Database-Archiving Products Are Gaining Market Traction

Gartner RAS Core Research Note G00161796, Carolyn DiCenzo, 13 October 2008

Few vendors are offering solutions for archiving data from custom-built database applications or from the breadth of ERP and CRM applications. Vendors in this market have been challenged in their attempts to gain market visibility and awareness for their products and for the benefits they can provide compared with just throwing more hardware at the problem.

Key Findings

- Database archiving significantly lowers storage costs for primary storage by moving older data to less-costly storage.
- Archiving reduces the size of primary storage, resulting in improved application performance and lower storage requirements for copies of the database for testing, backup and other purposes.
- It isn't necessary to have terabyte-size database applications to benefit from archiving. Performance and cost improvements can be sizable, even with applications that have less than 200GB of data.
- Although the number of solution providers is small, the breadth of coverage and the approach to solving the problem varies slightly with each, providing enough competition to give companies a choice when selecting a vendor.

Recommendations

- Every application should have a retention policy that includes when to delete the data. When data is being kept long term, there should be a plan to archive older data. Examine the database-archiving solutions that are available on the market, and determine how these solutions can accelerate the deployment of archiving.
- Find one application or module that is experiencing performance or backup problems, and use it as your first archiving project. Success with this project (and with the selected vendor) can be used to convince other application owners to implement archiving.
- Require all new applications to have an archiving plan, because it's easier to implement upfront when the business user and the database administrators are engaged with the application design and deployment than it is once the application is running and everyone has moved on to other projects.

Gartner

ANALYSIS

Database archiving continues to be a reluctant market, with customers adopting archiving solutions only when adding hardware no longer fixes the problem. Oracle provided an early push for the market with the release of Oracle E-Business Suite 11i in 2000. Organizations found that the upgrade required as much as 50% more storage. With that growth in the data store came performance issues, problems meeting backup windows (as well as recovery objectives), increased time for database maintenance and demands for additional storage that stressed IT budgets. The success of early adopters of database archiving in addressing their Oracle application growth issues is now providing references and opportunities for solution providers as organizations are seeing a steady growth in data stored as part of the normal growth of business information.

The entry into this market by HP and IBM has raised buyer confidence in the value and viability of the approach. Although the impression is that the benefit of archiving is realized in archiving terabytes of data from a database that has grown to hundreds of terabytes, real benefits can be seen by removing data from data stores that are as small as 50GB to 200GB. Deployment of archiving for applications where the vendor has prebuilt archive templates can take as little as two weeks – it will take longer if the organization has customized the application in a significant way.

Database-archiving products need to extract data from various database columns, rows and tables potentially across multiple table spaces on multiple physical disks. The extraction must be done at the application level so that associated business logic can be referenced to perform the extraction. This extraction leverages the purge facility provided by the database management system (DBMS) or application or uses the vendor's optimized routines, which leverage the same business rules as used by the application. Extracted data is often stored in archive database tables, but it can also be stored in a file-based structure – either the vendor's proprietary format or XML.

Gartner does not include in the database-archiving market special-purpose solutions that work at the application layer, but not at the database layer. An example of this kind of solution would be Open Text's Livelink ECM – Data Archiving for SAP, which operates on the application layer, leveraging application programming interfaces (APIs) provided by SAP. Gartner also does not include project-based solutions created by a vendor as part of a consulting project for a specific organization.

Data that can be deleted should be deleted, but some historical data is needed for audits, to meet retention policies or for business intelligence (BI) use. Archiving products enable administrators to move data out of the production database, based on business policies, saving the data and metadata on lower-cost storage while providing a way for users to access the archived data on demand. Referential integrity of the data is maintained, regardless of where the data is stored. Data can be brought back into the operational database in bulk or selectively, if required. The benefits of implementing database archiving for packaged or custom-built applications are:

- Improved application performance
- Decreased backup and recovery times
- Less high-performance (expensive) storage needed for the operational data with older data moved to lower-cost media
- Less storage for database copies, backups, clones and mirrors
- Lower licensing costs for applications and databases that are tied to total CPU storage use
- Less time spent tuning performance
- Less time spent on application or database upgrades
- Opportunity to move historical data or data from retired applications to another format for longer-term retention and access

Most vendors in the database archiving market leverage their archiving technology to provide tools for data subsetting, usually with data masking. These subsetting options speed the process and reduce the amount of data included in copies of the production database used for activities such as testing. The data masking helps meet regulatory requirements for data privacy.

