



The New Enterprise Blueprint

Featuring the Gartner Magic Quadrant



Featuring research from

Gartner

“By 2016, 30% of businesses will have begun directly or indirectly monetizing their information assets via bartering or selling them outright.”¹

Solix delivers the first unified platform for Enterprise Archiving, Enterprise Data Lake, and Analytics for all enterprise data.

Executive Summary

CIOs are in a bind. Demands for operational efficiencies and better business intelligence (BI) seem to be at odds, and CIOs are caught in the middle. They are expected to not only manage data growth and infrastructure costs effectively, but also provide modern tools to analyze and extract more value from their information assets. Now more than ever, CMOs, CFOs, and VPs of Sales are counting on the CIO to help them succeed.

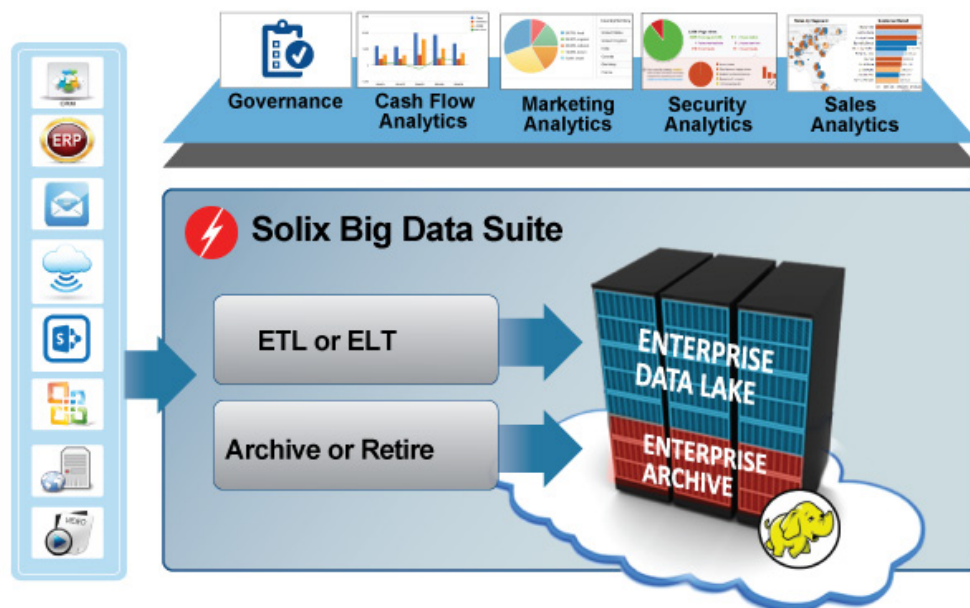
Gartner states that: “By 2016, 30% of businesses will have begun directly or indirectly monetizing their information assets via bartering or selling them outright.”¹ Other businesses simply want to mine enterprise information to improve effectiveness and bring more value to the business. Either way CIOs now hold the keys to organizational success and have gained responsibility for business growth. Caught in the bind are challenging requirements to control infrastructure costs, meet compliance objectives and also deliver better business intelligence. Process duplication, incremental hardware requirements and scarce resources

combine to threaten IT success. Few enterprise blueprints are capable to dramatically improve enterprise analytics and also reduce the cost of information assets.

Now more than ever, CMOs, CFOs, and VPs of Sales are counting on the CIO to help them succeed.

Solix Technologies has found a new way to deliver dramatically better business intelligence and still meet process, compliance, and infrastructure cost challenges. Apache Hadoop not only offers an exciting new framework for enterprise analytics, it also stores and processes large data sets at the lowest possible cost. According to a survey by Gartner, “big data investments in 2013 continued to rise, with 64% of organizations investing or planning to invest in big data technology compared with 58% in 2012.”²

FIGURE 1 Solix Big Data Suite



Source: Solix

^{1,2} “Big Data Benefits are Hampered by ‘Culture Clash,’” Gartner, September 12, 2013

Big data helps CIOs out of the bind, and we've built the Solix Big Data Suite to end competing interests between operational efficiency and BI. Built on Apache Hadoop, the Solix Big Data Suite is the first unified platform for Enterprise Archiving and Enterprise Data Lake to deliver analytics applications across the enterprise. Secure, compliant, low-cost, bulk storage of both structured and unstructured data is now available on a petabyte scale and capable to meet the most demanding requirements for next generation analytics.

Gartner states; "By 2016, 75% of structured data archiving applications will incorporate support for big data analytics."³

Enterprise Data Management solutions are changing rapidly and offering important new opportunities. IT costs may now be managed by the same enterprise blueprint that delivers improved BI. Analytics are the foundation upon which organizations are making business decisions to not only survive, but to thrive. Gartner believes BI has become a business asset in itself, and effective data management is at the heart of it all. The new enterprise blueprint manages all enterprise data at the lowest possible cost, improves operational efficiency and enables improved BI as well.

Source: Solix

³ "Magic Quadrant for Structured Data Archiving and Application Retirement," Gartner June 11, 2014

The Data Growth Crisis

CIOs are in a difficult position. The demand for operational efficiencies and the demand for Business Intelligence seem to be odds. Traditional thinking says the processes and infrastructure required to control costs, meet compliance objectives and keep costs in check do not mesh with the processes and structures needed for BI. Duplication of effort, hardware and operational costs could cripple IT.

CIOs must change the paradigm to ensure organizations can compete in today's economic climate as structured and unstructured data continue to explode. The time has come for a new enterprise blueprint. The tools are here. Apache Hadoop makes it possible. Solix Big Data Suite makes it a reality.

The Paradox of Moore's Law

While data growth explodes exponentially, Moore's Law continues to astound as processor and integrated circuit performance doubles every few years, just as the visionary Intel founder predicted. The benefits of such dramatic technological advances cannot be overstated: Today, more processing power is packed into a smartphone than yesterday's mainframe, and the cost to store a terabyte of data in the cloud has fallen to as low as \$10 per month.

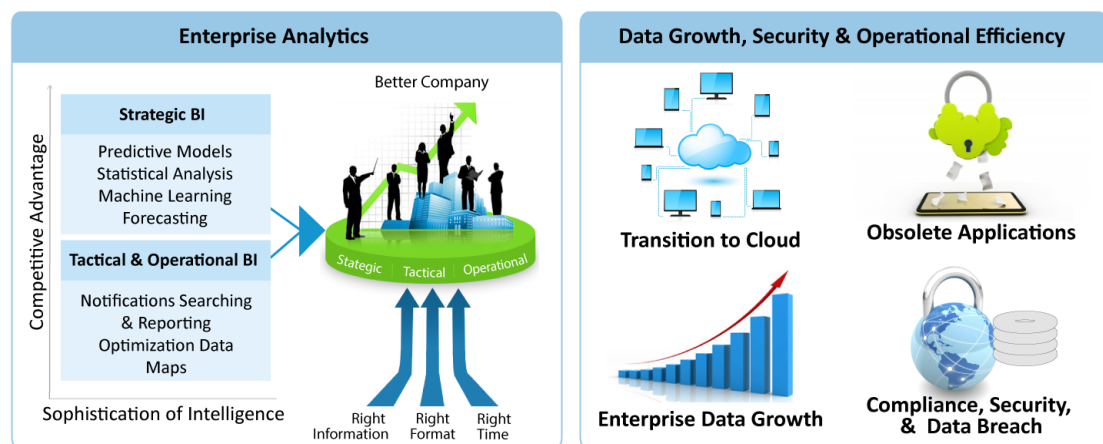
But rather than saving, organizations are spending more to gain improved business value by processing mission-critical enterprise data faster and faster. Until recently, few ERP users would ever have imagined paying the high cost to process enterprise data using full flash arrays, yet today, most have either already installed or are evaluating premium performance enterprise platforms with solid state disks (SSD) such as SAP Hana and Oracle Exadata.

Despite such spectacular gains in semiconductor price/performance, the overall cost of IT continues to rise, largely because we are now processing so much more data. Challenged by this dilemma, CIOs must continually find new ways to reduce the cost of data growth, so they may afford to fund more mission-critical applications that improve business results.

A New Enterprise Blueprint

Advances in semiconductor technology have indeed enabled "commodity" hardware to process and store extraordinary amounts of data at lower unit costs. Through virtualization, this low-cost infrastructure may now be utilized with extraordinary efficiency.

FIGURE 2 Information Management Challenges



Source: Solix

Apache Hadoop is a framework for distributed processing and storage of large data sets on virtualized commodity hardware. Hadoop has been designed to provide analytics over petabyte-scale data sets, and the infrastructure can be leveraged for long-term storage and analytics.

The Hadoop Distributed File System (HDFS) has rapidly emerged as the leading storage platform because it provides secure, stable storage for structured and unstructured enterprise data with enhanced access. Hadoop's MapReduce framework can process large data sets across distributed compute nodes in parallel, allowing "commodity" hardware to be used as the most efficient and cost-effective bulk data storage solution available.

