the two environments.

Elimination Module	
Environment	**Time (minutes)**
Non-SSD	2.15
SSD	0.54

The mean job completion time in the SSD environment is only 25% of that required in the non-SSD environment.

Selection Module

The following table shows the time taken for 50% of jobs to be completed in the two environments. The time taken in the SSD environment is almost a third compared with that in the non-SSD environment.

Selection Module	
Environment	**Time (minutes)**
Non-SSD	7.40
SSD	2.50

Appendix C: Technology Evaluation

Before beginning the proof of concept using SSDs, we evaluated several other hardware and software-based technologies, focused specifically on improving disk I/O latency in an Oracle environment. These technologies were the Oracle TimesTen In-Memory Database (IMDB) and Sun Oracle Exadata.

The Oracle TimesTen In-Memory database[2] is a memory-optimized relational database that can significantly reduce latencies associated with traditional disks and thus enhance performance of I/O-bound sampling system processes. Deployed in the application tier as an embedded database, it operates on databases that fit entirely in physical memory using standard SQL interfaces. The Oracle In-Memory Database Cache (IMDB Cache) uses the Oracle TimesTen In-Memory Database as its RDBMS engine.

This technology was deemed unsuitable for the client's needs because:

- Hardware support for the TimesTen database is rather limited as high-end blade servers provide only 192 GB RAM, compared with the terabtye of RAM or more required to achieve the performance improvement seen with SSD.
- Only a fraction of the client's database (projected to grow to 12 TB) can be put in one TimesTen database instance (being completely memory resident).
- TimesTen would have required a complicated database distribution among several server instances in order to achieve the desired performance increases.
- TimesTen would have caused a significant increase in complexity, resulting in higher hardware, administrative and total cost of ownership (TCO).

The Sun Oracle Exadata technology is based on smart database software that addresses key database I/O issues to speed application performance. It is based on a massively parallel architecture, which provides more pipes (or data pathways) to deliver more data faster between the database servers and the storage servers. It also uses wider pipes than those usually used with Oracle databases to provide extremely high bandwidth between the database servers and the storage servers. This technology is database aware and can ship just the data required to satisfy SQL requests, resulting in less data sent between the database servers and the storage servers. It also overcomes the mechanical limits of disk drive technology by automatically caching frequently accessed data.

However, Oracle Exadata was not considered a viable option because:

- It could only work with Oracle 11g, not the Oracle 10g the client is using and has no plans of upgrading.
- Its performance gains were significantly lower than the TimesTen or SSD options.
- The total cost of ownership would be high due to training and deployment issues.
- The use of Oracle Exadata would lead to greater vendor lock-in, which the client wanted to avoid.

Appendix D: Proof of Concept

The application for which we performed this proof of concept is used by a leading market research company in the U.S. to sample panelists' data to determine who will receive follow-up survey questionnaires. A suite of sampling system applications supports the process. This suite comprises two different technologies: Oracle (procedures and Forms) and Java/J2EE (used for balancing). The suite of applications also supports Panelist's Points Import Process.

This diagram shows the client's overall process flow with the sampling and points import processes highlighted.