1.0 Market Education Remains a Challenge for Vendors

The challenge facing vendors in an emerging market goes beyond delivering well-designed solutions. The biggest challenge is to educate the market of the value of changing established procedures and to get organizations to try a new approach to managing a data growth problem. Convincing organizations to look at database-archiving solutions has been particularly challenging, because of the resistance to touching working, mission-critical applications. Records in database applications are not deleted by individual users as they are for files and e-mail, so legal discovery and regulatory requirements have not been the strong forces for archiving that they have been in other areas. For databases, performance and application management have been the primary drivers for archiving.

Vendors entering this market need expertise in application development and database management, which results in a high barrier to entry and, thus, fewer solution providers. By 2000, OuterBay (now owned by HP) and Princeton Softech (now owned by IBM) had launched database-archiving solutions. Princeton started out with a solution for IBM's mainframe DB2 DBMS and then joined OuterBay in supporting open-systems database archiving. Applimation and Solix joined the market in 2003. Getting organizations to archive data out of custom applications proved the most difficult, because the original architects of the application were often unavailable and, thus, resistance was strong.

To provide a quick path to a successful initial archiving implementation, the vendors moved into support for the packaged application market, with Oracle E-Business Suite proving to be the most ready market, because upgrades resulted in larger data stores. Support for a broader set of the E-Business modules and for other packaged applications has been a way for vendors to differentiate themselves.

Partitioning is seen by some as an alternative to archiving. It can be used to address performance and storage cost issues; however, it is not functionally equal to database archiving and, thus, does not address the full range of customer requirements. Database partitioning provides a method of scaling certain applications and provides the ability to leverage tiered storage to take advantage of lower-cost storage options. However, it only works at the table level and for simple partitioning keys such as date.

Complex databases are harder to partition, especially if they were not architected with partitioning from the start. The partitioned data continues to reside in the database, so it doesn't significantly improve database performance the same way that moving complete sets of data out of the database accomplishes, especially if open, rather than indexed, queries are frequently run against the data. Database archiving is about moving complete transactions across sets of tables and maintaining the integrity of production and archive data. Archiving and partitioning can be complementary, not competing strategies. Partitioning, if desired, can be used to address performance and storage cost problems, with archiving then used to remove historical records when overall data size becomes a problem. Partitioning can also facilitate the application retirement process.

Based on total number of database-archiving customers as of year end 2007, IBM is the leading vendor, with 45.5% of the market (see Figure 1). With the acquisition of Princeton Softech in 2007, IBM joined its mainframe database-archiving business with Princeton's mainframe and distributed system solution.

If you look at only the packaged application part of the archiving market (see Figure 2), then Applimation leads with 33.3%. Applimation reports the largest number of Oracle E-Business Suite customers using its archiving solution. IBM has the broadest application archiving portfolio with penetration into the PeopleSoft, Siebel, Amdocs and JD Edwards customer bases. California-based Solix was later to market, but it has worked with partner Wipro Technologies to gain early traction in India, as well as the U.S. Neon Enterprise Software entered the market in 2007 with a mainframe archiving solution, but it later exited the market. HP lost market share to Applimation and Solix while preparing for an April 2008 product release that improved development and deployment tools and expanded platform support beyond Oracle to include Microsoft SQL Server.