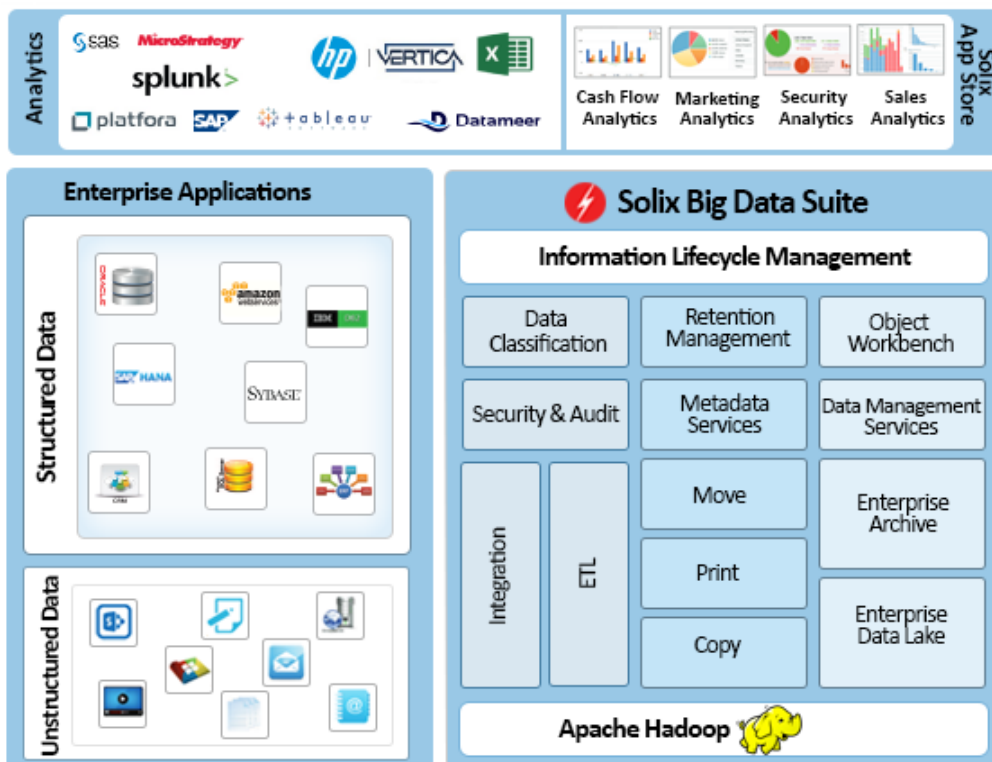
Experts agree that as much as 80 percent of production data in ERP, CRM, file servers and other mission-critical applications may not be in active use, and both structured and unstructured data becomes less active as they age. Large amounts of inactive data stored online for too long reduces the performance of production applications, increases costs and creates compliance challenges.

Apache Hadoop is a better alternative to storing inactive enterprise data online. By moving inactive data to nearline storage, application performance is improved and costs are reduced because data sets are smaller and workloads are more manageable. Universal data access is maintained by analytics applications, structured query and reporting, or just simple text search.

This new enterprise blueprint reduces infrastructure costs and enables organizations to gain improved value from their data. Enterprise data warehouse (EDW) and analytics applications leverage big data for better described views of critical information. As a low-cost data repository to store copies of enterprise data, big data is an ideal platform to stage critical enterprise data for later use by EDW and analytics applications.

Source: Solix

FIGURE 3 The New Enterprise Blue Print



Source: Solix

Solix Big Data Suite provides an Information Lifecycle Management (ILM) continuum that ensures an efficient and compliant extract transform and load (ETL) process.

Solix Big Data Suite — The New ILM Continuum

Capture. Organize. Analyze.

The Solix Big Data Suite delivers a new enterprise blueprint to help CIOs drive down infrastructure costs and make BI flourish, all while ensuring data is safe and secure. Solix Big Data Suite provides an Information Lifecycle Management (ILM) continuum that ensures an efficient and compliant extract transform and load (ETL) process. CIOs no longer need to choose between compliance, efficient processes and BI.

ILM Best-Practice

Gartner says “any organization thinking of simply applying existing information governance practices to big data will likely fail — not least because much data is ungoverned; or governed by others according to a different set of objectives.”⁴

Solix Big Data Suite provides the first true ILM continuum to address the complexity of data governance as data is copied and moved between production and Big Data. The ILM framework not only manages the ETL process, it also manages the data within HDFS and HBASE and provides an integrated retention management and legal-hold

capability for Apache Hadoop. Data from structured and unstructured sources is classified, migrated into HDFS/HBASE and validated with audit reports. These reports provide defensibility and chain of custody for compliance and data governance.

Solix’s ingestion engine captures application metadata and integrates with enterprise authentication systems, such as Active Directory, LDAP and others to establish role-based security. Solix provides a single user interface to create a seamless user experience for administrators and end-users. Through role-based-security, the Solix interface allows users to access different facets of the product. For example, a compliance officer could be given access to create retention policies, apply legal-holds or generate reports on cases. End-users, however, might only be allowed to perform keyword-searches on projects assigned to them.

This extensive ILM framework allows Solix Big Data Suite to create a unified repository to **capture** all enterprise data and optimally **organize** it for **analytics** tools offered through the Solix App Store.

FIGURE 4 Solix Big Data Suite - Advantages of Apache Hadoop

Scalability	Supports petabyte-scale enterprise data
Cost	Built on low-cost commodity storage
Analytics	1000s of tools to unlock value of enterprise data
Modern	Supports petabyte-scale enterprise data
Comprehensive	Supports structured, unstructured, social data, etc.

⁴“Big Data Governance From Truth to Trust,” published July 18, 2013

The highly scalable Solix Big Data Suite offers an extensible connector framework to ingest all forms of enterprise data from any source. Archiving, application retirement and flexible ETL capabilities improve the speed of deployment, reduce costs and optimize available infrastructure. Solix supports on-premise and cloud-based deployment on a variety of Hadoop distributions.

The Solix Big Data Suite harnesses the capabilities of Hadoop to create a comprehensive and efficient platform that produces unified and cost-effective ETL, ILM and BI infrastructures for all data, requiring smaller teams with fewer IT skills, while allowing quicker rollouts and faster results.

The Solix Big Data Suite includes:

- Solix Enterprise Archiving to improve enterprise application performance and reduce infrastructure costs. Enterprise application data running online is first moved, and then purged from its source location according to ILM policies to ensure governance, risk and compliance objectives are met.
- Solix Enterprise Data Lake reduces the complexity and processing burden to stage enterprise data warehouse (EDW) and analytics applications, and provides highly efficient, bulk storage of enterprise data for later use when it is needed. Solix Data Lake provides a copy of production data and stores it “as is” in bulk for later use.
- Solix App Store offers pre-integrated analytics tools for data within the Enterprise Archiving and Enterprise Data Lake.

Source: Solix

FIGURE 5 Information Lifecycle Management Framework



Source: Solix

Experts estimate that up to 40 percent of applications are candidates for retirement, migration or rationalization.

Solix Big Data Suite — Enterprise Archiving

Solix Enterprise Archiving offers a unified ILM framework to archive and retire all enterprise data — structured and unstructured — to optimize application performance, reduce cost and achieve compliance goals. Solix's ILM framework meets all the necessary aspects of governance, including legal hold, retention management, eDiscovery, and auditing. Solix takes governance to the next step by including validation of all archiving and retirement actions taken within the platform.

The Enterprise Archive has two components, Data Archiving and Application Retirement.

Data Archiving

Database archiving has emerged as a key component to an ILM best-practice framework to meet data growth challenges.

Data archiving best-practice requires that move and purge processes be coordinated and validated. All data deletions are done in a defensible manner with audit records and chain of custody. Solix Enterprise Archiving ensures proper data governance since enterprise data is ingested and stored based on retention management policies with support for custom business rules. Archive data is classified for security and compliance requirements, such as legal hold, and universal access is provided for business users through structured reports and full text search for business objects.

Application Retirement

Retiring legacy applications creates instant cost savings by freeing staff for more important assignments, eliminating licensing and maintenance fees, and enabling IT to shut down unnecessary hardware saving power, floor space and air conditioning load.

Retired data — both structured and unstructured — is compressed by as much as 90 percent and stored at the lowest possible cost in an immutable format which is easily accessed by users. Yet, active data can be migrated to the production database and accessed by new core applications. Hundreds of legacy applications, in a variety of configurations and on different platforms, can be decommissioned.

Print and Purge for Database Archiving and Application Retirement

Solix Print Report Archiving offers the capability to generate printed reports from native applications in readable formats such as PDF or Text for archiving or application retirement. The tool allows import of previously created reports from folders. Users can access these reports for compliance, eDiscovery and other regulatory requirements through Solix's search and reporting interfaces.

Solix Enterprise Archive Benefits

The features of database archiving and application retirement ensure that Solix Enterprise Archive meets ILM best practice for both structured and unstructured data. All move and purge processes can be automated based on organization policies and schedules. Only Solix Big Data Suite ensures all move and purge actions are validated to ensure the proper data has been processed.

User access to the data occurs through reporting and search tools utilizing metadata tags and keywords. Furthermore, archived data is classified for security and compliance requirements and access is role-based and predicated on clearance.

Solix Enterprise Archive also benefits organizations through:

- Improved application performance by reducing the amount of production data core applications must process.
- Smaller pools of active data means faster backups, minimized downtime during upgrades and faster return time for disaster recovery. Cloning time will also be reduced.
- Reduced operational costs and the ability to reallocate staff time to more mission-critical operations.
- Decommissioning legacy applications, which allows for the elimination of infrastructure, maintenance, and support costs.

