

A well known U.S. based financial institution asked Zettapoint, TriCore Solutions and EMC to conduct a Virtual Intelligent Storage Tiering Analysis of their current storage subsystem. They were so impressed with the results that they came back for more.

## Case Study

### Introduction

Corporate management at a well known U.S. based financial firm asked Zettapoint, TriCore Solutions and EMC to provide a data classification assessment to review and validate the organization's Oracle RAC infrastructure and performance as it relates to their Oracle database system supporting an important production database.

In order to better understand their data and business transactions in the "real world", TriCore consultants implemented the Zettapoint DBclassify™ software for a three week data collection period on the target production database instance selected by the customer.

#### The Right-Tiering Concept: Higher Performance and Lower Cost

*Right-Tiering is where data is classified by actual business usage and value, and then placed on the appropriate storage tier to consistently deliver the highest possible service at the lowest possible cost.* Right-Tiering helps optimize the placement of data on the right devices, on the right tier and at the right cost. With Right-Tiering, customers realize significant cost savings while maintaining or improving end-user performance goals. In general, Right-Tiering provides two main business benefits:

1. Improved SLA compliance for business users by significantly reducing I/O wait time for the most critical applications.
2. Reduced TCO by placing the right objects on the right tier based on their usage patterns (e.g. highly used "hot" data on fast devices; seldom used "cold" data on slower devices, etc.).

### Methodology

The EMC Virtual Intelligent Storage Tiering Analysis methodology focuses on the delivery of storage infrastructure Right-Tiering and Oracle best practice performance recommendations. The toolset and methodology employ an iterative process that captures business utilization patterns over time for new and existing data. The process is subdivided into six phases: Data Collection, Discovery, Analysis, Draft, Review, and Remediation Roadmap.

### Understanding the Strategic Initiative

A pressing strategic initiative for the organization is the migration of their Oracle databases from the current EMC DMX Fibre Channel drives to more modern Enterprise Flash Drive (EFD) technology. They believe this will help them meet their Service Level Agreements (SLA), address rapidly growing capacity demands and respond to increasing performance demands from the business. Prior to this migration, the company wanted to better understand how to leverage EMC EFD technology, as well as identify

*The company wanted to better understand how to leverage EMC EFD technology, as well as identify target database objects and files as candidates for a database consolidation/migration to EFDs.*

target database objects and files as candidates for a database consolidation/migration to EFDs. The company is interested in proactively understanding those areas in their database that are “hot” with high business impact for maximizing their performance while implementing total cost of ownership (TCO) design considerations with storage “Right-Tiering”.

### **Preliminary Results**

The analysis showed the following preliminary results:

- Total Capacity of Usable Database Storage was equal to 4.7 TB
- They require 246 GB Flash (Tier 0), 1.8 TB (Mid-Tier 1), 458 GB Slow/Archive (Tier 2)\*
- They suffered 4500 Hours of Total Wait time with 59% attributed to I/O Wait alone

\* These are minimum numbers that were subjected to further analysis through DBclassify filters and capacity growth assessments over time to determine more precisely the size of the programs representing the “hot” objects.

### **General Database Capacity and I/O Wait Analysis**

The DBclassify Capacity Dashboard view showed the total capacity of active data is 3.98 TB with a total of 729 GB of inactive data (15%). Further investigation of the inactive data showed a total of 458 GB (i.e. ~10% dormant objects) data objects such as tables, indexes and partitions, with the remaining inactive data as Oracle tablespace “free space” used to support data volume growth. Dormant data may be archived to slower, cheaper tiers of storage, such as large SATA drives in a RAID5 configuration or the company may wish to maintain an older storage class technology to be used as the archival tier of storage for their database.

The DBclassify I/O wait analysis showed peak wait times with their associated I/O wait peaks for all SQL statements across the selected database. During the entire collection period, 4500 hours of SQL waits occurred on the selected database with 59% (2663 hours) due to I/O wait alone. Further DBclassify analysis of the Top 5 wait types, SQL statements and Database Objects during the collection period showed that the largest waits are Oracle I/O waits that comprise approximately 39% of waits including RMAN, SBT Backup I/O and DB Sequential Read waits.

### **“Hot” Data Object Analysis: Candidates for Tier 0**

DBclassify Tablespace reports were produced at the tablespace, datafile and object level to help identify candidates to be placed on Tier 0 storage to alleviate I/O waits and improve overall performance without requiring any application changes. Estimated performance time savings for those files and objects were provided using published Oracle database EMC Enterprise Flash Drive response time metrics<sup>1</sup>.

---

<sup>1</sup> “Implementing EMC Symmetrix DMX-4 Flash Drives with Oracle Databases”, EMC Corporation’s Applied Technology Division, September 2008.

| Tablespace Groups                     | Tier Zero Time Saved (Hours) | Tier Zero Storage Target (TB) |
|---------------------------------------|------------------------------|-------------------------------|
| Tablespaces with > 8ms response times | 1227 – 1432 Hours            | 2.4                           |
| Heavy read/write tablespaces          | 899 – 1050 Hours             | 1.56                          |
| Totals:                               | 2126 – 2482 Hours (80-93%)   | 3.96                          |

Table of Tier 0 Tablespace time savings and EFD storage requirements

The DBclassify Object Analysis capabilities allowed the analysis team to take a closer look at the object level (Table, Index, and Partition) waits to determine if they could improve performance gains at the object level and reduce the final EFD footprint and cost. As it turned out, the top ten objects consumed approximately 17% of all I/O waits or 456 hours of wait time out of a total of 2663 hours during the data collection period. This led the analysis team to a powerful recommendation to get this object list (~246 GB\*\*) reorganized over to the EFD Tier 0 storage group. The reorganization of these objects would drive significant performance gains as the I/O behavior is predominantly Single Block I/O (i.e. 401 hours).

