



Cost-Efficient Data Management of
Large OLTP and ERP Environments:
An Integrated Approach to
Information Lifecycle Management

Ixion, Inc.
10805 Sunset Office Drive
Suite 300
St. Louis, MO 63127
314-725-0200
info@ixionsoftware.com
www.ixionsoftware.com



Information Lifecycle Management

The concept of information lifecycle management (ILM) is nothing new. However, the importance of a definitive plan for managing the efficient and accessible storage of data has become more and more urgent. Companies have had to face the challenge of numerous compliance and information security regulations that require data to be retained for longer periods of time than originally budgeted for. Regulations such as Sarbanes Oxley, HIPAA, Gramm-Leach-Bliley, etc. have forced companies to reevaluate their budgeted storage requirements forecasts, and on average have had to nearly double these estimates.

Most companies are still in need of a method for cost-effective management of this ever-growing amount of data. In most cases, data is stored in "on-line" memory that is easily accessible by the source applications that created it. On-line storage is perfectly suited for data that is accessed on a regular basis. But studies show that 80% of all stored on-line relational data is not utilized after 10 days. To maintain this type of inactive data in on-line storage on a primary database instance is very costly and unnecessary. Ultimately, this growth in data volume will have an adverse affect on not only the performance of the applications, but on the ability to provide adequate disaster recovery, and timely nightly batch processing.

Many companies are beginning to realize the added cost and risk involved with having too much data in on-line storage and storage area networks (SAN), and will soon have to take corrective action to manage storage in a more efficient manner or face spiraling storage costs. In environments where data volumes and data updates are very dynamic, such as OLTP or ERP applications, it makes sense to incorporate configurable, automated processes for monitoring the lifecycle of data in these systems: This type of monitoring should include audits of how data is created, modified, and by whom; the age of data; and procedures for purging or archiving certain types of information to less expensive storage.

There are numerous options for companies to utilize physical storage, or SAN's, to meet various data retention requirements. Companies can select from EMC, Veritas, Hewlitt-Packard, Hitachi, and many others. However, incremental storage costs can run into the millions of dollars per year for companies that create a terabyte or more of additional data per year. In addition, there are only so many I/O slots that can accommodate expansion before further hardware upgrades are necessary. Currently, there are only a small - but growing - number of storage and software vendors that are beginning to offer configurable, automated software solutions that manage data intelligently throughout the data lifecycle and minimize SAN growth.

©2005 Ixion, Inc.

10805 Sunset Office Drive • Suite 300 • St. Louis, MO 63127 • 314-725-0200
Info@ixionsoftware.com



Automated Data Lifecycle Management - *without* Performance Degradation

There is valid concern that any storage management processes that are “layered” on top of what may be an already-overburdened OLTP or ERP application can only result in additional performance issues. Therefore, there is a real need for a solution that will not impact mission-critical applications.

Some storage companies and a few small software companies are just now beginning to look at relational databases and developing automated processes to facilitate data lifecycle management to minimize reliance on expensive on-line storage and maximize the use of cheaper third tier, “near-line” storage alternatives such as ATA or WORM (write once, read many) storage media. The key is in providing an integrated solution that has minimal impact on your applications. Ixion’s industry-proven data capture, transmission, and apply technology for High Availability and Disaster Recovery provides a critical “first step” in providing just such a solution. Ixion’s ILM solution, known as IxILM, incorporates this technology to deliver the industry’s first integrated and automated high-performance ILM solution for Oracle and other RDBMS’s.

How can Ixion provide such a solution? To begin with, there is no product currently on the market that can replicate Oracle-based data faster than Ixion’s IxPropagator. IxPropagator has been proven time and time again in some of the largest, most complex and most challenging Oracle environments in the world. Low-overhead, high-speed data transmission is a critical first step toward achieving a fully-integrated ILM solution within a large ERP or OLTP environment. Concerns related to performance and speed of a “layered” ILM process that might slow a mission-critical application would be unfounded with IxPropagator. With IxPropagator it becomes possible for a configurable and automated ILM archival architecture to reside not on the primary database instance, but on a replica instance that is at all times fully-synchronized with the primary instance in real-time regardless of peak volumes or any batch processing windows.