Source: Solix

Solix Big Data Suite — Enterprise Data Lake

Solix Enterprise Data Lake is a repository to integrate all formats of data from structured and unstructured sources. The Solix Enterprise Data Lake offers data integration and ETL capabilities on a low-cost, highly scalable platform.

Traditional EDW platforms deliver highly specific data views based on corporate strategy. This canonical top-down enterprise approach is deductive and has merits, yet it does not allow for the inductive analytical style that drives many BI proponents. The fact is, the volume, velocity and variety of data in organizations is simply too much for the human mind to grasp. Yet, Gartner says these two approaches are locked in a “culture clash” that must be resolved.⁵

The Solix Enterprise Data Lake reduces the complexity and processing burden to stage EDW and analytics applications, and it provides highly efficient, bulk storage of enterprise data for later use when it is needed. Proponents of both deductive and inductive analytics can be satisfied.

With the Solix Enterprise Data Lake, the data is stored on “commodity” servers that are far less expensive than the Tier 1 infrastructure enterprise data traditionally utilizes. These commodity servers

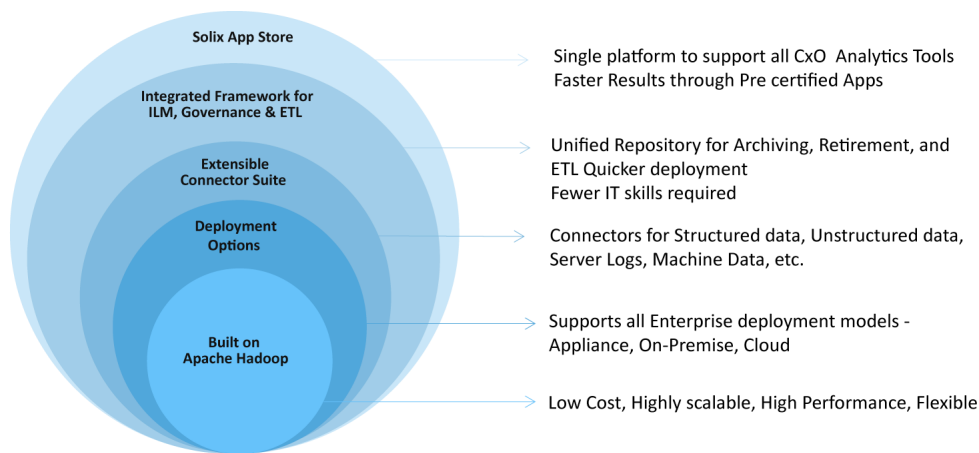
can save organizations more than 80 percent on infrastructure costs alone. Add in the reduced expense for licenses and utility costs to keep Tier 1 infrastructures running and the Data Lake pays for itself quickly.

The Solix Enterprise Data Lake provides a copy of production data and stores it “as is” in bulk to be better described and distilled later. This simple copy process eliminates the need for heavy ETL processing during ingestion. Once resident within the Hadoop file system (HDFS), enterprise data can be better described or transformed for use with business analytics applications based on user need, which allows organizations to get the most out of the data.

Utilizing the Solix Enterprise Data Lake, organizations can have access to larger pools of data to drive business decisions at far lower cost. Simple tools could leverage server and machine logs to predict system failures, alerts and the like. However, organizations need to mine data at a far more strategic level. Sophisticated BI tools can become key decision drivers and bring

Utilizing the Solix Enterprise Data Lake, organizations can have access to larger pools of data to drive business decisions at far lower cost.

FIGURE 6 Solix Big Data Suite - Key Benefits



Source: Solix

significant value to impact the revenue and growth of the organization. Solix Enterprise Data Lake also minimizes the “culture clash” with BI users where access to large quantities of data once considered unmanageable is now possible to drive deeper levels of analysis without sacrificing overall organizational needs.

Solix Enterprise Data Lake also offers the same extensive Solix ILM framework to ensure data retention and classification policies are deployed and data meets governance, risk and compliance objectives. Gartner analysts say “today’s

information governance practices focus on controlling various information quality attributes, typically for information governed behind the firewall. Organizations are either following the same approach with big data (which is proving impractical) or looking for new ways to govern that data to reduce risks.”^{5,6}

Solix Big Data Suite creates the new standard for governance in a Big Data world. ILM policies and business rules may be pre-configured to meet industry standard compliance objectives, such as COBIT, or custom designed to meet more specific requirements.

Source: Solix

^{5,6}“Big Data Governance From Truth to Trust,” published July 18, 2013

Solix App Store

The Solix App Store makes inductive BI user-friendly. The App Store offers out-of-the-box analytics through pre-integrated applications and offers the opportunity to utilize third-party apps. Organizations can also build custom and mobile apps through the App Store.

The App Store offers:

- **Cash-flow analytics**
- **Sales/pipeline analytics**
- **Marketing and social analytics**
- **Integration with Tableau, Splunk, etc.**
- **An executive dashboard**
- **Ad-hoc reporting and searching tools**

Connector Framework

Solix Big Data Suite has an extensive connector framework that allows the ingestion of data from all sources. The framework can take heterogenous data streams and allow the organization to manage the data according to business policies.

The connector framework encompasses:

- **Structured Data:** Core enterprise applications including Oracle E-Business Suite, PeopleSoft, Baan, JD Edwards, Siebel and custom applications
- **Unstructured Data:** MicroSoft Sharepoint, file systems, web servers, desktops, laptops etc.
- **Social feeds:** Including facebook, twitter, Instagram
- **Log files:** Security Logs, Network Server Logs, Mobile Logs, etc.
- **Print stream:** Native Report Archiving, Fax/Scan Data, etc.

The various streams of data are brought together through Solix Data Integration Studio, which provides a comprehensive set of ETL capabilities to ingest any enterprise data into the Solix Big Data Suite.

Source: Solix

Conclusion

Solix Big Data Suite brings the new enterprise blueprint within the grasp of every CIO.

With Solix Enterprise Archiving and Solix Enterprise Data Lake, CIOs find a path to start the enterprise Hadoop journey without having to choose between operational efficiency or BI. Solix Big Data Suite manages the entire big data lifecycle without sacrificing governance, compliance or performance.

Solix supports a variety of commercial Apache Hadoop distributions including Cloudera and Hortonworks. Deployment methodologies include a hardware appliance and public/private/hybrid cloud to fit into all enterprises.

The challenge to exploit Apache Hadoop technologies has been solved. Solix has the products and expertise to complement CIOs and their teams to leverage emerging Apache big data technologies and stay ahead of their competition.

CIOs are in a bind no longer.

Case Studies

Enterprise Archiving Customer:

Publically Traded Telecom Company

Business Objectives:

The customer has a fast growing Siebel application. The performance and cost of managing the environment is growing year-over-year. Customer would like to reduce the size of the production database through archiving, maintain access to the data, and apply retention policies to the data.

Solix Solution & Benefits:

Solix has archived more than 50 months of data, amounting to more than 19 Billion records. Due to the high transaction volume of the telecom customer, Solix archives at a rate of 300,000 records per minute. With Solix's archiving solution, the customer has recovered a significant portion of their expensive Tier-1 storage, along with improved performance. Cost of operations have been reduced by including infrastructure and support costs. Solix's Application Portfolio Manager (APM) provides a reporting and search capabilities for the data with full ILM capability for compliance and regulatory purposes.

Enterprise Archiving Customer:

Fortune 100 Technology and Manufacturing Company

Business Objectives:

The customer has more than 50 legacy enterprise applications, including JD Edwards, which are eligible for decommissioning. The primary objectives are to reduce IT cost, centralize retired data, and maintain regulatory compliance for data retention.

Solix Solution & Benefits:

Solix Enterprise Archiving platform was leveraged to migrate and consolidate data from all the enterprise applications into a single repository for centralized management. The archive provides retention and legal hold capability to address governance requirements. Additionally, for business users, an integrated user interface provides seamless role-based access through searches and reports. With Solix's solution, the customer is projected to save \$12 Million over 3 years.

Other notable use cases of the Solix Big Data Suite

- **Enterprise Data Lake** - One of Solix's manufacturing customers is deploying the Solix Big Data Suite to aggregate security log files from heterogeneous data sources to analyze them for unauthorized access to intellectual property and prevention of IP theft.
- **Enterprise Archiving** - A Fortune 500 distributor of building supplies is leveraging Solix Big Data Suite to archive SharePoint sites and consolidate documents into a unified repository for easier access and also retire an ERP application.

From the Gartner Files:

Magic Quadrant for Structured Data Archiving and Application Retirement

Structured data archiving applications help IT leaders retire legacy applications, reduce costs, and meet governance and compliance requirements. We evaluate vendors offering products and services that provide archiving for databases and data from enterprise applications such as SAP.

Strategic Planning Assumptions

By 2017, structured data archiving will represent one-fifth of the information governance efforts in enterprises. By 2016, 75% of structured data archiving applications will incorporate support for big data analytics.

Market Definition/Description

This document was revised on 12 June 2014. The document you are viewing is the corrected version. For more information, see the Corrections page on gartner.com.