Figure 1. Relative Position of Database-Archiving Market Leaders

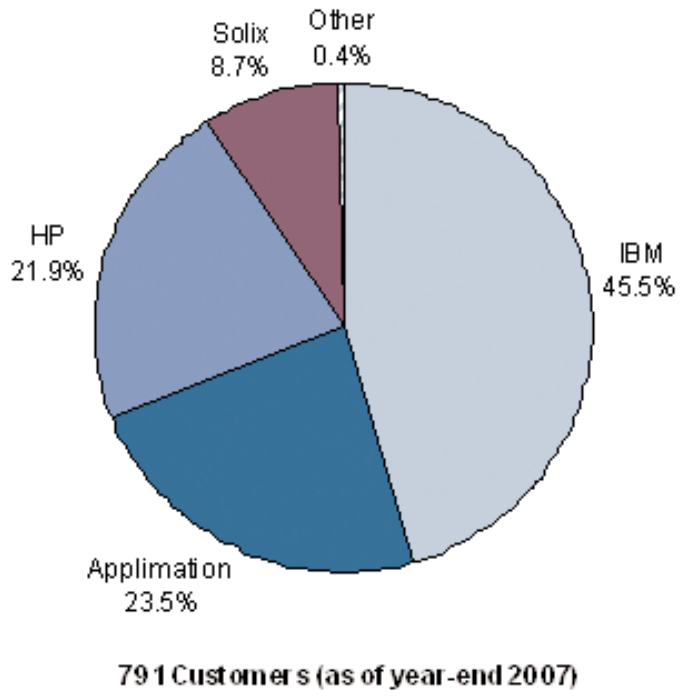

Source: Gartner (October 2008)

Figure 2. Packaged Application Archiving

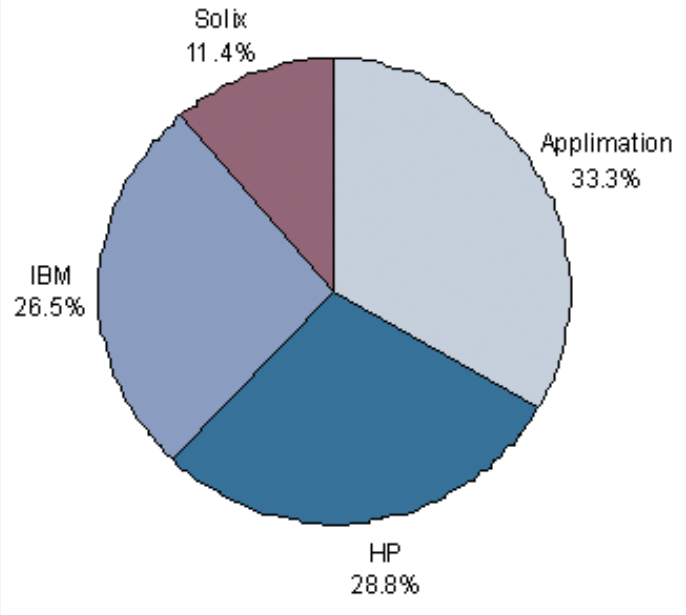

Source: Gartner (October 2008)

2.0 Key Vendors in the Database Archiving Market

Gartner began covering the market for database-archiving products in 2002 and first invited users to share their experiences with archiving tools at the 2003 Gartner PlanetStorage conference. Client inquiry between 2003 and 2007 was light, averaging one or two calls a month. However, market interest in database archiving is now growing, and organizations are becoming more interested in including database-based applications in their information life cycle management projects. The number of companies that responded to a Gartner survey saying that they have deployed some sort of archiving for database applications increased from 24% in June 2006 to 30% in June 2008. Only 10% of users surveyed in the 2008 report indicated no interest in database archiving.