Client's Process Flow

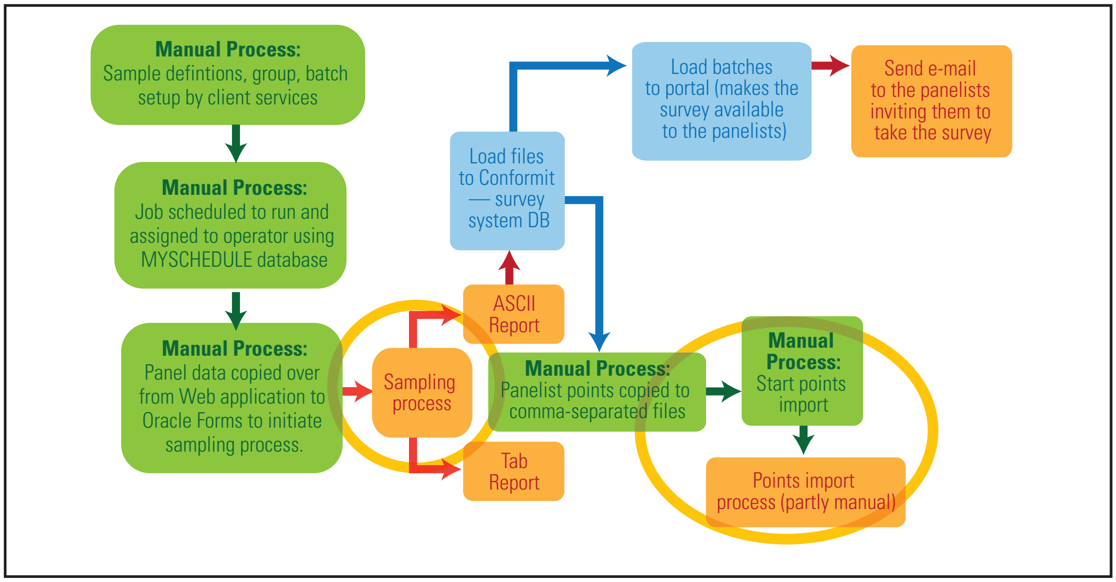

Figure 2

Applications and processes that are required to read or write frequently from the database are identified. The sampling process was considered a candidate due to several I/O-bound applications. The sampling process runs in three different regions:

- **North America (U.S.):** Balancing is performed using standard balancing module (PL/SQL), as well as BPD balancing module (Java).
- **Europe (EAP):** The DRAMA balancing submodule of the sampling system is written in PEARL (Process and Experiment Automation Real-time Language).
- **Asia Pacific (ALM):** Balancing is performed using standard balancing module (PL/SQL), as well as BPD balancing module (Java).

The following diagram details the sampling process. Selection, elimination and standard balancing applications (highlighted above) of the sampling process were identified as I/O-intensive modules suffering from poor performance.

The selection program of the sampling process selects, or quarries, data (also known as the panel data) from the client's database tables. The balancing module then calls the elimination program to screen the selected panelist data, based on predefined criteria. After elimination, the balancing logic distributes the remaining panelist data as per the batch and quota definition of the job, and the final output is stored in the database table for further processing.

The following factors were identified:

- The balancing process is also CPU intensive, and therefore, use of SSD alone would not improve performance significantly.
- The sampling process quarries data from database tables and writes processed data at various stages to the database.
- Sample data has to be stored permanently to produce reports, send e-mails to sampled panelists and to partially re-run the process in the future on the client's request. This requires that any size and cost calculation of SSDs needs to take into account the volume of data