Structured data archiving describes the ability to index, migrate and protect application data in secondary databases or flat files typically located on lower-cost storage for policy-based retention. It makes data available in context and protects it in the event of litigation or an audit.

Structured data archiving addresses:

- Overall storage optimization — It can reduce the volume of data in production and maintain seamless data access. The benefits of using this technology include reduced capital and operating expenditures, improved information governance, lower risk, and access to secondary data for reporting and analysis.

- Governance — The technology preserves data for compliance when retiring applications. Structured data is often transactional and related to financial accounts or back-office functions (e.g., HR, patient enrollment in healthcare and other use cases that might be regulated) that require information governance, control and security, along with the ability to respond to related events such as audits, litigation and investigation. These and other requirements, such as maintaining information context, can prevent organizations from moving data to lower-cost tiers of storage, or adopting other do-it-yourself approaches.
- Cost optimization and compliance — Structured data archiving and application retirement can result in significant ROI. Structured data in legacy systems, ERP and databases accumulates over years — and, in some cases, over decades — driving up operational and capital expenses.
- Data scalability — The technology can manage large volumes of nontraditional data resulting from newer applications that can generate billions of small objects. Scalability to petabytes of capacity is required in these cases.

The desire to leverage archives as a secondary data store for big data analytics is driving the growth of the structured data archiving market. Newer market participants are offering alternate ways for managing archived data that involve virtual copies of databases, extreme compression and native SQL access.

Magic Quadrant



Vendor Strengths and Cautions dataglobal

Dataglobal is a privately held company headquartered in Heilbronn, Germany, with offices in Boston, Massachusetts, and Romania. The company has been in the digital archiving business for 22 years. It offers dg hyparchive as its archiving backbone, with dg erp supporting connectors for SAP ERP systems and dg connect as a generic API for archiving additional application content (e.g., JD Edwards). Dataglobal includes SAP and other structured content as part of its enterprise information archiving platform, which includes records management and e-discovery capabilities such as tagging and legal hold. SAP content is supported via SAP ArchiveLink and stored in the dg repository along with other unstructured content, such as files and email. The product provides functionally rich support for SAP, including archiving data from any SAP Business Suite application, SAP NetWeaver Business Warehouse (BW) and other content associated with SAP.

Common services include automatic-retention policies, and search inside and outside the archive. Some analytics capabilities are available.

Dataglobal sells its product directly and through partners, including Dell. Most applications other than SAP are sold via partners that offer more customized solutions using dg connect. Eighty-five percent of the company's sales are in Europe, where its majority of SAP archiving installations are located. The product is sold as a cloud archiving service via T-Systems in Europe, and dataglobal is extending its support for the cloud via additional partners. Support is generally provided via partners that are experienced and know the product well. Pricing is based on the number of connectors (i.e., the number of SAP instances) or on volume for dg connect. To continue to grow in SAP archiving, dataglobal will look to expand partnerships through consultancies and resellers, such as the one established with

Dell. There is no indication that dataglobal will expand beyond SAP archiving to areas such as big data analytics or application retirement.

Strengths

- Strong classification, content analysis and metadata management capabilities are available through dg classification, supporting organizations that are looking for information compliance and governance as a primary use case.
- Support services, intuitive ease of use and stability garner positive responses from designated references.

Cautions

- The vendor does not support features such as graphical modeling of application data, creation of application-specific business rules and policies, data masking or database-to-database archiving.
- Most organizations are using the dg hyparchive for smaller structured data archive environments. The product has no references for archives ranging in size from hundreds of terabytes to a petabyte or more.

Data Migration

Data Migration, located in Kreuzlingen, Switzerland, has been selling its application archiving product, JiVS, since 1997. Data Migration's core competencies are in SAP application retirement. The product provides a platform approach with options such as retention management and data masking supported through configuration. The company takes a different approach to SAP archiving than other vendors; rather than use standard SAP ArchiveLink interfaces, JiVS employs unique plug-ins, but achieves similar results with SAP Archive Development Kit (ADK) files stored in a file system. Data Migration has added support for GDPdU, a standard for enforcing the audit ability of archived data and documents. The firm supports custom legacy applications, offering consulting or training services. JiVS supports a number of structured data use cases, including archiving, application retirement, data integration and validation.

Data Migration gains 70% of its revenue through value-added resellers (VARs) and the rest through direct sales; 90% of its total revenue is generated in Europe. The firm targets large enterprises in

Europe and relies on partners like T-Systems for support and deployment. Pricing is flexible, based on three models: per application plus database volume, per defined scope of applications to retire (fixed pricing), or per terabyte for retirement programs managed with JiVS. Data Migration's plans during 2014 include adding enhancements for application retirement management and common document archiving functions. In terms of its road map, Data Migration plans to strengthen its out-of-the-box configurations for nonretirement functions; this is an area that is highly needed.

Strengths

- Data Migration has a history of supporting application retirement for legacy applications — including SAP R/3, SAP CRM, SAP Utilities (IS-U), SAP Supplier Relationship Management (SAP SRM), JD Edwards, Siebel, Oracle Financials and BaaN — and is viewed favorably for its total cost of ownership (TCO) and fast time to value.
- For application retirement use cases involving large numbers of applications, Data Migration has demonstrated scalability.

Cautions

- JiVS is a platform, rather than an application, and features such as complex retention management and data masking are available through programmatic and customized means, causing a reliance on professional services.
- Data Migration is a small firm with the strategy to deliver projects via VARs. Needed resources, such as project management and support, can be stretched thin.

Delphix

Delphix was founded in 2008. Its initial products focused on testing/development use cases facilitated by database virtualization. Delphix recently announced support for application and database archiving, retirement and migration, in large part because its customers discovered that the product's database virtualization extended naturally to these use cases. The company offers the Delphix Agile Data Platform and Delphix Modernization Engine for archiving and retirement.

The product creates a full copy of all data associated with an application, including application binaries, associated files, virtual machines, the database and DBMS binaries and license keys, as a single secondary copy, which can be preserved for compliance or other purposes, and mounted as a virtual copy for near-immediate access. The product runs on VMware (delivered as a software appliance), and supports application migration to private, hybrid and public clouds, including IBM SoftLayer and SunGard. Because Delphix appears as a storage volume to application and database servers, access to archived data is via native database utilities. Retention is supported through the Live Archive feature. The TimeFlow capability supports transactionally consistent point-in-time access to historical database versions. Support for integrated masking was released in 1H14. The product can be used to preserve database data during a migration from one application version to another, or from Unix to Linux. Archiving SAP data is supported in the same way as other application data. Delphix's largest SAP customer supports 270TB of data as an archive. Delphix is an Endorsed Business Solution partner of SAP.

The product is predominantly sold directly. However, the company is increasing its archiving-specific channel with partners like Accenture, which is leveraging the product as part of its Application Retirement practice, and cloud-focused partners such as Cloud Technology Partners. Pricing is term-based (one year or three years) and per CPU. Customers license servers from Delphix and run as many copies as the server will support.

Strengths

- Delphix's data virtualization approach to application archiving with compression and filtering, coupled with integrated data protection features, enables significantly reduced storage costs (20-to-1 consolidation across database copies is typical).
- Delphix's responsiveness and technical acumen are strong points. Delphix has approached problems in creative ways to deliver what customers need.

Cautions

- Database support is currently limited: The product supports only Oracle, Microsoft SQL Server, PostgreSQL and Sybase Adaptive Server Enterprise (ASE).

- Delphix's solution is not appropriate for looking to store database content in enterprise information archiving systems that support flat-file repositories containing multiple content types. Delphix operates at the block level, rather than the database schema level, so archive granularity is per database.

EMC

EMC offers EMC InfoArchive, a new product entrant in the structured data archiving market. EMC InfoArchive became available as a product in 2013, but has existed as a service offering since 2009. InfoArchive contains elements of the Documentum enterprise content management (ECM) technology, and takes the approach of delivering an archive that brings unstructured and structured content together. InfoArchive addresses four distinct use cases: (1) archiving unstructured content, such as images and files, which includes report content such as print-stream-generated statements and bills through its partnership with Crawford Technologies; (2) archiving structured database content for which a partnership with Solix provides connectors to databases and prebuilt integrations to common packaged applications; (3) archiving relational database management system (RDBMS) tables; and (4) archiving complex information sets that aggregate multiple data records and content files to a single business record.

Structured data is archived to InfoArchive in a compressed XML format, and is searchable via a standard Web services API that leverages XQuery. Database-to-database archiving is not supported. InfoArchive has found initial success with application retirement in financial services, retiring legacy applications such as mainframe reporting systems, and is broadening its support for production applications and additional content sources via connectors built by its partner ecosystem. EMC provides Archive Services for SAP to deliver SAP ArchiveLink-based archiving of SAP data, documents and print lists as part of the Documentum portfolio.

EMC sells InfoArchive via its Information Intelligence Group sales team. EMC leverages its storage sales forces, including the Data Protection and Availability Division group, to identify opportunities that can be combined with archiving platforms such as EMC Centera and Data Domain. Pricing is per terabyte. The long-term vision for InfoArchive is to enhance and deliver

better analytics through integration with Pivotal, an EMC company. This holds promise as customers look to leverage, repurpose and use big data as a competitive advantage.