Starting in 2007, Gartner's inquiry call volume on database archiving began to escalate, averaging five or six calls per month for 2007 and doubling that volume in 2008. Initially, interest was focused on archiving data from commonly deployed packaged applications, but calls regarding archiving for application retirement have increased in the past year. Gartner conversations with users of archiving products indicate that all deliver on their promises, and purchase decisions are usually based on vendor experience with a specific DBMS or application module, the vendor's willingness to work with the company on a unique archive requirement, or the vendor's ability to build archive templates for an unsupported application or application module. Each vendor has a different look and feel to its implementation tools, and users will gravitate to the approach that most appeals to their way of working.

2.1 Applimation Informia Archive

Applimation is a 150-person, private company that was founded in 1998. With the merger of Tickmark Solutions and Crystallize in July 2002, the company took the name Applimation for the resulting merged entity. Applimation is based in Chicago, with international headquarters in Switzerland and India. Applimation launched Informia in 2003, leveraging established technologies into a new product set under the new brand. Although Applimation supports archiving from DB2 on distributed and mainframe systems, as well as Microsoft SQL and Sybase Adaptive Server Enterprise, the bulk of its business has centered on applications built on the Oracle relational DBMS (RDBMS), including Oracle E-Business Suite, PeopleSoft and Siebel applications, as well as custom-built applications.

With the acquisition of Gamma Technologies in January, 2008, the company expanded its data-subsetting capabilities. It also acquired expertise with the SAP ERP application, enabling the company to expand into the SAP archiving market in September 2008. Special solutions for Deltek's Costpoint ERP solution, which is specifically designed for use by project-oriented businesses, and for the PeopleSoft Campus Solutions for Higher Education Universities were released in January 2008. Approximately one-third of the company's business comes from outside the U.S. Key channel partners include CSC, Intelligroup, Infosys and Sybase.

The Applimation Informia Archive offering uses the built-in application's purge routines or its preferred purpose-built routines, which extend the same business rules in the application's native purge. Data is archived to another database instance, or to XML, and deleted from production. Both production and archive data are transparently available to application users. The company's Data Growth Analyzer tool is used to help prospects gather trending and sizing information about tables and databases to determine the potential impact of archiving. The tool can also be used to construct a business case and determine the return on investment (ROI) of the archiving project.

Within the product, the Enterprise Data Manager is viewed as the single biggest differentiator. This component enables a customer to review, modify and extend Applimation application archiving templates, called accelerators, and enables customers to build their own accelerators for custom applications or applications that do not have Applimation-produced accelerators. This Java-based tool has a graphical interface that is used to display tables and relationships and provides visibility into every aspect of the Informia Archive product. Enterprise Data Manager works across all supported DBMSs and applications.

2.2 HP Database Archiving

In February 2006, HP entered the database-archiving market with its acquisition of California-based OuterBay Technologies. Initially managed out of the HP storage business group, the product is now part of the HP Software Information Management business unit. Although OuterBay was a visible player in the early market, the offering has become less visible as a part of HP, as the company focused on a few select new customer deployments in 2006 and 2007, worked to re-engineer the product to HP standards, and enhanced the design and development tools. The company also needed to build the necessary sales and support organizations. With the release of version 6.0 in April 2008, the company has relaunched the product for more-aggressive sale by HP sales and partners. It has also developed a set of services for delivery by HP's professional services organization. Most product deployments by HP have been in North America and the Asia/Pacific region.

HP Database Archiving supports database-to-database archiving for Oracle E-Business Suite with a growing business in PeopleSoft Enterprise applications. Application-level access is retained for archive retrieval. About 20% of customers are using the product to archive custom or other packaged applications. With v.6.0, the company added extended capabilities for Oracle, including partitioned-based archiving and Windows platform support, and added Microsoft SQL Server archiving of custom applications. The product is made up of a set of components, including a new visual designer tool used to model data and build business rules in support of archiving from third-party or custom applications.

Other components include prebuilt application archive packs and tools to handle data movement, process management and archive access. A module for data subsetting (but with no data masking capability) is also available. The product supports migration of the archive data to XML format for long-term retention. A bulk archive option is used for first-time archiving of large data volumes to achieve desired retention levels. This is run at times when the access to the production database is low and includes tools to reorganize and shrink database size allocation.