Sampling Process Flow

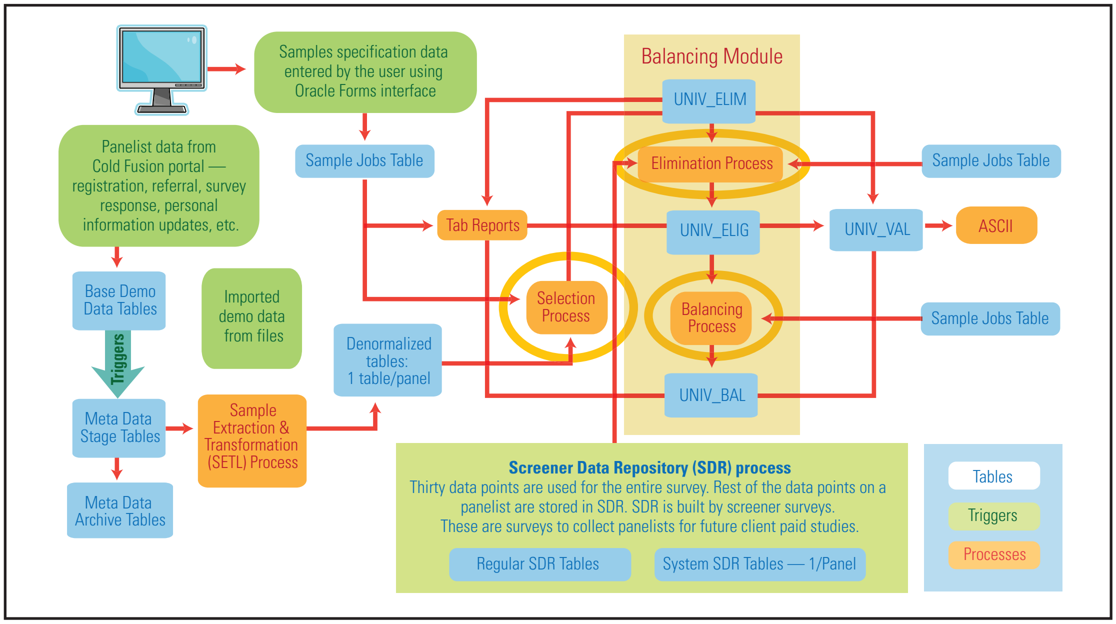

Figure 3

being generated, moved to HDDs and, upon clients' requests, restored to the SSDs.

- The volume of data generated by sampling is very large, and this was originally identified as one of the problem areas. The redefined sampling process cut the data volume by 50% (with the new sampling process using only two tables instead of four).

The following are the major steps undertaken to devise a strategy for SSD deployment in the European environment, with a goal of minimizing the use of SSD to strike a balance between cost and performance:

- Identify the I/O-bound processes/programs.
- Identify the tables read by these programs and the tables to which the data is written.
- Plan to transfer only the input and output data of these processes to the SSD.
- Move output data frequently out of SSD to HDD to save SSD space.
- Move the tables that are read during sampling process to SSD.
- Move the indices to the SSD to speed up the selects from the input tables and updates/inserts on the output tables.
- To improve the speed of database writes, move the redo logs to SSD.
- To handle concurrent processes better, move the undo table space to the SSD.
- To make the standard balancing program run faster, move the temporary table space to the SSD, improving the performance of sorting logic in the balancing process.

Based on this analysis, we calculated the amount of data to be stored on the SSD to balance cost and system performance.

As mentioned above, the sampling process generates data that is to be stored in the system on a permanent basis. The amount of data generated depends on the number of job runs. Each job produces a different quantity of data. However, over a period of one week, total data produced for the entire weekly job runs is found to be a good representation of a weekly average.

Sampling data can be grouped into the following categories:

- **Offline Data:** This is relatively old data (more than three months old) and is stored on backup tapes. When needed, it is restored to disk for further processing. The proposed SLA is 72 hours.
- **Online Data:** This data is generated in the last three months and is kept on the hard drive itself.

With the introduction of SSD, the system will now have online data (generated in the last few days) residing on the SSD, as well as data (older than a few days) residing on the conventional hard drive. The analysis showed that the data to be stored on SSD for at least one week comprises:

- Sampling process peak data output per week in Europe during November-December, 2009: **100 GB**.
- Number of monthly sampling jobs run during November-December, 2009: **1,400**.
- Expected monthly jobs for the next three years (obtained from business heads in Europe): **4,500**.
- Expected sampling job output data volume per week for 4,500 jobs: **100*4,500/1,400 = 320 GB (approximately)**.
- In addition to the output data, it is calculated that the total space required for all the tables and indices that the sampling process uses as input will have an approximate volume of: **119.5 GB**.
- Total data volume to reside on the SSD is: **320 + 119.5 = 440 GB (approximately)**.
- Therefore, a 500 GB SSD would suffice for the European proof of concept, before rolling out a permanent and more comprehensive solution to Europe, North America and Asia Pacific.

It was planned to initially move sampling data out of the SSD once a week. As the workload increases, the frequency of data moves can be increased to once or twice a day. This will eliminate the need to purchase additional SSDs in the near future.