The combination of real-time replication with data archival will enable Ixion to provide the only fully-integrated ILM solution for Oracle that both maximizes performance of the primary instance and allows for archival processing logic to reside on the secondary database instance. This secondary instance could be fully dedicated to data archival or it could be more efficiently-utilized for read-only operations such as reporting or warehouse feeds. Further, with the IxPropagator Disaster Recovery features available, it can also serve as a disaster recovery site in

©2005 Ixion, Inc.

10805 Sunset Office Drive • Suite 300 • St. Louis, MO 63127 • 314-725-0200
Info@ixionsoftware.com



the event of unplanned outages and a secondary processing site during planned maintenance of the primary. Such maximum utilization of the secondary database instance will allow companies to avoid additional investments in yet more redundant hardware and Oracle licenses to accomplish their data archival requirements.

Overview of IxILM - the Ixion ILM Solution

The Primary Database Instance

IxPropagator, Ixion's proven data replication solution for Oracle, is used to replicate data in real-time from the primary database instance to a secondary instance, which may not only be used for ILM, but also for Disaster Recovery, Reporting, and other read-only processes that may overburden the primary instance. Once the data is replicated, Ixion utilizes IxPurge, a user-configurable purging program, against the primary instance. IxPurge runs outside of the replication process to automatically delete selected records from the primary instance at scheduled times. Note that the deletions are not replicated to the secondary instance. The data that is selected for purging will be completely based on rules input during the configuration for the application. It may be based on the type of data, application source, age of the data, and periods of inactivity. This purging will allow the primary instance to maintain only current, active data, which will provide for the most efficient use of storage and processing capacity.

The Secondary Database Instance

The secondary database instance can serve multiple purposes. First, it is a full-replica of the primary instance and can be utilized for disaster recovery. It is also open for read-only operations and can therefore be used to off-load reporting applications or data warehouse feeds that would otherwise over-burden the primary instance and conflict with the primary database's main purpose of hosting the application that creates the data. Finally, the secondary instance can now be used for ILM archival purposes. Because of the integrated approach to replicating the data for not only disaster recovery and read-only access, but also ILM, customers gain the advantage of a fully utilized database instance and hardware platform, and avoid costly investments in unnecessary and redundant software and hardware.

The secondary instance utilizes Ixion's partition management solution, IxPartitionMgr. The partitioning logic is rules-based, depending on configurable parameters set by the customer. Parameters may include the age of the data, data

©2005 Ixion, Inc.

10805 Sunset Office Drive • Suite 300 • St. Louis, MO 63127 • 314-725-0200
Info@ixionsoftware.com



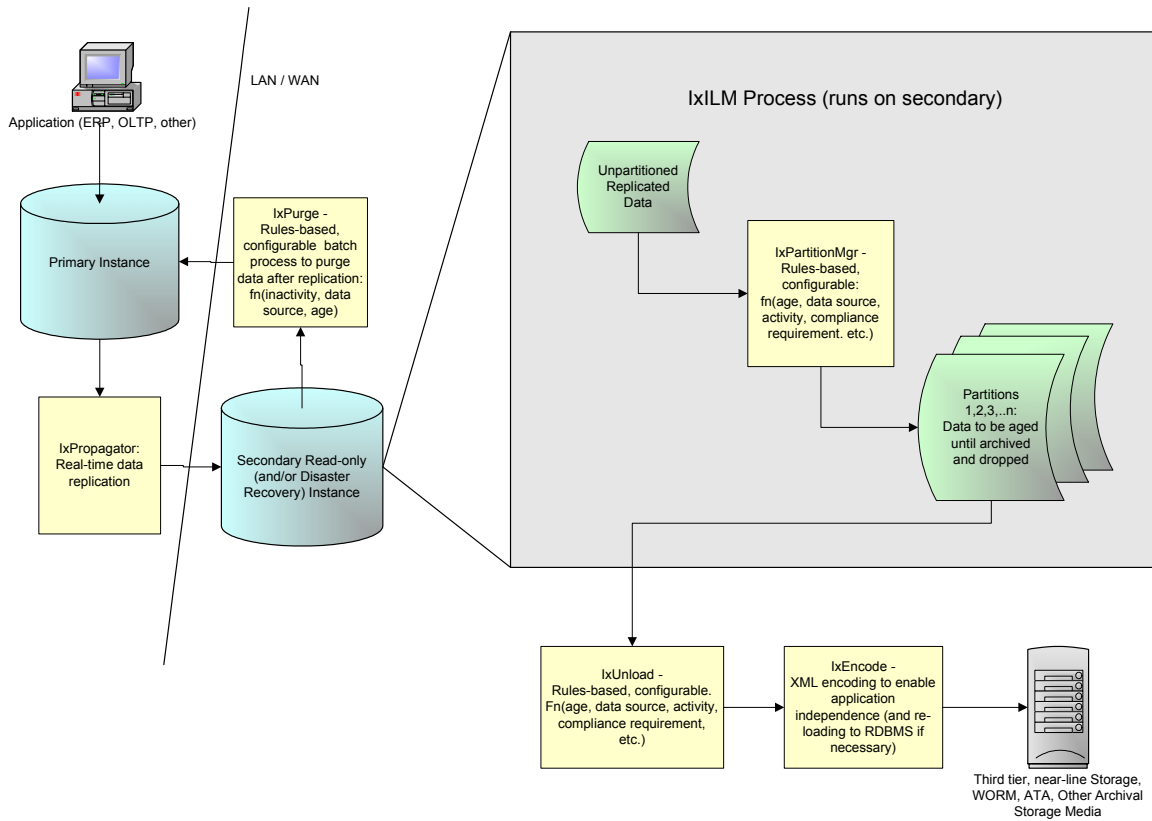
update/access activity, user activity, various compliance regulations, and information on the application from which it was originally sourced. This will allow selected data to be partitioned off and staged in multiple partitions for eventual archival to less expensive "near-line" storage such as write-once, read many (WORM) and/or ATA storage. This archival process utilizes Ixion's own existing data unload technology available in IxUnload. Once the partitioned data has been archived, it can be dropped from the secondary on-line storage, thus enabling greater usage efficiencies of on-line storage to accommodate data replicated from the primary instance.

