An Oracle White Paper
January 2011

Using Oracle's StorageTek Search Accelerator

ORACLE

Executive Summary....2
Introduction....2
The Problem with Searching Large Data Sets....3
The StorageTek Search Accelerator Solution....3
StorageTek Search Accelerator Implementation....4
Example Using a Basic Grep Test....4
Example Comparing SSA to IBM Hardware Assisted Search....5
Conclusion....6

Executive Summary

Oracle's StorageTek T10000C tape drive is the first to offer the StorageTek Search Accelerator (SSA). SSA uses tape drive hardware to speed data searches, and offload that processing requirement from other data center resources.

Introduction

In today's world of e-discovery, data transformation, encryption and large capacity tape cartridges it is becoming more and more important to improve the searchability of tape.

To improve data accessibility and search for tape storage, Oracle is embedding a hardware search capability in the StorageTek T10000C tape drive that allows applications to offload search to the drive. Using this feature can ensure expensive compute and storage resources remain dedicated to critical business needs.

The Problem with Searching Large Data Sets

Without the right tools, it is a daunting task to find specific files, or individual records, on tape cartridges with large capacities. Many products are offered to solve this problem. Typically these solutions read an entire tape across an interface and create an index for efficient search operations. These indices and even tape records are stored on disk and searching them is processor intensive and time consuming. The applications that specialize in these search functions are expensive. They often require constant update as formats change and new digital applications are created. There are also Hierarchical Storage Management (HSM) audit solutions enhancing tape search with specialized hardware on proprietary equipment. This certainly increases performance, but it can come at a high system-level cost, usually with a vendor specific implementation. Because of the cost and resources needed to perform search operations many enterprises obtain this capability only after an event like a lawsuit or disaster occurs. The data on tape is seldom, if ever, searched once it on tape.

The StorageTek Search Accelerator Solution

To solve search problems associated with tape, Oracle is offering the StorageTek Search Accelerator (SSA) on the StorageTek T10000C tape drive. All tape drives use digital logic to check and generate format specific Cyclic Redundancy Check (CRC) or other data check information. Oracle expanded this existing capability to support searching for user provided strings. This search is performed after the data records have been decrypted and decompressed, so it is performed on the original records as sent to the tape drive. SSA allows any application to search data records on any StorageTek T10000C written tape cartridge, and return only those meeting specific match criteria. When this feature is enabled there is no performance loss. In fact, depending on system configuration, there might be a small increase in performance.

To use SSA the application provides a binary string, and the tape drive returns only those records containing a match to that string. The search string can represent names, words, numbers, labels or any marker that the application stored in a record. More than one binary string can be provided, and the search length is bounded by a record count, reaching a File Mark or End of Data (EOD). An offset is also provided to the drive to specify where to begin the search in each record.

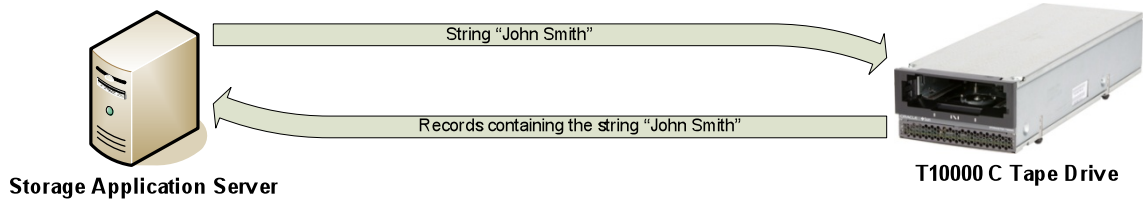

Figure 1: A basic search

In Figure 1, the application needs all the records that contain the binary string representing the name "John Smith". The binary string "4a6f686e7f536d697468" is sent to the drive. A record offset is not

specified, therefore each record is searched from beginning to end, and the search length is limited by the EOD. In this example, the tape drive returns all the records containing the string "John Smith".

In another example, a marker pattern "fffe23457edfafffff000abab" identifies metadata records, and is located at offset 1000. These metadata records must be checked to audit data sets on a tape. This is shown in Figure 2.

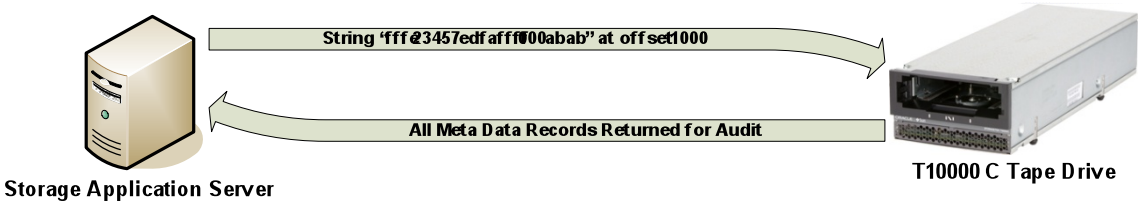

Figure 2: Metadata Record Search

In this search, the drive looks only at offset 1000 for an exact string match. This search would be much faster than the basic search shown in Figure 1. Again the search starts at the beginning of tape and ends at the partition EOD. All the metadata records for this tape would be returned to the application and could be used to audit the tape or perform some other service.