SSD Vendor & Cost

Texas Memory System (TMS) and EMC were considered as potential vendors. EMC was interested in supplying SSD as a bundle with EMC conventional platter disk arrays, while Texas Memory System was ready to supply a 1 TB SSD within a short timeframe. For the purpose of the POC, a RamSan-620 SSD supplied by TMS was used, at a cost of USD $61,450 (plus shipping and handling).

EMC supplied the SSD for the final, more comprehensive system implementation as part of a 12 TB disk bundle costing USD $500,000.

Deployment Details

The movement of data from SSD to HDD is planned on a weekly basis over each weekend. Unused disk space is available for future growth.

It is also planned that with the increase in demand on the system, the movement of data can be accomplished more frequently. A daily move of data to HDD will effectively support five times higher job volume. This same technique can be used for the permanent solution in all three regions. It is also planned that after deployment of the permanent solution in Europe, the interim hardware can be used for staging and development environments.

In addition to the SSD deployment, several other hardware enhancements were completed for the European system, at a total cost of approximately $200,000. This included a database server upgrade, purchase of a new operating system (Linux) and new Oracle licenses. Details are shown below.

Name	Amount	Comments
2x Server 2 CPUx6 Cores	$20,991	
RAM SAN 620: 1 TB	$61,450	Taxes & shipping
Linux license	$1,000	Per year
Oracle database licensing for 8 additional CPUs	$113,000	Plus tax
Total	**$196,441**	

Appendix E: A Comparison Between SSD and HDD

Reliability: Because SSDs contain no moving parts, they last longer than hard drives. SSDs using flash memory can sustain almost 100,000 write cycles per write cell.

Read/Write Speeds: The typical access time for a flash-based SSD is about 35 to 100 microseconds, compared with 5,000-10,000 microseconds for a rotating disk. That makes a flash-based SSD approximately 100 times faster than a rotating disk. While hard drives typically transfer data at 80 MB/sec., SSDs typically transfer data at 170 MB/sec. Flash SSDs are slower than spinning disk for small file writes due to their large erase block size of 0.5 to 1 MB, but they have a constant read and write rate. At Cebit 2009, OCZ demonstrated a 1 TB flash SSD using a PCI Express x8 interface, achieving a maximum write speed of 654 MB/sec. and a maximum read speed of 712 MB/sec. Flash SSDs are faster than HDDs for small random reads due to their lack of moving parts and negligible seek time.

Noise: HDDs produce audible clicks and crunching sounds. SDDs are often quieter because they have no mechanical parts.

Size: Flash-based SSDs are manufactured in standard 2.5" and 3.5" form factors. 2.5" SSDs are normally used in laptops or notebooks, while the 3.5" form factors are used in desktops.

Vibration: SSDs are naturally more rugged than HDDs. SSD drives can sustain up to 1,500 Gs of shock before sustaining damage or a drop in performance, while HDD drives can withstand up to 350 Gs while operating and 800 Gs when turned off.

Reliability: SSDs can deliver high performance even in extreme conditions, such as military applications that require, in most cases, an operating temperature range of -60°C to +95°C. Shock, vibration and temperature ratings of HDDs cannot comply with military standards.

Power: SSDs consume less power than HDDs.

Heat Dissipation: Along with the lower power consumption, SSDs produce less heat due to the absence of rotating/movable media. This makes SSDs more attractive for mobile systems such as PDAs and notebooks.

Cost: SSDs typically cost $3 per gigabyte, while traditional hard drives cost about 20 to 30 cents per gigabyte.

