SOLIX®
Empowering Data Management

Enterprise Archiving with Apache Hadoop

Featuring the 2015 Gartner Magic Quadrant

Gartner®

Executive Summary

Every CIO wants to know if their infrastructure will handle it when data growth reaches 40 zettabytes by 2020.

When data sets become too large, application performance slows and infrastructure struggles to keep up. Data growth drives increased cost and complexity everywhere, including power consumption, data center space, performance and availability.

System availability is impacted as batch processes are no longer able to meet scheduled completion times. The "outage windows" necessary to convert data during ERP upgrade cycles may extend from hours to days.

Other critical processes such as replication and disaster recovery are impacted because more data is just harder to MOVE and COPY.

Left unchecked, data growth may also create governance, risk, and compliance challenges. HIPAA, PCI DSS, FISMA, and SAS 70 mandates all require that organizations establish compliance frameworks for data security and compliance.

According to Gartner, data growth is the No. 1 infrastructure challenge for data centers.

We know the value of data declines with age because it becomes less active. Enterprise data must be managed so inactive data doesn't clog the infrastructure and impact critical processing.

On the other hand, enterprise data is no longer confined to enterprise data centers.

Business critical data grows outside of the firewalls — with social media sites, blogs, hosted CRMs, etc. Enterprises must manage these data sources in order to stay relevant in the competitive world.

The tension between wanting more data to drive the organization successfully into the future and the need to keep infrastructure running efficiently and cost effectively has never been greater. How do organizations harness all the necessary and complex strains of data without over burdening infrastructure and personnel?

The solution is Information Lifecycle Management (ILM), which is the data management best practice to manage the lifecycle of data from creation to deletion and disposal.

The goals of ILM are:

- **Optimize application performance,**
- **Manage data security, risk, compliance,**
- **Reduce infrastructure costs.**

ILM achieves these goals by assigning retention policies to data based on business rules. Solix ILM moves data to the most appropriate infrastructure tier based on retention policies such as the age of the data. Since older data is less frequently accessed; it is therefore less valuable and less deserving of limited tier one performance and capacity.

Enterprise applications such as ERP, CRM, and HCM represent an excellent opportunity for improving performance and reducing costs through application tiering with Apache Hadoop.

FIGURE 1 Best Practice for Application Tiering

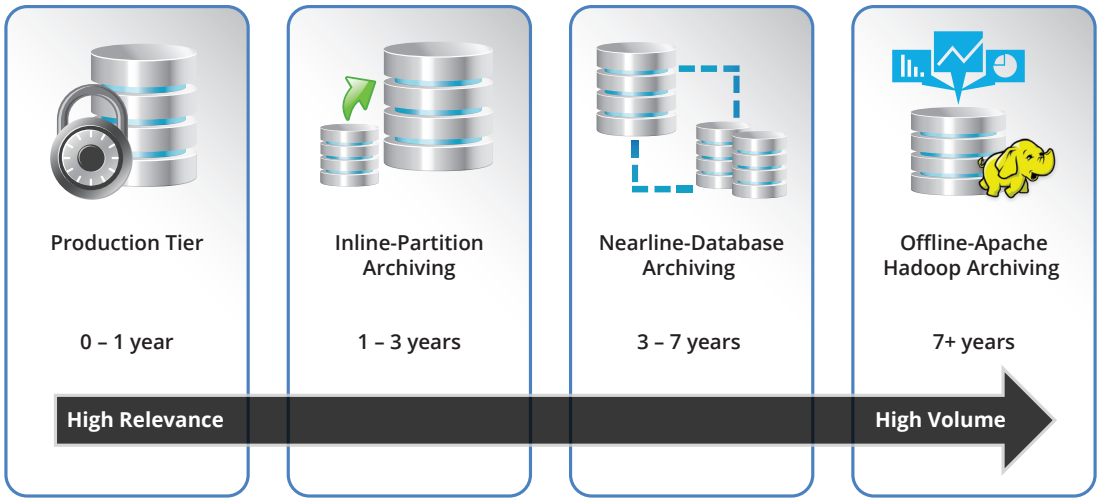

Source: Solix

Effective use of social data is vital in delivering a high level of customer satisfaction.

A first class passenger on an international flight tweeted about the bad food on the plane. This tweet was read by a ground crew member who relayed the feedback to the flight crew within minutes. This passenger, a high value customer with a lot of social media "followers," was pampered for the rest of the flight, prompting a glowing review of the airline to his followers.