Strengths

- InfoArchive can support many different and unique content types, an important quality for application retirement use cases.
- The product offers strong retention/records management and indexing, and user-friendly advanced search.

Cautions

- InfoArchive is a new product that requires a solution-based approach that may need to rely on partners for success.
- Customers have cited issues with setting up backups.

HP

HP offers HP Structured Data Manager for database archiving and application retirement. The product is mature, due to its long history in this market dating back to 1998. HP had let the product languish for a few years; but, due to its belief in the current market opportunity, HP is making significant investments to regain market share in database archiving and application retirement.

Initially built for Oracle Databases, the product now supports a broader variety of applications and underlying databases, including Microsoft SQL Server and IBM DB2. Prebuilt Integration Packs support common enterprise applications; other apps are supported as custom implementations. Database-to-database and database-to-file archiving are both supported, and the product is integrated with HP Intelligent Data Operating Layer (IDOL), enabling search and retrieval across production and archive databases. Native applications, Open Database Connectivity (ODBC)/Java Database Connectivity (JDBC)-compliant reporting tools and HP's viewing technologies provide data access. The designer interface provides business rule modeling. Data masking is provided natively. On-premises, private, hybrid and public cloud (such as HP Cloud Services) implementations are supported. Apache Hadoop (Hadoop Distributed File System [HDFS]) is supported as an archive target. The product is

integrated with HP Records Manager (formerly HP TRIM software) for records management. SAP archiving is supported via SAP ArchiveLink.

HP Structured Data Manager is sold directly from an information management and governance sales team formed in late 2013, as well as through partners. It is offered as part of the HP Autonomy information governance portfolio, or stand-alone. Sales teams supporting HP's backup, recover and archive solutions, as well as the newer HAVen technology stack, are also enabled to sell and support the product. Pricing is per source database, regardless of source type or data volume. HP likely will have to attract and retain talent in professional services in structured data archiving to be successful, as the firm neglected this area during the 2009 to 2012 time frame.

Strengths

- The product has demonstrated success with cloud implementations. A number of reference customers have deployed it alongside HP Autonomy's enterprise information archiving products in the HP Autonomy cloud.
- Customer support, responsiveness, knowledge and product flexibility are positive attributes, particularly compared with competitive offerings. HP runs the solution in its IT department to support these claims.

Cautions

- No native deduplication capabilities are provided. Deduplication must be provided by the archive target.
- Structured Data Manager's performance with respect to archiving very large volumes of data was not as fast as small implementations. Performance is typically faster if data is limited.

IBM

IBM delivers application retirement, structured data archiving, data masking and test data management through InfoSphere Optim. SAP data archiving is available through IBM Content Collector for SAP. These products are often packaged in conjunction with other InfoSphere offerings as part of information life cycle governance (ILG), a holistic data governance strategy typically aimed at large enterprises.

IBM InfoSphere Optim provides structured data archiving for hundreds of common off-the-shelf and custom applications, including distributed systems and mainframes. Archived data is stored in highly compressed binary files supported by a large selection of storage repositories; including write once, read many (WORM) storage platforms. InfoSphere Optim provides e-discovery functionality, including legal hold, and integrates with IBM Atlas eDiscovery, as well as a variety of third-party e-discovery offerings. IBM offers robust indexing and access to data via a Google-like interface for searching, ODBC and JDBC reporting tools, or its native UI. InfoSphere Optim ships with a supported version of IBM's Hadoop distribution InfoSphere BigInsights, and can be integrated with third-party Hadoop distributions. Tight integration between IBM PureData Systems for Analytics and InfoSphere Optim provides built-in capabilities to archive directly to IBM PureData Systems for Hadoop. InfoSphere Optim analyzes relationships across various repositories to provide a unified view of data called Complete Business Objects. The portfolio includes dormant data analysis that allows administrators to identify which data and applications should be archived. Integrated data masking is supported. IBM InfoSphere Optim's largest customer has archived more than 1PB of data. The IBM Content Collector for SAP, a component of IBM's enterprise content management portfolio, leverages the SAP ADK and SAP ArchiveLink functionality to store data and documents in an IBM FileNet repository.

In an effort to simplify pricing, IBM recently moved to a capacity-based pricing model, and offers InfoSphere Optim in an Enterprise Edition and a Workgroup Edition. InfoSphere Optim is primarily a direct sale through IBM and is available on-premises, in a hybrid model and in a cloud-based model with third-party cloud offerings or as part of IBM SoftLayer. An appliance is available for archiving high data volumes. IBM likely will look to simplify its cloud delivery options and seek to get greater leverage from its SoftLayer acquisition.

Strengths

- The IBM InfoSphere Optim team has broad vertical industry awareness, which has translated into focused implementation and consulting engagements.
- The product is mature and feature-rich, including support for a broad set of data sources, ease of deployment, and support for business rules as part of the archive and data access processes.

Cautions

- Training, documentation and online resources are areas in need of improvement.
- Database-to-database archiving is supported; however, it requires a two-step process whereby data is archived to proprietary compressed immutable binary files first, then to another database of the same or different type.

Informatica

Informatica sells Informatica Data Archive. The product supports a wide array of applications and underlying databases resident on distributed systems or mainframe platforms, including Oracle, SQL Server and IBM DB2. The product also supports archiving from Apache Hadoop, IBM PureData System for Analytics and Teradata, as well as SaaS applications, such as from salesforce.com. Informatica includes a capability it calls Smart Partitioning, which allows administrators to place related records across tables into a single database table space, and then effectively exclude this partition from database production operations, including queries and reports. This "archive in place" ensures that the database remains intact, retains data for compliance and can improve performance.

The product supports database-to-database and database-to-file archiving (leveraging a columnar compressed/deduplicated format). The archives can be accessed via the original application, SQL/ODBC/JDBC reporting tools, keyword-based search or an integrated report builder. Retention, legal hold and masking are supported at the record level. Utilities for monitoring data growth and application performance are also available. The latest version of the product has been integrated with the company's flagship Informatica platform. SAP archiving is via SAP ADK and SAP ArchiveLink. Informatica ILM Nearline is available to archive data from SAP NetWeaver Business Warehouse to the Data Archive repository. Private cloud deployments are supported. The company plans to release support for the public cloud in 2014.

Informatica benefits from having a large customer base using its PowerCenter products. Customers choose the data archiving product based on their relationship with Informatica, as well as the new product integration. The majority of sales are direct, and Informatica has partnered with Oracle for a joint solution focused on the Data Archive Smart Partitioning functionality. The product is available in a Standard Edition and an

Advanced Edition differentiated by features, with a number of additional cost options available. In 2014, Informatica will look to deliver vertical applications in areas such as healthcare and cloud-based offerings.

Strengths

- Informatica offers enterprise structured data archiving for a broad set of applications, packaged and customized. The solutions have the ability to archive on-premises and cloud applications, such as from salesforce.com.
- Database connectivity, reporting and visualization, and compliance functionality, legal hold, retention management and auditing are strong points.

Cautions

- Informatica Data Archive has complex pricing. The total cost of the overall solution can be an inhibitor to potential buyers.
- Informatica's Data Archiving implementation and consulting teams are areas with room for improvement, specifically with respect to knowledge transfer.

OpenText

OpenText offers OpenText Archiving for SAP Solutions, including OpenText Archiving for SAP and OpenText Document Access for SAP, for archiving SAP data and content. The SAP archiving application is based on the 14-year-old archiving technology the company acquired when it bought IXOS Software in 2004. OpenText complements SAP archiving by offering archiving solutions for unstructured content (including email, files and SharePoint). In addition to the SAP archiving products, OpenText offers the InfoFusion Integration Center and InfoFusion Discovery Platform for application decommissioning, data archiving and information governance for non-SAP RDBMS content. The products offer consolidated extraction, transformation and loading (ETL), advanced search, content analytics, and connectors to a wide variety of structured and unstructured data sources. Content from all products is stored in the OpenText Content Suite. Retention management, auto-classification, sampling for quality assurance, legal hold and various workflows are among the features included in the solutions. Direct access to archived SAP data is supported via SAP ArchiveLink. Access to data stored in the InfoFusion database is accessible using

various methods, including browsing, search and reporting. Comprehensive records management is included in the solution. The archiving products can be deployed on-premises or in a hybrid cloud. Amazon Web Services and Microsoft Azure are supported as target repositories.

OpenText is the market share leader for SAP archiving and has over 2,000 organizations using these products. Data archiving and document access are sold directly by OpenText, and resold by SAP as SAP Archiving by OpenText and SAP Document Access by OpenText, providing a go-to-market advantage when buyers seeking archiving solutions work directly through SAP. OpenText's SAP archiving products are priced per named SAP user. InfoFusion Discovery Platform pricing is based on a combination of connectors and seats for various functions, such as search, classification and content remediation. OpenText will look to build more connectors and integration points to expand into application retirement use cases, a market driver for structured data archiving.

Strengths

- OpenText has strong records and retention management capabilities aligned with its mature SAP products.
- OpenText provides an integrated enterprise content management and enterprise archiving (both are structured and unstructured content) approach that enables enterprises to manage a wide array of content types and applications through common interfaces.

Cautions

- For organizations seeking an archive system without complementary, comprehensive ECM, OpenText software and professional services are expensive.
- For non-SAP structured archiving, there is no support for archiving based on the determination of application business logic. Knowledge of the database schema is required.