#### “Warm” Data Object Analysis: Candidates for Tier 1 or Tier 2

DBclassify Table Usage features and drill down capabilities helped identify those objects that fall into normal data utilization patterns. These were objects found to have less than 1% execution time across all DML for the duration of the data collection period. DBclassify filtering capabilities were used to query “< 1%” executions. This helped drive the decision to take these objects, totaling approximately 1.8 TB in size, and reorganize them onto Tier 1 (e.g. 15K RPM Fibre Channel HDD) or Tier 2 (e.g. 10K RPM Fibre Channel HDD) mid-tier storage infrastructure.

#### “Cold” Data Object Analysis: Candidates for Archival

Standard DBclassify reports were used to help identify dormant objects (i.e. Tables, Indexes, and Partitions) that are candidates for a lower storage tier or even off line archival. From the list of objects that were found to have no activity during the measurement interval, it was determined that:

- **Dormant Tables** - Recommendation to relocate these tables (128 GB) to cheaper large SATA drives in a RAID5 fault tolerant configuration.
- **Dormant Indexes** - Recommendation for these indexes is for them to be reorganized to a separate tablespace(s) on a cheaper storage tier consisting of SATA drives with RAID5 fault tolerance levels. Further analysis is required to positively identify if these indexes map back to active base tables. If the base tables are found in the frequently active list of objects, the company should consider capturing the Index DDL to recreate these indexes, and permanently drop these indexes (105 GB) from the database to minimize I/O during DML operations (Inserts, Updates, Deletes) to their respective base tables.
- **Dormant Partitions** - Partitioning is moderately leveraged in this database with 225 GB of dormant partitions that could be re-located to a lower tier of storage.

## Other Considerations: Query Performance and Applications

In order to properly “Right-Tier” the database with the application patterns, SQL execution times were also reviewed in this process. High execution SQL or SQL with high I/O waits indicated high resource consumption that would benefit from having Tier 0 placement. Standard DBclassify features allowed the architect to sort production SQL by highest executions and highest I/O to identify the subset of queries that would benefit from Right-Tiering, Data Model adjustments or SQL tuning.

DBclassify can also be used to track database utilization mapped back to an application, set of program(s) or business process(es). From discovery sessions with the company’s administrative team and further analysis, the analysis team identified the programs having potentially significant I/O in the database. DBclassify filters were used that associate all SQL activity and I/O patterns as they pertain to the respective program processes. By filtering, noise from other SQL queries was eliminated and a tight focus was placed on critical areas of fine tuning these programs for the business. Filters result in the ability to create a manageable and auditable end-user experience that drives high performance under a repeatable, managed framework such as a SLA. By Right-Tiering onto EFD as Tier 0, it provides an immediate return to these critical areas of the database while minimizing the need for code changes.

## Tier 0 Return on Investment Analysis

Values were taken from DBclassify reports for Tier 0, Tier 1 and Tier 2 over a four year infrastructure plan with Right-Tiering using industry benchmarks for the price per tier of: Tier 0: \$60k/yr, Tier 1: \$20k/yr, and Tier 2: \$5k/yr. ROI calculations assume the following over a four year plan:

- 10% dormant data
- 38% mid/low range
- 52% Tier 0 range
- 20% annual data growth

The investment to place the entire database on Tier 0 over a four year plan would be approximately \$1.5 million. By leveraging the DBclassify Right-Tiering strategy, an estimated 28% return on investment of \$427K can be realized over the same four year period and includes a much more efficient use of this expensive resource.

## Conclusions

After careful review of the Zettapoint DBclassify reports, the analysis team determined that there would be significant performance gains and savings by implementing “Right-Tiering” in the company’s selected database environment. Over a short window of data collection of a few weeks, 2,663 hours of waits (59% I/O bound combined waits) were visible that would benefit from EFD. Based on the data classification analysis, the team provided a roadmap to the company for the next steps to implement the necessary changes.



## Epilogue

As a result of this Intelligent Storage Tiering Analysis, members of the original TriCore Solutions analysis team along with members of the EMC Solutions team were invited back to the company for a Proof of Concept demonstration utilizing an EMC V-MAX configured with two engines, 64GB cache, 32 400GB Enterprise Flash Drives, 8 4GB Fibre Channel Director ports, 160 LUNs, and 4X40 Raid Groups using RAID5. DBclassify has become a regular part of the company's ongoing Right-Tiering initiative.

Ultimately, the company purchased a new EMC VMAX and asked EMC to provide Data Layout Services to leverage the information gleaned from DBclassify reports to more precisely determine the objects that should be reorganized onto Tier 0.

The company recently configured their SAN to map to the Data Layout, reorganized their production database to match the tiering recommendations, and are now migrating their database to the new and newly acquired technology.

### For more information, contact:

Zettapoint. Inc.  
45 Walden St.  
Concord, MA 01742

Phone: 978.369.9900  
[info@zettapoint.com](mailto:info@zettapoint.com)

[www.zettapoint.com](http://www.zettapoint.com)