Proactive tracking and management of social data was critical to rectifying the concerns of a high-value social influencer and improving the Net-Promoter-Score.

FIGURE 2 Enterprise Application Tiering

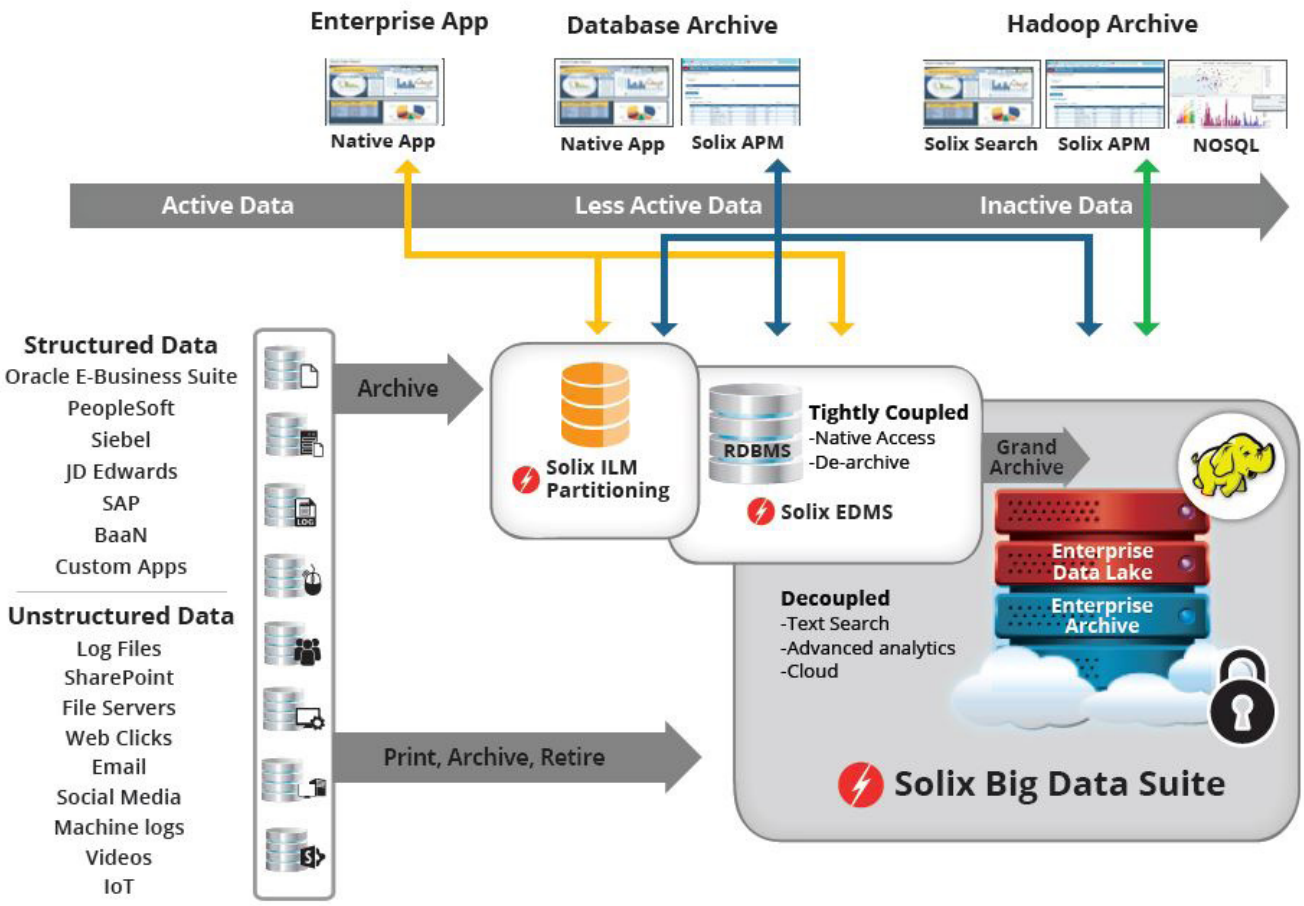

Source: Solix

Solix is the only vendor today with a solution that provides comprehensive ILM for all enterprise data — structured and unstructured.