PBS Software

PBS Software (PBS) has been offering SAP archiving solutions since 1991. PBS is focused 100% on supporting SAP and does not offer solutions for other application types or platforms. It provides a broad set of solutions and utilities for system decommissioning, data extraction, data

analysis and dedicated industry-specific solutions in SAP environments. PBS ContentLink supports access to archived content directly through the native SAP interface, is SAP ArchiveLink and WebDav certified by SAP (BC-ILM 3.0) for SAP Information Lifecycle Management (ILM)-aware storage, and maps directly to all SAP ERP modules. PBS CBW NLS and NLS IQ provide “nearline” support for SAP NetWeaver Business Warehouse. Archived data is stored in a compressed, deduplicated format in SAP IQ or Actian Vectorwise as a nearline repository. It offers retention management and legal-hold support in addition to integrating with SAP ILM. PBS products are deployed in midsize and very large enterprises.

PBS sells its offerings mainly to Europe-based organizations, with 68% of its customers located in EMEA, 25% in North America and 7% in Asia/Pacific region. PBS is supported by VARs such as Dolphin Enterprise Solutions (aka Dolphin) in the U.S., through which it has a number of large corporate clients. Pricing is per SAP production system, client and named user for the ERP offering, and per SAP NetWeaver Business Warehouse users for that offering. In 2014, PBS will look to enhance its Nearline Analytic Infrastructure solution; this proposed offering will support SAP systems with classical databases, help prepare migration to Hana and work with Hana as a nearline database.

Strengths

- PBS has a great deal of knowledge about SAP archiving, with more than 20 years of experience working in and selling to the SAP community.
- User adoption of PBS’s products is high, as it leverages SAP user experiences and ties directly with SAP.

Cautions

- Some reference customers would like PBS to provide more round-the-clock hours of technical support, rather than only European time zone support.
- PBS products require more deployment time than average, compared with structured data archiving vendors for similar capabilities and capacities.

RainStor

RainStor is a privately held company founded in 2004. It offers what the company calls an “analytical archive” for compliance, historical analysis or application retirement. At the core, the product provides a highly compressed row/columnar hybrid repository that enables up to 20x to 40x compression along with fast access and retrieval. A wide variety of underlying databases are supported, as well as Teradata, SAP Sybase IQ, IBM PureData System for Analytics and Oracle Exadata data warehouses. Access to archived data is via native SQL with database-specific extensions. The product supports data retention, tagging and legal hold at the record level, as well as various WORM storage platforms. It currently offers native support for Apache Hadoop (HDFS) as an archive target, and data masking/encryption is included with the product. Private, hybrid and public clouds (Amazon Simple Storage Service [S3]) are supported. SAP archiving is not supported, except through partners.

The product is typically deployed in very large environments. Customers’ archive capacity is generally more than 300TB. Many of the company’s customers are managing more than a petabyte of data either across a variety of applications or within one application. Applications supporting large volumes of small objects, such as call data records or financial trades, are well-suited for RainStor.

The company has historically sold almost exclusively through partners and OEMs (as an embedded repository component for other archiving products), but is moving to increase its direct sales channel in 2014. The company supports deployments of prepackaged application archives with Informatica, IBM and Solix. RainStor has a sales partnership with Teradata. Its sales partnership with EMC was recently enhanced due to integration with EMC Isilon to support a SQL database running on scale-out network-attached storage (NAS) or native Apache Hadoop. Pricing is by capacity (raw TB, uncompressed). In its latest product release, RainStor added support for Securities and Exchange Commission (SEC) Rule 17A-4 WORM-compliant storage, which should align well with RainStor’s strategy to work with EMC and support financial services.

Strengths

- RainStor's market-leading compression, along with its integration with leading storage platforms, data warehouses and Hadoop, makes it particularly suitable for leading-edge big data environments with large volumes of data.
- RainStor gets extremely high marks from reference customers for customer service and support.

Cautions

- While RainStor provides retention, legal hold and other repository management functions, it does not provide front-end application business logic for identification of data for archiving, and must rely on partners for this capability, if required.
- No UI is available for management operations, including archiving processes. Interaction with the system is via command line interface or Java API. ODBC/JDBC interfaces are supported for query.

RSD

RSD is a 40-year-old privately held firm headquartered in Switzerland. It offers RSD Glass, which is marketed as an information governance solution, and RSD Glass Repository (previously named RSD Folders), which is embedded as the optional archiving repository. RSD has more than 1,200 customers using its archiving solutions, with a small subset storing SAP archived data.

RSD ingests data from legacy applications, ECM systems, SharePoint and other systems into RSD Glass for classification, information governance and information lifecycle management (ILM). RSD leverages the SAP ArchiveLink protocol to ingest SAP data and to provide reporting access by SAP for the archived data. Archived data is stored in the original file format, and the repository provides legal hold capabilities for archived information. The RSD policy engine enables the enforcement of policies across RSD and non-RSD repositories. RSD Glass can be installed on-premises, hosted by RSD or deployed in a cloud-based model.

RSD's licensing is based on a number of variables that includes users, platforms and connections. Fifty percent of its revenue comes from EMEA, with 65% sold direct through RSD. The largest RSD SAP

archive has approximately 70TB of data, and RSD has clients with multiple petabytes of information archived. Clients have cited the reliability of RSD's products as a strong point. RSD doesn't have tangible improvements planned for structured data archiving in its road map, and is not a primary focus for the company.

Strengths

- RSD Glass has the ability to view data natively from applications.
- RSD's ability to manage structured, unstructured and physical records is a unique strength.

Cautions

- RSD's stand-alone archiving business is largely a legacy one, and the company is more focused on archiving through its information governance strategy with RSD Glass.
- RSD's interface lacks some user-friendly features, such as the ability to drag and drop files into and out of RSD.

Solix Technologies

Solix Technologies (Solix), based in Santa Clara, California, sells an archiving product called EDMS. Solix has demonstrated an ability to scale down to meet small application archiving opportunities while also actively competing for large-enterprise deals. Solix supports numerous databases, with an emphasis on active archiving in Oracle environments. It provides out-of-the-box support for many packaged and custom applications, including JD Edwards, Oracle E-Business Suite and PeopleSoft. It has a small SAP customer base, and plans to increase support for SAP archiving in 2015. Solix supports database-to-database and database-to-file archiving, the latter to both comma-separated values (CSV) and XML file formats. Columnar compression/deduplication is available via integrated Sybase IQ or RainStor, and archiving of database partitions is supported. Solix can support Apache Hadoop as a repository for unstructured and structured data archiving and application retirement. Archived data is accessed via the original application when data is stored in an archive database via full-text search or the Solix Application Portfolio Manager, which provides access to data from retired applications without the requirement for SQL queries. Integrated masking and encryption are available.

In 2014, Solix is adding support for unstructured content and looking to grow its customer base through partnerships with EMC for its InfoArchive product, and with Kronos, which embeds Solix EDMS into its workforce management suite. Solix can be delivered as an appliance called the Solix ExAPPS or as software, and includes support for cloud-based archiving through Rackspace or Amazon. The product price is based on volume, with an additional charge for prebuilt archiving templates; subscription pricing is available. A free downloadable version is available on the Solix website.

Strengths

- Solix's ease of use, technical support and strong professional services are all strengths — key areas when dealing with legacy application environments.
- The company offers the widest variety of deployment models (including on-premises, cloud and appliance) and pricing options (including perpetual, subscription and appliance) compared with other vendors in this market.

Cautions

- For custom enterprise applications, some references have cited slow ingestion speeds for EDMS as an area in need of improvement.
- Solix's EDMS is a platform and needs more out-of-the-box configurations for integration with databases and applications.

ZL Technologies

ZL Technologies (ZL) focuses primarily on database-to-database archiving, with a history of retiring custom legacy applications, in addition to IBM Notes (formerly Lotus Notes) applications and databases. ZL Unified Archive is built with a grid architecture that supports the large-scale environments the company targets. The product includes support for archiving Oracle and Microsoft SQL Server data as part of its single unified repository strategy. The company believes that archive repositories will increasingly support structured and unstructured content, and continues to increase the variety of content types supported, including a Hadoop-enabled infrastructure for easy interface with the Hadoop ecosystem.

Support for application data is via an API enabling the creation of connectors into structured data sources or by using ZL DirectExtract. ZL DirectExtract is a utility that can be used to identify metadata, database schemas and other database constructs to create a metadata model to extract data from source systems. Archived data is stored in ZL's file system, which provides support for compression, encryption and various security features. For all archived data, ZL Unified Archive provides comprehensive records management, compliance, analytics and e-discovery support in one unified platform.

ZL targets large-enterprise organizations as customers. ZL Unified Archive is available direct from the company and through channel partners, as well as a number of large, specialty cloud providers (including RenewData, SunGard and Viewpointe) that utilize ZL as the technology powering their cloud archiving services. Pricing is per GB for on-premises deployments, and either monthly per user or per GB for cloud deployments.

Strengths

- ZL's retention management, legal hold and left-hand-side e-discovery capabilities are well-suited for regulated environments such as financial services, where audits, litigation and investigation are commonplace.
- ZL Technologies is very responsive to customer issues, and product support is rated highly according to reference customers.

