

Controlling data growth and system performance through an Information Lifecycle Management (ILM)–oriented archiving approach with HP StorageWorks Reference Information Manager for Databases white paper



- Introduction 2
- Tektronix, Inc. overview 2
- The problem: Data growth affects performance 2
- The solution: Purging 3
- Proof of concept project..... 4
 - Test environment 4
 - Tests 4
 - Results 5
- Proof of concept conclusions 6
- Archive implementation objectives 6
- Data retention issues 7
- Defining a data retention policy 8
- Archive implementation 9
- Conclusion..... 9
- For more information..... 10

Introduction

Between 1996 and 1999 Tektronix successfully implemented Oracle® Applications on a single instance for its entire business spanning 23 countries. Tektronix has since added four more countries for a total of 34 legal entities.

This paper discusses some of the issues generated by this Oracle Applications implementation and the steps taken to mitigate the challenges presented by growing application data. The paper primarily focuses on issues related to data growth, data availability, system performance, and system availability in a 24 x 7 environment.

To address these growth-related challenges, Tektronix implemented a worldwide data archive solution from HP in June of 2002. This archive solution proved invaluable in its efforts to lessen the burden of data growth, mitigate performance problems in the company's OLTP system, and still satisfy worldwide financial and reporting requirements.

Based on Tektronix's initial proof of concept testing and current experience, the top data growth areas for the company are AR, Order Entry, Inventory Transactions, Cost, Price List Lines, and GL. In addition to these vanilla modules, Tektronix addressed several customizations.

The issues involved and the decision process used to help determine data retention periods, archive and purge cycles, and selecting a standard archive solution are explained.

Tektronix, Inc. overview

Tektronix, Inc. is a technology company focused on providing test, measurement, and monitoring solutions to enable and sustain the explosive growth of global communications networks and Internet technologies. With headquarters in Beaverton, Oregon, Tektronix maintains operations in 27 countries outside the United States. Founded in 1946, the company's measurement business had revenues of \$150 million in fiscal 2002.

A leader in its target markets, Tektronix is universally recognized for high-quality products, technology innovations, and partnership with customers. The company holds more than 1,000 patents worldwide and receives approximately 60 more each year. Tektronix is dedicated to leveraging this vast store of technology and expertise to deliver products that allow its customers to design microprocessors and computers and deploy communications systems that drive the ongoing revolution in global communications networking.

The problem: Data growth affects performance

When Tektronix went live on Oracle in 1999, its production database size was 60 GB. Data growth quickly surged to a rate of about 1.25 GB per month. Despite tuning exercises and hardware upgrades, the increased growth rate caused performance to decline. Run times for batch programs increased despite fewer numbers of executions.

Figure 1 shows the steady growth of space allocated for data and indexes in the database from August 2000, through December 2000. Figure 2 shows the cumulative run times in minutes increasing and execution count for all concurrent programs decreasing slightly for the same period.

Figure 1.