Source: Solix

Best Practices for Application Archiving

The tremendous growth in volumes of data — both traditional structured data and new data types, such as Internet-of-Things (IoT) — and the advent of in-memory database technologies like SAP's HANA and NAND flash storage, which are faster but more expensive, has made data archiving mandatory. Companies simply cannot afford to operate as they once did, allowing years of data, much of it seldom used, to accumulate in single tier databases. The old data clogs systems, hurting performance, and, when that database is running on flash or in-memory, it also becomes prohibitively expensive.

For too long, organizations have debated the best way to manage the lifecycle of application data. Organizations want to implement true ILM to ensure governance, data security, and operational efficiency.

While unstructured data archiving is relatively simple as it is primarily based on age, structured data archiving is complex requiring that multiple criteria be factored into the process.

The best way to improve the management of enterprise data is to create tiers of data based on value. Our recommended ILM best practice is to leverage four processing tiers integrated with Apache Hadoop:

Production Tier: 0 – 1 year

Highest performance infrastructure is reserved for high value, active data. Large flash arrays manage OLTP processing loads in-memory for maximum performance.

Partition tier: 1 – 3 years

In-line ILM partitions (still running on tier one infrastructure) allow a table or index to be subdivided into ranges based on parameters such as the age of the data. Older, less valuable data may be placed in partitions to exclude them from causing processing overhead. Each partition may be assigned its own storage characteristics.

Database archive tier: 3 – 7 years

Data which is moved and purged from the source database is called an archive. A tightly coupled archive retains native access to the application as well as the ability to de-archive back to into the source production database if necessary.

Apache Hadoop tier 7+ years

Apache Hadoop is the ideal platform for a grand archive because it offers the lowest cost solution for bulk data storage. Hadoop provides a point-in-time snapshot of a business record. Because the data represents a complete business object decoupled from the application, data no longer must be upgraded in synch with the application. Big data analytics tools — text search as well as traditional structured query tools — provide enhanced access to the data.

Source: Solix

Application Archiving at a Glance

	Partition Archive	Database Archive	Hadoop Archive
Data Age	***1 – 3 years***	***3 – 7 years***	***7+ years***
Software Cost	High	Medium	Low
Hardware Cost	Medium	Medium	Low
Impact of Upgrades	High	High	None
Impact of Patches	High	High	None
Scalability *(Volume)*	Low	Low	Very High
Archive Structured & Unstructured Data *(Variety)*	✗	✗	✓
IoT & Streaming data Capture *(Velocity)*	✗	✗	✓
Designed for Advanced Analytics	✗	✗	✓
Support for Full-Text / Content Search	✗	✗	✓

Source: Solix

Why Apache Hadoop

Apache Hadoop is a free, open source computing framework designed to operate powerful, low-cost infrastructure at a lesser tier while still delivering massive scalability and performance.

Using the MapReduce programming model to process large data sets across distributed compute nodes in parallel, Hadoop delivers highly scalable workload performance and very low-cost, bulk data storage.

All this means that Hadoop offers dramatic cost savings over traditional tier one infrastructure.

Consider the following comparison: According to Monash Research, the cost of tier one database infrastructure is more than $60,000 per TB. At the same time, 1TB of S3 bucket storage at Amazon Web Services is $30 per month according to their recent price list.

Conclusion: Hadoop is 55.5X cheaper than tier one infrastructure.

Recent Gartner research states that by 2017 enterprise archiving will represent 25% of the information governance efforts in enterprises. By 2016, 75% of enterprise archiving solutions will incorporate support for big data analytics.

The cost to store 1TB of data in Hadoop...

... is 55.5 X cheaper than tier one infrastructure.