Third-tier, Near-line Storage

Ixion utilizes a cataloging process to track data that resides in third tier, near-line storage for eventual purging as well as subsequent access that may be required at some future date. While the data that resides in near-line storage may have been created by various applications, it is important to store this data long-term in an "application-independent" format for several reasons. First, the very nature of low-cost near-line storage requires that there be no costly database or application licensing associated with it. Secondly, there is no guarantee that the application that created the data will still be available to the company for the extended durations that are required to maintain the data for compliance purposes or other archival requirements. Therefore, the archival process includes an additional process, IxEncode, to "encode" the data in XML format to enable it to be accessed by multitudes of current and future web-based reporting tools. Encoding will also include necessary schema information to allow for the archived data to be "re-loaded" to a relational database if necessary for analysis or compliance purposes.

The Ixion ILM process is illustrated in the figure below.

Ixion Information Lifecycle Management



Monitoring Tools

Additional monitoring tools will be provided to track replication rates, the growth rate of archived and unarchived data, to monitor purging, to track data movement through the archival process, to ensure that mandated storage requirements are being met, and to help optimize the rate of data archival.



Summary

The amount of data in on-line storage continues to increase, and the rate of increase has only been exacerbated by federally mandated compliance regulations that require extended data retention. As the amount of data in expensive on-line storage increases, not only do storage costs escalate, but companies become more vulnerable to data corruption and loss, decreased performance of their applications, increasingly long batch processing windows, and extended periods of downtime associated with unplanned as well as planned outages. While costs of additional on-line storage are directly measurable, the costs associated with downtime and lost data can also be directly measurable by not only the additional consulting costs associated with recovery, but also the amount of lost business incurred and any fines for non-compliance. Indirect costs can certainly result from loss of reputation due to both perceived and real risk associated with these data vulnerabilities. Since companies can afford to keep only so much data in on-line storage, the need for high-performance replication solutions integrated with configurable and automated ILM solutions is presenting itself as a necessary investment to mitigate risk, maximize application performance, and minimize storage costs and vulnerabilities associated with this scale of data retention.

Copyright © 2005 Ixion. All rights reserved.

Ixion believes the information in this publication is accurate as of its publication date. The information is subject to change without notice. This document is provided for information purposes only. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied by law, including implied warranties and conditions of merchantability or fitness for a particular purpose.

©2005 Ixion, Inc.

10805 Sunset Office Drive • Suite 300 • St. Louis, MO 63127 • 314-725-0200
Info@ixionsoftware.com