



EMC DatabaseXtender for Archiving Data — And Actually Quite a Bit More

Analyst: Michael Fisch

Management Summary

While there is no such thing as a *cure-all*, there are things around that are useful for a variety of different purposes. You might say they are *cure-many-things*. Take aspirin, for instance, a common staple in medicine cabinets. Aspirin is good for treating the occasional headache or backache, of course. It reduces inflammation in sports injuries like muscle strains, helping relieve pain and speed recovery. More recent medical evidence shows taking aspirin is also a good preventative measure for certain kinds of heart disease and cancer. Who would have thought this humble pain reliever could do so much?

EMC announced a new solution called *DatabaseXtender* which, like aspirin, can solve a variety of different problems in an enterprise's IT environment. It is part of the up-and-coming product category of Application Data Management, which dynamically manages database records through their lifecycle through policy-based archiving to secondary storage tiers. Benefits of implementing DatabaseXtender include:

- **Faster application performance** by pruning the database on an ongoing basis,
- **A defined and enforced data retention policy** for meeting regulatory, legal, and/or operational requirements of the business,
- **Lower storage and server hardware costs** by streamlining the production database and taking of low-cost storage for inactive data, and
- **Faster and less-disruptive IT operations** like backup, software upgrades, and testing and development due to a smaller database.

This may seem like a wide variety of benefits for an archiving product, but that is why it may be better to think of it as a “cure-many-things.”

DatabaseXtender preserves transparent, online access to the archived data. Though it is located on another storage system, users can still access it by normal means. Furthermore, it includes components for monitoring and analyzing data growth and a feature for creating database subsets that is useful for speeding up application testing, development, and upgrades.

DatabaseXtender initially supports Oracle Applications on Oracle Database, with additional applications and databases forthcoming in the future. Read on for more details.

IN THIS ISSUE

➤ Getting It Straight.....	2
➤ Problems DatabaseXtender Solves.....	2
➤ EMC DatabaseXtender.....	3
➤ Conclusion	4

Getting It Straight

The first thing to do is bring some order to the chaos – or definition to the buzzwords. EMC *DatabaseXtender* falls into a product category called Application Data Management (ADM). **ADM is the dynamic management of structured data through its lifecycle, from creation to deletion.** More specifically, it archives data by automatically sorting and moving it to secondary storage tiers, while preserving the ability to access and search it. The term *structured* implies data under management is contained in an application database. ADM does not apply to files or to raw blocks apart from their application context. It can manage data intelligently, based on business policies, because it is cognizant of the specific structure and relationships of data within the application.

Furthermore, the phrases *dynamic management* and *lifecycle* suggest that ADM is a form of Information Lifecycle Management (ILM), which it is. A frequently-mentioned phrase nowadays, **ILM is a concept for data** (both structured, as in a database, and unstructured, like email) **and storage management that takes advantage of the fact that the value of data changes.** Not all data is created equal, nor does its worth remain constant over time. Its initial value depends on the application or user that created it. Value tends to erode over time, though unexpected events, like a tax audit or lawsuit, can suddenly increase it again. An ILM solution seeks to apply the optimal quality of service (i.e., price/performance characteristics or backup and recovery requirements, for example) to data at each point in its lifecycle, from cradle to grave. It classifies data based on policies and moves it among multiple tiers of storage over time. The primary benefit is a minimized storage TCO while precisely meeting the requirements of the business. It avoids two common extremes: (1) Spending too much on storage by assuming the value of most data is high and (2) spending too little and suffering performance bottlenecks or crippling data loss.

In a nutshell, ILM is a lifestyle and ADM is a fashionable accessory. Keep in mind that there is no silver bullet for ILM – no single product that encapsulates what it is or can be. But EMC's *DatabaseXtender* ADM could be your entry point into the ILM lifestyle.

Problems DatabaseXtender Solves

As mentioned, *DatabaseXtender* can solve a variety of problems. Any of the following could be a trigger to consider it:

- **Application performance is slow because of a large and fast-growing database.** Databases are like cargo ships in that they become slow and unstable when overloaded (i.e., too large). For example, user response time increases, which can negatively impact worker productivity. Instability can ensue and cause downtime, also impacting productivity. By automatically pruning records from the database and moving them to a secondary storage tier, *DatabaseXtender* offloads the extra “cargo”, streamlines the database, and makes it faster and more nimble. The alternative is to continuously tune the database and periodically throw more hardware at it in the form of server processors, memory, and top-tier storage capacity. Or, one can manually prune and purge, but this is labor-intensive and can lead to inadvertently discarding the “wrong” data.
- **Data retention is necessary to meet regulatory, legal, and/or operational requirements.** The recent trumpeting of regulatory compliance concerns has pushed the issue of data retention into the spotlight, but it has always been an important business issue. Best practice suggests that all enterprises should define and enforce data retention policies, if only to meet the legal, tax, and operational requirements of the business. In reality, though, it may be the need to comply (or fear of noncompliance) with regulations like Sarbanes-Oxley, HIPAA, SEC 17a-4, etc., that cause many to take action. By

archiving records based on business policies and preserving access, DatabaseXtender can be the tool that allows an enterprise to set and enforce retention policies for structured data.