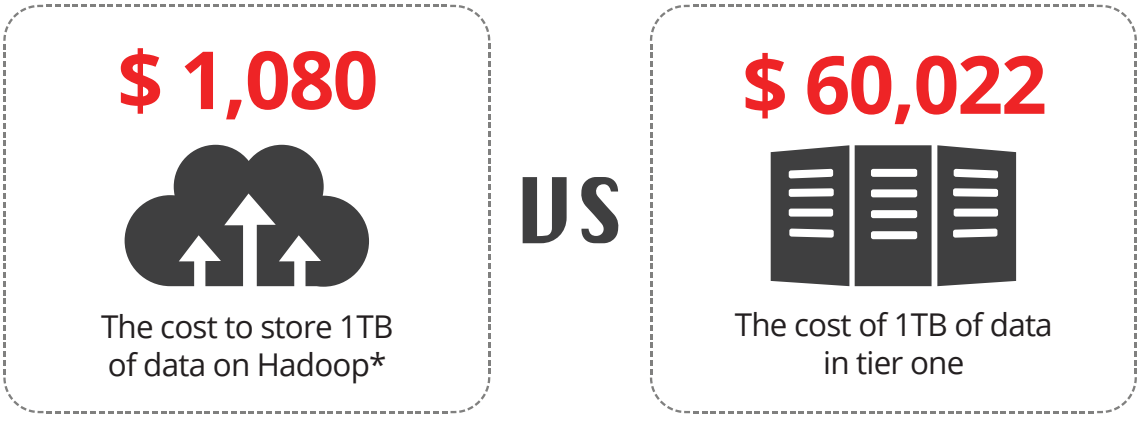

Source: Solix

Source: Solix

Solix Big Data Suite

The Solix Big Data suite provides the framework for an ILM continuum that ensures CIO's don't have to choose between application performance, operational efficiency, and cost.

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."

The Solix Big Data Suite provides the first true ILM continuum that addresses the complexity of governance in the Big Data world while ensuring governance for core enterprise applications is not sacrificed.

The Solix Big Data Suite's ILM framework manages the data within HDFS and HBASE. The Solix ILM framework also provides an integrated retention-management and legal-hold capability for data within Apache Hadoop.

Structured and unstructured data from other data sources are migrated into HDFS/HBASE with full data-validation and audit reports. These reports provide the necessary defensibility and chain of custody for compliance and data governance.

This extensive ILM framework allows the 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.

The suite is highly scalable, with an extensible connector framework to ingest all the enterprise data. The integrated suite allows seamless archiving, retirement, and flexible extract transform load (ETL) capabilities to improve the speed of deployment, decrease the cost, and optimize infrastructure. Solix also 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 creates unified and cost-effective 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 is first moved and then purged from its source location according to ILM policies to ensure governance, risk, and compliance objectives are met.

- The Solix Enterprise Data Lake reduces the complexity and processing burden to stage enterprise data warehouse (EDW) and analytics applications and provides highly efficient, low-cost, 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.

- The Solix App Store offers pre-integrated analytics tools for data within Enterprise Archiving and the Enterprise Data Lake.

FIGURE 1 Solix Big Data Suite

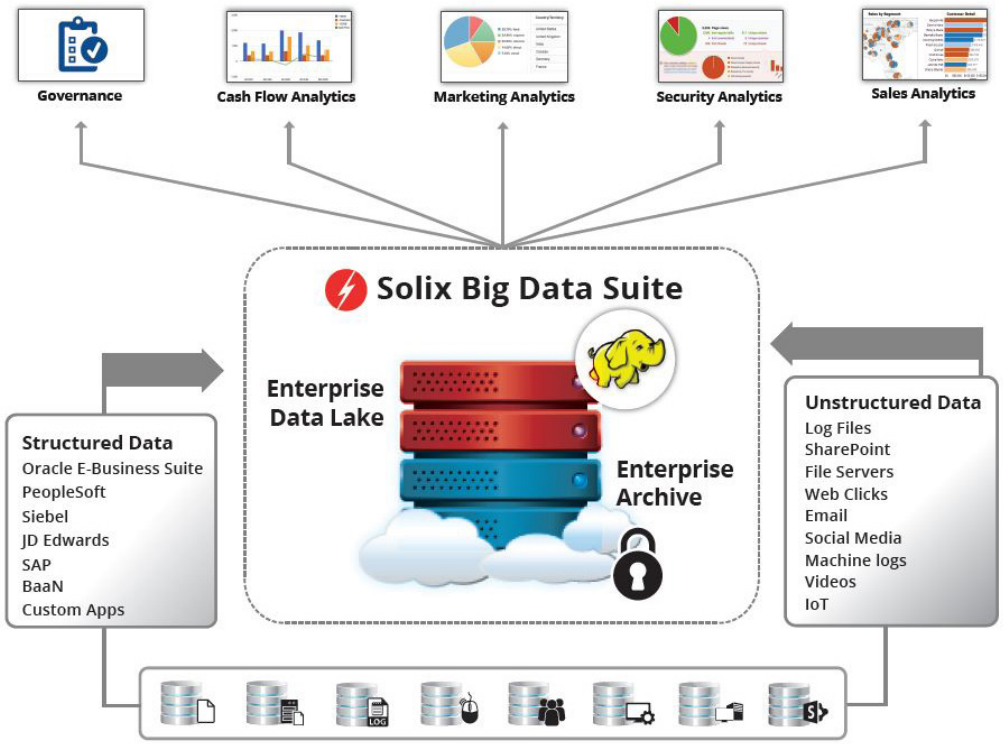

Source: Solix

Source: Solix

Conclusion

The landscape of Enterprise data is changing with the advent of Enterprise Social Data, IOT, Logs, Clicks. The reason this is called big data is because this exceeds the processing capacity of conventional database systems. The data is too big, moves too fast, or doesn't fit the strictures of your current database architectures. To gain value from this data, you need new infrastructure to manage it, and that is Apache Hadoop.