2.3 IBM Optim Data Growth Solution

In September 2007, IBM acquired Princeton Softech, a private company based in New Jersey, for its Optim product suite of database-archiving, test data management and data privacy software. Princeton Softech first released its mainframe Archive for DB2 in 1999, then expanded it to include open-systems DBMS and packaged application support. The company was previously known for its portfolio of test database management tools. Princeton Softech brought all its solutions under the Optim brand in October 2005, integrating them into a single solution framework.

IBM offered a basic archiving solution to its mainframe customers prior to acquiring Princeton, but now leads with the Optim solution. Optim has had little competition in IBM's mainframe and open-systems DB2 DBMS market, which now represents slightly more than 40% of its installed base. It also was early in providing the widest range of application archiving solutions, even as it reports more than half of its customer base is using the product to archive custom applications.

As part of IBM's Information Management business unit, which includes the DB2 group, the product benefits from the established sales and delivery expertise supporting DBMS sales and project implementations, adding to the over 200 employees that Princeton Softech brought with it to IBM. Because of the number of employees transferring to IBM, the established links within IBM Global Services, and the fact that IBM allows the product line to operate as a mini-company within the larger group, the transition to IBM has not slowed the growth or market awareness of the offering.

Optim supports Oracle E-Business Suite, PeopleSoft Enterprise, JD Edwards EnterpriseOne, Amdocs CRM and Siebel. It also supports custom and packaged applications based on DB2 for z-series, i-series and distributed platforms; Oracle; Microsoft SQL Server; Sybase Adaptive Server Enterprise; and IBM Informix Dynamic Server. It also reports support for IMS, VSAM/SEQ, Teradata and Adabas. Optim supports a range of archive formats, including archive to another database instance, compressed file or XML.

When archiving to the Optim proprietary file archive format, access to the archive is via any ODBC/JDBC compliant application, including the Open Data Manager supplied by Optim. The archive, or any portion thereof, can be restored to the originating application or database, or to any replacement application or database. As part of its sales and deployment activities, IBM is able to offer a set of assessment and design tools, services and professional services resources, which are especially helpful with complex projects, or where the prospect is looking for help in managing the project.

2.4 Solix Enterprise Data Archiving

Solix is a privately owned, California-based company that was founded in 2002. The company introduced ARCHIVEjinni in 2003, rebranding it as part of the Solix Enterprise Data Management Suite (EDMS) in 2007. The suite includes database archiving, cloning, subsetting and masking, and application sunsetting and migration capabilities. More than 90% of its business is centered on the Oracle Database platform, but it also has customers using the suite with Microsoft SQL Server, DB2 and Informix. Application coverage includes Oracle E-Business Suite, PeopleSoft applications and Baan ERP. The company also supports JD Edwards and Siebel applications. Most customers are in North America and India, where Solix has a customer support center, in addition to the one in Santa Clara, California. Global services provider Wipro has deployed the solution to customers in the U.S. and India as well.

With EDMS version 4.0 released in April 2007, archiving for multiple applications can be designed, deployed and managed with a single instance of EDMS, using a Web-based management interface. The metadata manager component of the suite captures metadata from packaged and customized applications leveraging pre-populated modules for applications supported by Solix. The metadata is defined once during the initial implementation for a given application and then can be shared across archiving, cloning, subsetting, migration and sunsetting activities. Data is usually archived to another database instance; however, the solution also supports the option to archive to XML.

3.0 Bottom Line

Archiving and deletion must be part of all application plans. Although database archiving is less known and understood, the tools available in the market provide viable solutions with references to help dispel the perception that this kind of tool is risky to implement. Organizations should find one application that is particularly problematic, in terms of performance or management, and use it as their first implementation of archiving. Use the success of this project as a model to break down roadblocks with other application owners. Organizations facing application retirement projects should look to these vendors for tools to provide a way to get data that must be retained into a format that can be accessed independently of the retired application.