Cautions

- Some capabilities for structured data archiving, such as support for SAP, are still on the company's road map and slated for release.
- ZL Technologies is a new participant in the structured data archiving market and has fewer reference customers in unstructured content archiving compared with other vendors in this market.

Vendors Added and Dropped

We review and adjust our inclusion criteria for Magic Quadrants and MarketScopes as markets change. As a result of these adjustments, the mix of vendors in any Magic Quadrant or MarketScope may change over time. A vendor's appearance in

a Magic Quadrant or MarketScope one year and not the next does not necessarily indicate that we have changed our opinion of that vendor. It may be a reflection of a change in the market and, therefore, changed evaluation criteria, or of a change of focus by that vendor.

Inclusion and Exclusion Criteria

To be included in the Magic Quadrant for Structured Data Archiving and Application Retirement, a vendor must:

- Offer products that meet the definition for structured data archiving and application retirement detailed in the Market Definition/Description section of this report.
- Be the developer of the product, and not just a reseller or VAR.
- Support a growing base of customers, including at least 15 enterprise customers that are using the software in a production environment.
- Have a presence in at least two geographies (North America, EMEA, Asia/Pacific region, South America) worldwide and be industry-independent.
- Provide its solution as an on-premises software product, a SaaS offering or some combination.
- Achieve more than \$3 million in new license or maintenance revenue annually.

Evaluation Criteria

Ability to Execute

Product: An evaluation of the features and functions of the vendor's structured data archiving solution, including those related to:

- Archiving to an alternate (nonproduction) database or file format.
- Maintaining referential integrity (even for the most complex data models).
- Seamless access to archived data from the original application or via alternate methods (search, reporting).
- Security, access control and audit logs.
- The road map should support plans for big data initiatives and analytics, including Apache Hadoop.

Higher ratings are:

- Assessed for support for data validation, broad application support (including for custom and legacy applications), data retention and purge management, data discovery, data masking and test data management, and support for legal hold.
- Assigned to solutions with strong archive architectures, policy-based archiving and storage management features, quality of user experience, and support for unstructured content.

Overall Viability: Includes an assessment of the vendor's overall financial health, the financial and practical success of the structured data archiving business unit, and the likelihood of the individual business unit to continue to invest in a structured data archiving solution.

Sales Execution/Pricing: The vendor's capabilities in all sales activities, and the structure that supports them. This includes pricing and negotiation, presales support and the overall effectiveness of the sales channel.

Market Responsiveness/Track Record: Includes the ability to respond, change direction and be flexible as market dynamics vary. This criterion also considers the vendor's history of responsiveness.

Marketing Execution: The effectiveness of the vendor's marketing programs, and its ability to create awareness and mind share in the market. It assesses whether the messaging is clear, whether the vendor provided references that used the unique features of the product in its target environment, and whether the promotion of the product on the company website is effective.

Customer Experience: The quality of the customer experience based on reference calls and Gartner client teleconferences (inquiry).

Operations: The ability of the organization to meet its goals and commitments in an efficient manner. Past performance is weighted heavily.

Table 1. Ability to Execute Evaluation Criteria

Evaluation Criteria	Weighting
Product or Service	High
Overall Viability	High
Sales Execution/Pricing	High
Market Responsiveness/Record	Medium
Marketing Execution	High
Customer Experience	High
Operations	Medium
Source: Gartner (June 2014)	

Completeness of Vision

Market Understanding: Ability of the vendor to understand buyers' needs, and to translate those needs into the appropriate features in the structured data archiving product, along with the ability to anticipate market trends (for example, the requirement to support heterogeneous applications and databases, including SAP, e-discovery or unstructured content) and to adapt quickly via new features, partnerships or acquisitions.

Marketing Strategy: A clear set of messages that positions the product and differentiates it from competitors, consistently communicated throughout the organization and externalized through the website, advertising, customer programs and positioning statements.

Sales Strategy: The vendor's strategy for selling to its target audience, including an analysis of the appropriate mix of direct and indirect sales channels.

Offering (Product) Strategy: An evaluation of the vendor's strategic product direction, including an analysis of its road map.

Business Model: The soundness and logic of a vendor's underlying business proposition.

Vertical/Industry Strategy: The vendor's strategy for meeting the specific needs of individual vertical markets and market segments (for example, financial-industry-regulated employee supervision, or state and local government information retention and disclosure requirements).

Innovation: The vendor's product leadership and ability to deliver archiving features and functions that distinguish the vendor from its competitors.

Geographic Strategy: The vendor's strategy for penetrating geographies outside its home or native market.

A vendor's Completeness of Vision is evaluated based on its ability to convincingly articulate its product direction and demonstrate innovation in meeting customer needs, enabling the vendor to more effectively compete in the market. The credibility of a vendor's vision is weighed against its past Ability to Execute and against previously stated plans. Market understanding should be the guiding factor in new product development to ensure that the engineered product meets customer needs. Managing the complexity of archiving environments requires innovative approaches that will distinguish leaders and delight customers.

Table 2. Completeness of Vision Evaluation Criteria

Evaluation Criteria	Weighting
Market Understanding	High
Marketing Strategy	Medium
Sales Strategy	Medium
Offering (Product) Strategy	High
Business Model	Low
Vertical/Industry Strategy	Low
Innovation	Medium
Geographic Strategy	Low
Source: Gartner (June 2014)	

Quadrant Descriptions Leaders

Leaders have the highest combined measures of Ability to Execute and Completeness of Vision. They may have the most comprehensive and scalable products. They have a proven track record of financial performance and an established market presence. In terms of vision, they are perceived as thought leaders, with well-articulated plans for ease of use, product breadth and how to address scalability. For vendors to have long-term success, they must plan to address the expanded market requirements for structured data archiving and application retirement, including support for Apache Hadoop and big data, support for the cloud, solid relevant SAP archiving functionality, and a strong administrative UI.

Leaders must not only deliver to current market requirements, which continue to change, but also anticipate and begin to deliver on future requirements. A cornerstone for Leaders is the ability to articulate how these requirements will

be addressed as part of their vision for expanded archive management. As a group, Leaders are considered part of most new-purchase proposals, and have high success rates in winning new business. There are four Leaders in this Magic Quadrant: IBM, Informatica, HP and Solix Technologies.

Challengers

Challengers can execute today, but have a limited or an evolving vision. They have capable products and can perform well for many enterprises. These vendors have the financial and market resources and capabilities to become Leaders, but may have elected to focus more heavily on one vertical industry or one structured data archiving use case. The sole Challenger in this Magic Quadrant is Delphix.

Visionaries

Visionaries are forward-thinking, but their execution has not propelled them into a leadership position. These vendors are differentiated by product innovation, but they have not achieved the sales and marketing success required to give them the high visibility of Leaders. In the case of this Magic Quadrant, they may be hampered by their product immaturity or lack of structured data archiving features and capabilities. The Visionaries in this Magic Quadrant are EMC and RainStor.

Niche Players

Niche Players are narrowly focused on an application type, such as SAP, offer some degree of structured data archiving as an adjunct to enterprise information archiving (unstructured content archiving, such as email, files or SharePoint) or offer broad capabilities without the relative success of their competitors in other quadrants. This is acceptable for a number of buyers, and some of the Niche Players' offerings are used successfully by very large global enterprises. Niche Players may focus on a segment of the market and do it well, or they may simply have modest horizons and lower overall capabilities compared with competitors. Others are simply too new to the market or have fallen behind, and, although they're worth watching, they have not yet developed complete functionality or the Ability to Execute. Niche Players in this Magic Quadrant are dataglobal, Data Migration, OpenText, PBS Software, RSD and ZL Technologies.

Context

Placement on the Magic Quadrant is based on Gartner's view of a vendor's performance against the criteria noted in this research. Gartner's view regarding vendor placement on the Magic Quadrant is heavily influenced by surveys completed by the vendors, and several hundred inquiries and one-on-one conversations at Gartner conferences conducted during the past 12 months with our clients on the topic of structured data archiving. The Magic Quadrant methodology includes the solicitation of references from each vendor, and Gartner then conducts reference checks from a set of those customers.

This Magic Quadrant does not rate only a product's quality, capabilities and features. The product is an important part of the rating, but the vendor's ability to acquire and support customers is equally important, as is its ability to grow product and service revenue. A vendor that offers a strong, technically elegant product, but is unable or unwilling to invest in marketing and sales to generate revenue and growth, will find itself unable to invest sufficiently in development.

Market Overview

Based on Gartner's estimates, the size of the structured data archiving and application retirement market is \$270 million, and growing at a compound annual growth rate (CAGR) of 10%. The use of this technology has long been viewed as a cost avoidance measure to contain operational and capital expenditures related to data growth, and as a measure to improve factors such as application performance. The market is changing and expanding due to growth in data, application retirement, information governance and big data analysis opportunities.