Footnotes

[1] "New Configuration for Sun Storage 7110 System," System News, June 29, 2009, http://sun.systemnews.com/articles/137/1/OpenStorage/21799

[2] Oracle In-Memory Database Cache Introduction, Oracle, 2010, http://download.oracle.com/otn_hosted_doc/timesten/1121/doc/timesten.1121/e14261/toc.htm#

References

"IBM User's Guide, Thirteenth Edition," http://web.utk.edu/~mnewman/ibmguide03.html

Jim Cooke, "Flash Memory 101: An Introduction to NAND Flash," EE Times.com, March 20, 2006, http://www.commsdesign.com/showArticle.jhtml?articleID=183700957

Lucas Mearian, "Review: Hard Disk vs. Solid-State Drive – Is an SSD Worth the Money?" Computerworld.com, June 18, 2009, http://www.computerworld.com/s/article/9134468/Review_Hard_disk_vs._solid_state_drive_is_an_SSD_worth_the_money_

"Enterprise FLASH Drives for DMX-4 (Completed)," Storagezilla, Jan. 14, 2008, http://storagezilla.typepad.com/storagezilla/2008/01/enterprise-flas.html

Chris Mellor, "EMC has Changed Enterprise Disk Storage Forever: First into the Enterprise Flash Breech," Techworld.com, Jan. 14, 2008, http://features.techworld.com/storage/3928/emc-has-changed-enterprise-disk-storage-for-ever/

Barry Burke, "1.040: EFD: What's in a Name?" The Storage Anarchist, Feb. 18, 2009, http://thestorageanarchist.typepad.com/weblog/2009/02/1040-efd-whats-in-a-name.html

Woody Hutsell, "An In-Depth Look at the RamSan-620 Flash Solid State Disk," Texas Memory Systems, July 2009, http://www.texmemsys.com/files/f000272.pdf

Michael Ault with Jamon Bowen, Matt Key and Joe Bromley, "Oracle Performance Tuning with Solid State Disk, Oracle In-Focus Series," Rampant TechPress, 2009, http://www.texmemsys.com/files/oracle_performance_tuning_with_ssd.pdf

Oracle Exadata Database Machine, overview, http://www.oracle.com/technology/products/bi/db/exadata/pdf/exadata-storage-technical-overview.pdf

Douglas Dumitru, "Understanding Flash SSD Performance," EasyCo LLC, Aug. 16, 2007, http://managedflash.com/news/papers/easyco-flashperformance-art.pdf

About the Authors

Abhijit Ray is an Architect in the Enterprise Architecture Center of Excellence, Global Technology Office, at Cognizant. He has a master's degree and a PhD in engineering and is TOGAF certified. Abhijit is an experienced enterprise architect with wide-ranging experience at established blue chip companies, as well as broad exposure to various industry sectors, include government, finance, telecommunications, retail and information media. He can be reached at abhijit.ray@cognizant.com

About Cognizant

Cognizant (NASDAQ: CTSH) is a leading provider of information technology, consulting, and business process outsourcing services. Cognizant's single-minded passion is to dedicate our global technology and innovation know-how, our industry expertise and worldwide resources to working together with clients to make their businesses stronger. With over 50 global delivery centers and more than 100,000 employees as of as of December 1, 2010, we combine a unique global delivery model infused with a distinct culture of customer satisfaction. A member of the NASDAQ-100 Index and S&P 500 Index, Cognizant is a Forbes Global 2000 company and a member of the Fortune 1000 and is ranked among the top information technology companies in BusinessWeek's Hot Growth and Top 50 Performers listings.

Visit us online at www.cognizant.com for more information.

Cognizant™

World Headquarters
500 Frank W. Burr Blvd.
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233
Fax: +1 201 801 0243
Toll Free: +1 888 937 3277
Email: inquiry@cognizant.com

European Headquarters
Haymarket House
28-29 Haymarket
London SW1Y 4SP UK
Phone: +44 (0) 20 7321 4888
Fax: +44 (0) 20 7321 4890
Email: infouk@cognizant.com

India Operations Headquarters
#5/535, Old Mahabalipuram Road
Okkiyam Pettai, Thoraipakkam
Chennai, 600 096 India
Phone: +91 (0) 44 4209 6000
Fax: +91 (0) 44 4209 6060
Email: inquiryindia@cognizant.com