StorageTek Search Accelerator Implementation

Like other Oracle tape innovations (StorageTek Tape Tiering Accelerator and StorageTek Data Integrity Validation), SSA has been designed for ease of use and flexibility. A set of vendor specific SCSI commands support this function in our fibre channel interface and an API is available that uses a C Library supporting these search functions.

SSA supports a binary search for a single string, up to 1024 bytes in length. The search operation can match as many as two search strings with a combined length of 2048 bytes. The search can begin at any record. It supports starting the search at the beginning of each record (offset 0) or at a specified offset from the beginning of each record. The search completes at the end of data; when a tape mark is encountered; or a when the set limit number of records have been searched. Only records matching the search criteria are returned to the application. It is important to note that the tape drive's processor and hardware are used to perform the search off line. This ensures critical business processing, storage and SAN resources are not burdened with records that don't contain the search target string.

Example Using a Basic Grep Test

To perform a basic test we created a data set of approximately 20 GBs on disk and wrote it to tape with a StorageTek T10000C tape drive. The target search string was located at the very end of the data set. A search was conducted for the target string on tape, and a "grep" (the UNIX standard command line text search utility) was performed on the original data set on disk.

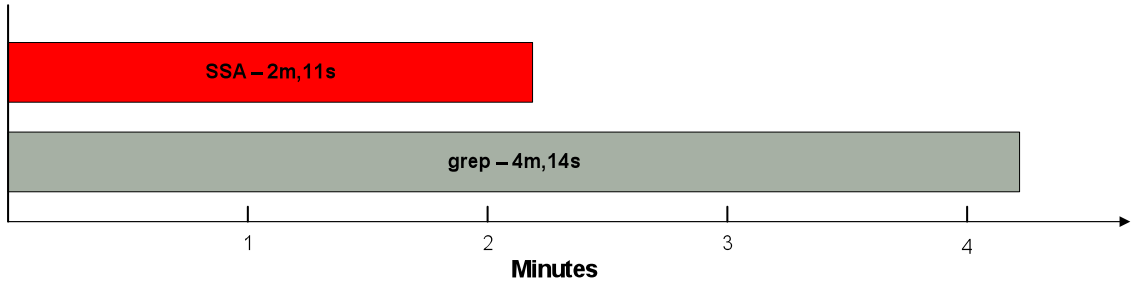

Figure 3: Grep Test Results

The SSA search took 2 minutes and 11 seconds to return the file containing the target string. The disk "grep" completed in 4 minutes and 14 seconds.

Of course "grep" is not an optimized search application, so performing twice as fast as "grep" is no big deal. However, this test does show that SSA is capable of searching at ~153 MB/s. Most importantly, several minutes of processor bandwidth and 20 GBs of storage were not needed find a specific tape record.

Example Comparing SSA to IBM Hardware Assisted Search

To perform a more targeted test of SSA we compared it to IBM Hardware Assisted Search, as provided with the TS1130 controller for HSM audit.

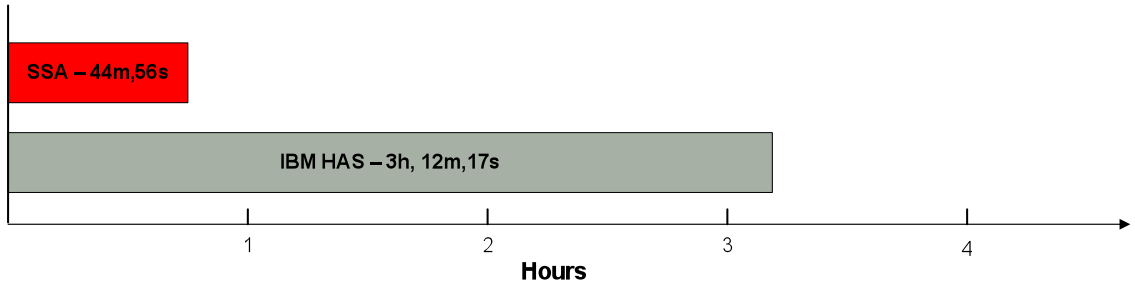

Figure 4: HSM Audit Test Results

In this test, both audits were performed against HSM migration tapes that contained 720 identical datasets. The audit looks for a specific signature that identifies a Control Data Set (CDS) record. The CDS records are then checked to ensure the HSM records are consistent, and up to date. The HSM audit took 44 minutes and 56 seconds using SSA. With HAS, the audit completed in 3 hours, 12 minutes and 17 seconds.

For SSA, a binary search string was provided to find records containing the CDS signature. Only CDS records were returned to the HSM audit application and the audit was completed very efficiently.

For the HAS test, all 720 datasets had to be read into the IBM controller where HAS was used to find records with the CDS signature. These CDS records were then provided to the HSM audit application to complete the audit.

Conclusion

Oracle's StorageTek Search Accelerator provides a new way for eDiscovery, and other applications, to search tape without using critical IT resources, and is only available on the StorageTek T10000C tape drive. This feature, and other Oracle innovations, is designed to support new methodologies and redefine tape usage. With the StorageTek T10000C tape drive, Oracle is redefining tape storage.

Oracle is committed to developing practices and products that help protect the environment

Using Oracle's StorageTek Search Accelerator
January 2011
Author: Dwayne Edling

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0111

Hardware and Software, Engineered to Work Together