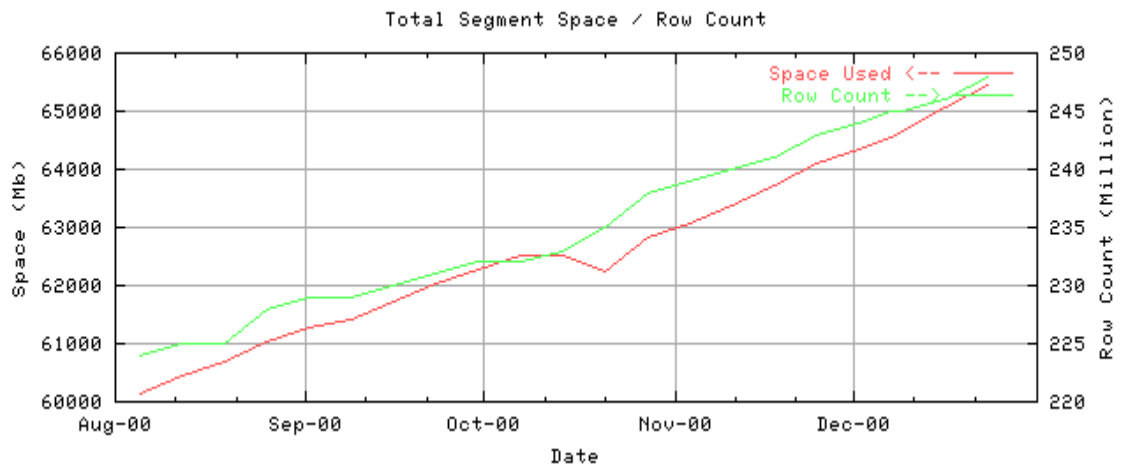
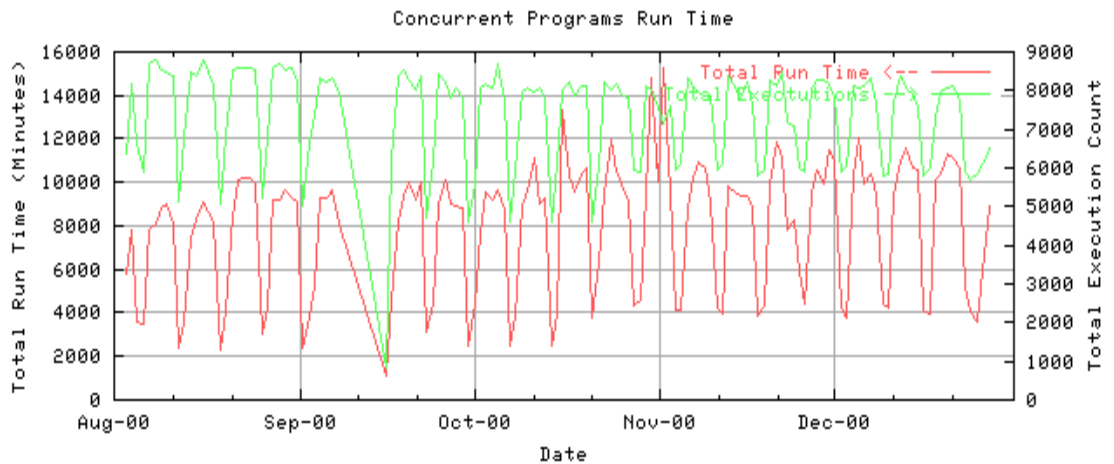


Figure 2.



The solution: Purging

With performance tuning efforts proving fruitless, Tektronix decided to address the root cause of declining performance by focusing on data growth. Purges seemed the best option for reducing the growing data footprint. However, because of international financial reporting regulations, Tektronix could not implement a purge strategy without also implementing some form of archiving.

Initially, Tektronix considered developing its own solution internally. However, after discovering HP during the Spring 2000 OAUG, Tektronix determined the HP StorageWorks Reference Information Manager for Databases (RIM for Databases) provided an elegant and simple approach to this difficult problem. RIM for Databases was already shipping with full support for Oracle Financials and Manufacturing and provided all the capabilities Tektronix deemed necessary, including:

- Support for standard Oracle purge programs
- A clean extension of Oracle Applications
- Support for customizations
- Online archive access using the same user interface as production
- An upgrade path for archived data that always stays in-sync with production
- No additional end-user training
- Minimal IT overhead for archive support
- No impact on business processes
- Application-centric approach
- Proven in real-world implementations

Despite the track record of success for HP, there were still some internal skeptics who remained unconvinced that purging would actually help improve performance. To help convince the disbelievers and validate the proposed archiving project, Tektronix performed a proof of concept project over a three-month period. The project endorsed the proposed archiving strategy and showed space savings and dramatic performance improvements directly linked to purging.

Proof of concept project

The proof of concept plan was fairly straightforward:

1. Create a test environment
2. Install the RIM for Databases solution
3. Design and execute tests

Test environment

RIM for Databases was installed in the test environment on a recent copy of the production database. The installation of the archive solution including several extensions to Oracle's vanilla purges to accommodate custom data and the complete development of one custom purge took less than two weeks.

Tests

Tektronix chose three purges for the tests that would exercise the purges and the archiving solution and could be used to demonstrate performance improvements in commonly used programs.