- **Storage hardware acquisition costs are too high.** In the spirit of ILM, archiving inactive data saves the cost differential between the primary and secondary tiers of storage. This is especially attractive if the amount of data is large. Platforms like ATA disk arrays can be a fraction of the price of high-end Fibre Channel or SCSI arrays, while still keeping data online. Savings accrue even more significantly if an enterprise uses multiple replicas of the database for purposes like disaster recovery, testing, and development – as many do. The total figure would be a multiple of the number of replicas.
- **A too-large database is an inhibitor to timely backup, database upgrade, or replication.** The bigger a database is, the longer it takes to backup, restore, replicate, or perform an upgrade. If any of these processes take an inordinate amount of time, it can cause unacceptable downtime, interfere with production operations, and/or delay testing and development cycles. There may be more than one way to solve this problem, but using DatabaseXtender to trim and keep the database at a reasonable size may be the simplest solution.

EMC DatabaseXtender

DatabaseXtender is comprised of four components.

DatabaseXtender Optimizer

Optimizer moves inactive records out of the production database and into an archive (i.e., secondary storage tier). For each record moved, *Optimizer* leaves a stub in the primary database that points to the secondary database. So, users can still access archived data in a transparent and online fashion. Nothing fundamentally changes from their perspective, except that

Defining Archival Policies

The key to successful policy definition is to allow the broader context of the business to shape them.

Business stakeholders and IT personnel need to be involved in the process because no single group has a full sense of what can or ought to be done. Functional personnel have the best sense for their particular business process and application requirements. IT personnel have the best understanding of technical issues and ramifications. Where regulatory and legal requirements are concerned, executives and even legal counsel will want to provide input and direction. How data is handled is critically important to an enterprise, so communication and a consensus-oriented approach are recommended to ensure policies reflect the needs of the whole business.

the data is read-only. Even patches and upgrades can be applied. This level of transparency is a unique feature, and it is an important one for winning approval from stakeholders in your organization who may be nervous about allowing their application data to be moved.

DatabaseXtender Analyzer

Analyzer is a monitoring and analysis tool that identifies the fastest-growing tables, records, and application modules. It assists customers with defining their policies by highlighting growth patterns and their impact on business operations (by linking them with specific application modules).

DatabaseXtender Subsetter

Subsetter is a tool for creating subsets of the database with fully-intact data relationships. The subsets are useful for common activities like testing, development, upgrades, and enhancements. The smaller size of the instances means they require less server and storage hardware and can speed up the processes in which they are used.

DatabaseXtender Archiver

Archiver performs the second and final stage in the archiving process (before deletion). After data has aged to the point where users no longer need transparent access, Archiver can send it to a low-cost storage media like tape or content-addressed storage (CAS) for long-term storage. It converts data in a self-describing flat file using XML/XSD. The context of the data is preserved but the application is no longer needed to view it. Users can search and access the data using third-party report tools. Archiver is scheduled to be available in Q3 of 2004.

Other Details

DatabaseXtender currently supports Oracle Applications on Oracle databases for Solaris and HP-UX. EMC plans to add support for PeopleSoft and custom applications on Oracle databases in Q3, and beyond that there are plans for Sybase, SAP, Siebel, UDB, and SQL Server. EMC also offers professional services solution in conjunction with Oracle for installing DatabaseXtender and setting up appropriate archiving policies based on business requirements. (See box on the previous page for more details on policy definition.)

Conclusion

The data archiving, analysis, and replication features of DatabaseXtender touch on many areas – more than just filing away data. It may be a performance issue that leads you to consider it, or a regulatory requirement, hardware infrastructure costs that are too high, or an IT process like backup or upgrading that is taking too long. An ADM solution that can address all of these issues is valuable in a similar way as aspirin, the *cure-many-things*. Because of the amount of innovation and competition in the IT industry, there are always other technologies and approaches to solve these problems. So, it always comes down to a cost/benefit analysis. **But you may find that EMC's DatabaseXtender solution for ADM is the most effective and economical of the alternatives.**



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Michael Fisch is Director of Storage and Networking for The Clipper Group. He brings over eight years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at Mike.Fisch@clipper.com or at 781-235-0085 Ext. 25. (Please dial "1-25" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, and "*clipper.com*" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "*Navigating Information Technology Horizons*", and "*teraproductivity*" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.