Trend Toward Big Data Analytics and Petabyte-Scale Archives

The growing use of Apache Hadoop, increasing data warehouse volume sizes and the accumulation of legacy systems in organizations are fostering structured data growth. These factors are leading enterprises to understand how to reuse, repurpose and gain critical insight from this data. Apache Hadoop is capable of storing large volumes of data. Thus, organizations are using HDFS to store structured data, as well as information such as social and machine data that doesn't fit into

databases. Many organizations are looking to add structure and meaning to this information repository, beyond just using it as a low-cost means of storage. Structured data archiving vendors have responded by adding support for Apache Hadoop as a data source and a target. Gartner expects to see this emerging requirement for Hadoop support going beyond baseline storage management to include support for more analytic tools (for example, from Tableau Software) and other reporting mechanisms to the point where the line between archiving and active use will blur. Big data analytic tools will become a baseline component of structured data archiving tools by 2016. The various distributions of Hadoop, such as Hortonworks and Cloudera, are increasingly emphasizing information life cycle and retention management in Hadoop, which will put pressure on structured data archiving vendors to innovate further in this area.

Growing Importance of Information Governance in Structured Data

Structured data from applications is an easy target for external auditors. They are experienced in identifying the relevant applications and lack of controls that may occur in protecting valuable financial data. In some respects, it's an easier task than identifying unstructured content, such as spreadsheets that contain financial data that may be scattered and managed lightly in the enterprise. Auditors may raise a red flag if the legacy application is so old that it's no longer supported or loosely managed. Migrating data to an upgraded version of the application or to an alternate format may mitigate this problem of maintaining structured data.

Most of the IT focus on preparing for and responding to e-discovery requests has been for unstructured data. However, there have been numerous cases where structured data has been a target for discovery requests. The discovery of structured data presents challenges, and organizations want to ensure they can respond quickly to requests for information when that information is not accessible from its native application. By taking an active and systematic approach to application retirement, organizations can purge data that no longer has business relevance, not only to reduce costs for maintenance, but also to reduce the cost of responding to e-discovery requests by making data more searchable, defensible and easier to preserve.

Application Retirement as a Leading Use Case for Structured Data Archiving

Organizational mergers and acquisitions, data center consolidation, and migration to cloud-based applications have accelerated the requirement to retire legacy and redundant applications. Application retirement presents numerous cost benefits and efficiencies that further fuel this trend.

Although the storage savings and positive effects of reduced complexity are highly attractive, the relationship between applications and data makes application retirement a highly complex task. Enterprises must understand and develop requirements for data access and long-term retention, and execute policies based on those objectives. Identifying candidates and developing a business case for retirement based on potential costs savings must accompany these efforts.

Structured data archiving solutions can help in application retirement. Application retirement typically involves the transfer and retention of the underlying database and requires consideration of a number of factors, including ongoing access requirements, preservation of data and its business logic, governance and retention requirements, and data storage. In response to application retirement trends, structured data archiving vendors have developed solutions to retire legacy applications and their associated infrastructure. Greater interest in application retirement is contributing to the growth of the structured data archiving market.

Role of SAP in Structured Data Archiving and Retirement

The structured data archiving market includes solutions that archive data from applications such as those from SAP. As is the case for any application, previous SAP instances need to be retired and the data managed systematically for cost and governance reasons, with support for ongoing access to data. In many instances, vendors such as IBM and Informatica provide solutions for archiving directly from databases, as well as active archiving, and application retirement for SAP and other ERP and CRM applications. Numerous vendors are certified and support SAP archiving for active archiving through the SAP ADK and XML Archive API. Gartner receives a steady stream of inquiries asking about alternatives, indicating that, although solutions exist, lack of credible SAP expertise and high prices have plagued adoption. Like the interest in application

retirement, some vendors have identified these gaps in a well-established market and are making inroads against long-established players. As part of the Magic Quadrant for Structured Data Archiving and Application Retirement, Gartner evaluated and identified SAP archiving and retirement solutions.

Vendors to Watch

In addition to the 13 vendors evaluated in this Magic Quadrant, numerous other vendors offer archiving products specifically for structured data. The following list includes vendors that provide, or have plans to provide, support for structured data archiving and application retirement:

- **Actifio** has a copy data virtualization platform called Copy Data Storage (CDS), which takes a different approach than other vendors to managing structured data. It looks at long-term retention of backup data that can be actively used and repurposed. Actifio supports databases such as MS SQL, as well as business applications such as SAP and Oracle E-Business Suite.
- **CommVault** provides Simpana, a single-platform approach to backup and archiving, and supports data and document archiving for SAP modules. CommVault has stated its road map intentions for supporting additional structured data archiving capabilities in future product releases.
- **Gimbal**, with its ERP-Link product, takes the approach of enhancing and using Microsoft SharePoint as a strategic enterprise repository for managing content, including SAP data and documents. Gimbal provides strong domain expertise and technology related to records and retention management that can be applied to structured data.
- **SAP** offers NetWeaver Information Lifecycle Management, which archives SAP data and provides retention management capabilities. NetWeaver Information Lifecycle Management helps organizations comply with audit and compliance requirements, and consolidates SAP instances.

Evidence

“Forecast: Enterprise Software Markets, Worldwide, 2010-2017, 3Q13 Update” and supporting Gartner research.

Evaluation Criteria Definitions

Ability to Execute

Product/Service: Core goods and services offered by the vendor for the defined market. This includes current product/service capabilities, quality, feature sets, skills and so on, whether offered natively or through OEM agreements/partnerships as defined in the market definition and detailed in the subcriteria.

Overall Viability: Viability includes an assessment of the overall organization’s financial health, the financial and practical success of the business unit, and the likelihood that the individual business unit will continue investing in the product, will continue offering the product and will advance the state of the art within the organization’s portfolio of products.

Sales Execution/Pricing: The vendor’s capabilities in all presales activities and the structure that supports them. This includes deal management, pricing and negotiation, presales support, and the overall effectiveness of the sales channel.

Market Responsiveness/Record: Ability to respond, change direction, be flexible and achieve competitive success as opportunities develop, competitors act, customer needs evolve and market dynamics change. This criterion also considers the vendor’s history of responsiveness.

Marketing Execution: The clarity, quality, creativity and efficacy of programs designed to deliver the organization’s message to influence the market, promote the brand and business, increase awareness of the products, and establish a positive identification with the product/brand and organization in the minds of buyers. This “mind share” can be driven by a combination of publicity, promotional initiatives, thought leadership, word of mouth and sales activities.

Customer Experience: Relationships, products and services/programs that enable clients to be successful with the products evaluated. Specifically, this includes the ways customers receive technical support or account support. This can also include ancillary tools, customer support programs (and the quality thereof), availability of user groups, service-level agreements and so on.

Operations: The ability of the organization to meet its goals and commitments. Factors include the quality of the organizational structure, including skills, experiences, programs, systems and other vehicles that enable the organization to operate effectively and efficiently on an ongoing basis.

Completeness of Vision

Market Understanding: Ability of the vendor to understand buyers' wants and needs and to translate those into products and services. Vendors that show the highest degree of vision listen to and understand buyers' wants and needs, and can shape or enhance those with their added vision.

Marketing Strategy: A clear, differentiated set of messages consistently communicated throughout the organization and externalized through the website, advertising, customer programs and positioning statements.

Sales Strategy: The strategy for selling products that uses the appropriate network of direct and indirect sales, marketing, service, and communication affiliates that extend the scope and depth of market reach, skills, expertise, technologies, services and the customer base.

Offering (Product) Strategy: The vendor's approach to product development and delivery that emphasizes differentiation, functionality, methodology and feature sets as they map to current and future requirements.

Business Model: The soundness and logic of the vendor's underlying business proposition.

Vertical/Industry Strategy: The vendor's strategy to direct resources, skills and offerings to meet the specific needs of individual market segments, including vertical markets.

Innovation: Direct, related, complementary and synergistic layouts of resources, expertise or capital for investment, consolidation, defensive or pre-emptive purposes.

Geographic Strategy: The vendor's strategy to direct resources, skills and offerings to meet the specific needs of geographies outside the "home" or native geography, either directly or through partners, channels and subsidiaries as appropriate for that geography and market.

About Us

Solix Technologies, Inc., the leading provider of Enterprise Data Management (EDM) solutions, is transforming information management with the first enterprise archiving and data lake application suite for big data: The Solix Big Data Suite. Solix is helping organizations learn more from their data with enterprise analytics and achieve Information Lifecycle Management (ILM) goals. The Solix Enterprise Data Management Suite (Solix EDMS) and Solix Enterprise Standard Edition (SE) enable organizations to improve application performance, meet compliance objectives and reduce the cost of data management across the enterprise. Solix Technologies, Inc. is headquartered in Santa Clara, California and operates worldwide through an established network of value added resellers (VARs) and systems integrators.



Biography

Vikram Gaitonde heads the technology leadership and product strategy for Solix as Vice President of Products. Vikram is responsible for product management and engineering for Solix's next generation information governance solutions. Mr. Vikram brings 15 years of product management, engineering, and software delivery experience and has held several leadership positions within technology companies. Prior to Solix he was the Director of Enterprise Solutions at EMC, where he led the product solutions and product management teams for the information intelligence, data governance, eDiscovery, and content management products. Additionally, he has held numerous technical and product leadership positions at Silicon Valley startups and industry giants including Hotmail/Microsoft, Wipro, Kazeon (acquired by EMC), and Cast Iron Systems (acquired by IBM). Mr. Vikram holds an MBA from The Wharton School of the University of Pennsylvania and a B.S. in Computer Science from the University of Pune, India.