- The easy inventory transactions purge
- The complex sales order purge
- A custom purge for the company's price list lines extension

To test the affects of purging on performance, key batch programs were identified that met the following criteria:

- Run on a regular basis
- Heavily used by the business
- Reference the tables addressed by the purges

Before purging, all of the reports were run in a test environment and the run times were recorded. Purges were run that addressed the following data.

Purge	Period	Remarks
Inventory Transactions	<= Dec 31 1998	Approximately 50% of all rows
Sales Orders	All eligible <= Dec 31 1998	Approximately 63% of all rows
Price List Lines	Expiration Date <= Dec 31 2000	

The DBA group was then tasked with reclaiming the space that had been freed up by the purges. The tables that were purged from were de-fragmented and their indexes were rebuilt.

The same reports were run for the final time after the data and indexes had been cleaned up. These reports provided the run times to measure the performance benefits to be realized as a result of purging.

Results

As can be seen from the following results, the average performance improvement was over 47%, with one report improving by 73%.

Data Purged

Purge	Rows Purged Millions	%	Space Saved GB	%
Inventory Transactions	8.2	52	2.6	53
Sales Orders	6.7	63	1.4	58
Price List Lines	1.7	17	2.7	67

Report	Pre-purge run time	Post-purge run time	Improvement
ICA Open Receiving Transactions	0:01:31	0:00:58	37%
Daily Transaction Register	0:15:23	0:08:12	47%
Closed Orders	0:14:29	0:06:43	54%
Unbooked Orders	1:53:28	0:43:07	62%
Hold Source Activity	0:15:21	0:04:10	73%
Price List Extract	0:03:31	0:03:06	12%

Proof of concept conclusions

The proof of concept and testing provided convincing evidence of the efficacy of the HP RIM for Databases solution. In particular, it established:

- RIM for Databases is a robust, reliable product that solves Tektronix’s archiving issue.
- The solution can be used to archive custom data.
- The vanilla purge programs could be extended to perform additional functions.
- The archived data was available in the archive instance without data loss.
- There were measurable performance improvements.
- Significant space savings could be realized.

Based on the proof of concept results, Tektronix elected to proceed with a full RIM for Databases implementation. To have the fastest impact on escalating application data growth, Tektronix prioritized the implementation of the Oracle Application archives based on their current space usage.

AR Purge	11.7 GB
Sales Orders	8.9 GB
Inventory Transaction	6.2 GB (2.4 GB custom)
Cost	5.3 GB
Price Line Lists	4.9 GB
GL	4.0 GB

Archive implementation objectives

With the effectiveness of the HP archiving solution confirmed during the proof of concept, Tektronix sought to establish a framework and set of objectives to guide the implementation efforts. At the highest level, the company wanted the production archive to deliver the same high-quality results exhibited in testing. In particular, Tektronix wanted to see a marked improvement in performance with a corresponding reduction in storage requirements. The solution needed to stabilize both performance and storage usage.

Beyond the basic demands for performance and reduced infrastructure costs, the company wanted the data archive to:

- Satisfy worldwide financial and legal reporting requirements.
- Identify the costs of maintaining portfolio company data and isolate those costs from Tektronix’s ongoing cost of operations.
- Create an ongoing, production-strength archive process that affords the company the ability to maximize the efficiency of its core production databases and deliver:
 - Increased efficiency of user-run inquiries
 - Reduced time to back up production databases
 - Reduced costs for storing multiple images of production
 - Reduced time to copy multiple images of production
- Provide users with an effective means for accessing the required data elements with minimal use of new/additional access tools and ongoing support costs.
- Ensure audit and reconciliation checks are performed before any data is purged from the company’s production instances.

Data retention issues

The obvious downside to removing data from your OLTP database is that business users will no longer have access to it. Implementing an archiving solution would help ensure that the transactional data was not lost forever but could still be accessed when required. The burden of data retention has grown significantly as government regulations, including the Sarbanes-Oxley Act of 2002, require longer holding periods and more stringent control of historical records. These new requirements are in response to recent financial and accounting improprieties and the ensuing litigation.

These new regulations impose three levels of compliance on the corporation. First, organizations are required to retain a wider range of financial records, including supporting documentation. Second, stricter management oversight demands a centralized, single instance of all financial data. And, third, the mandated reporting periods have been shortened, demanding greater performance and consistency from the financial system. Each of these regulations requires optimal performance from the production financial system as well as greater capacity and accessibility from the archive.

Currently, there are several options for archiving data, including:

- Paper
- CD
- Microfiche
- Disk
- Data warehouse

During archive planning and strategy discussions with the business users and IT staff, Tektronix quickly determined that an online archive solution was the most practical for all parties. This solution would allow the implementation of an aggressive purge schedule that would maximize performance benefits and disk savings, while allowing the business users instant access to the historical data without assistance from the IT groups. The flexibility gained by leaving the data intact and online outweighs the cost savings the company might realize by moving the data to lower cost offline medium such as reports on a CD.

Defining a data retention policy

The company's data or record retention policy is the foundation for the categorization and management of the company's data or information assets. Tektronix believes that data is an information asset and that all information assets have a useful life. Assets can become expenses and liabilities after their useful life expires. Information assets have a three-phase lifecycle that includes Current State, Online Archive, and Final Form. An example of categorization within the information asset life cycle is represented in the following table.

Current State	Online Archive	Final Form Retention
Current Operation	-2 year	-7 year and beyond
-1 year	-3 year	
	-4 year	
	-5 year	
	-6 year	
Online production transaction processing	Read only, no transaction processing	Read only, no transaction processing
Release dependent	Release compatible Complete copy of what was in current state	Independent of production versions/releases
No long-term history	Less costly form of storage	Summarized data, reports, specific tables, and so on stored offline on CDs, fiche, delimited files, and so on
Most costly form of storage Administered by IT	Administered by IT	Least costly form of storage Administered by functional data owners, not IT

When developing the archive plan, Tektronix established several data retention guidelines to govern the execution of its biannual archive processes. These guidelines established the following framework:

- Financial and order data would be archived online for a maximum of six years. During that period, Tektronix has extensive legal and financial reporting obligations throughout the world.
- After the six-year retention period, transaction data is removed from the company's transaction processing and archive systems. There are few countries where Tektronix does business that require detailed transaction data beyond six years.
- Data retention beyond six years becomes the responsibility of the functional data owner. The cost of the final retention format—paper, CD, or microfiche—is borne by this data owner.
- For most permanent retention demands, paper is the preferred format.

Archive implementation

In June of 2002, Tektronix moved HP RIM for Databases to production. The Order Entry archive was run in June of that year with the first Accounts Receivable archive run four months later. Since then, Cost Records has moved to production and Order Entry and Accounts Receivable have each run an additional archiving process.

The initial Order Entry and Accounts Receivable archive operations yielded results in line with those demonstrated in the proof of concept. Initial database space for AR was reduced by 48% with a corresponding 40% improvement in performance. Likewise, Order Entry saw a 56% reduction in space and a 52% initial performance improvement.

Data Purged

Purge	Rows Purged Millions	%	Space Saved GB	%
Accounts Receivable	8.2		7.2	48
Order Entry	8.2		7.2	56

Conclusion

The initial proof of concept made a clear argument for the validity of an application data management solution with the relocation of historical data into an online archive. The subsequent results of the HP RIM for Databases implementation and archive operations further reinforced this conclusion.

Based on Tektronix's firsthand experience, the company has determined:

- Purging data improves system performance
- Purging data reduces database server resource consumption by:
 - Reducing hardware upgrade costs
 - Consistently executing batch jobs on time
 - Improving online response time
 - Reducing disk growth
- RIM for Databases reduced the cost and time for additional development and test environments.
- RIM for Databases reduced backup and copy time as well as disaster recovery time.
- RIM for Databases acts as an extension of Oracle Applications to eliminate the need for additional end-user training.
- HP provided excellent partnership and support.

For more information

For more information on HP StorageWorks Reference Information Manager for Databases, visit:

www.hp.com/go/rim4db

© 2006 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Oracle is a registered US trademark of Oracle Corporation, Redwood City, California.

4AA0-7801ENW, September 2006

Tektronix[®]
Enabling Innovation