Big data technologies are being marketed to CIOs as a platform for BI and analytics. However, that is only part of the Big Data potential. With Solix Big Data Suite, CIOs can harness Apache Hadoop by using it for application archiving in addition to BI and analytics.

Our advice to CIOs is to explore enterprise archiving on Apache Hadoop as the first step. This introduces big data technology to the enterprise, delivers immediate ROI, and can be leveraged to expand into big data analytics.

The Benefits of
Enterprise Archiving with Apache Hadoop

- Improve application performance
- Allows faster backups and minimizes downtime
- Eliminates infrastructure, maintenance & support costs
- Reduces operational complexity

Source: Solix

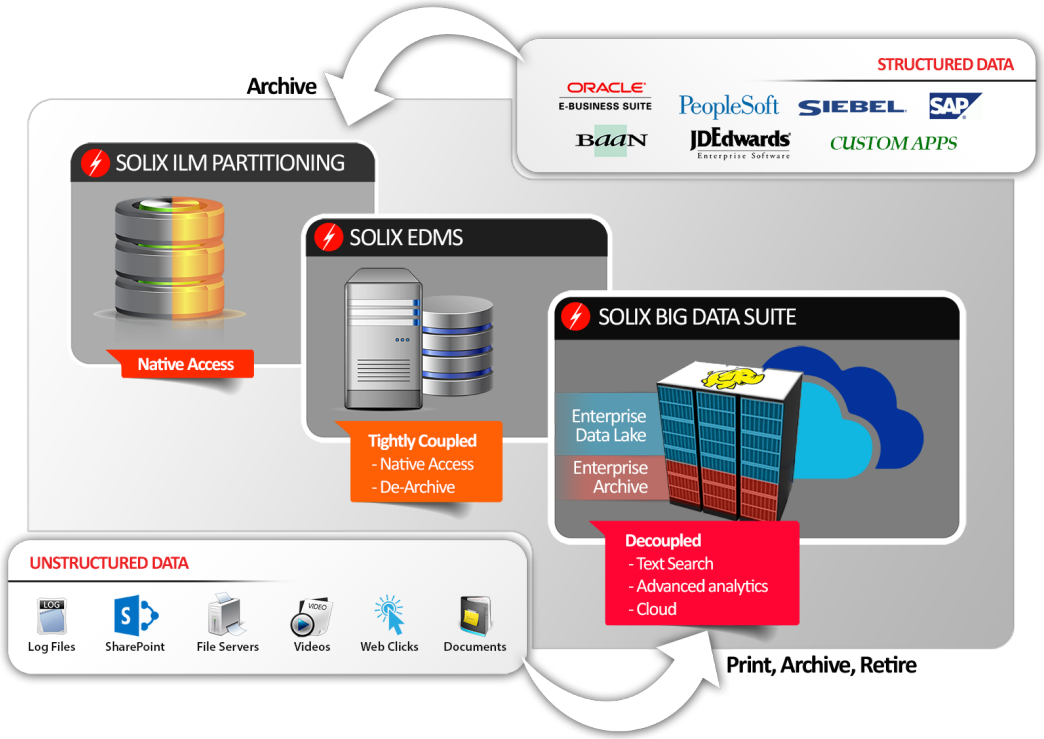

Source: Solix

Structured data archiving technologies help IT leaders retire legacy applications, reduce capital and operating expenses, and meet governance and compliance requirements.

Source: Gartner, Inc. *2015 MQ for Structured Data Archiving and Application Retirement*

Infographic: Solving the Data Growth Crisis with Apache Hadoop

SOLVING THE DATA GROWTH CRISIS WITH HADOOP ENTERPRISE ARCHIVING & DATA LAKE

THE CRISIS

Did you know?

Data growth is the #1 infrastructure challenge for data centers.

Why?

The world is experiencing so much data growth, that by 2020, the amount of generated data is expected to be **40 Zettabytes!**

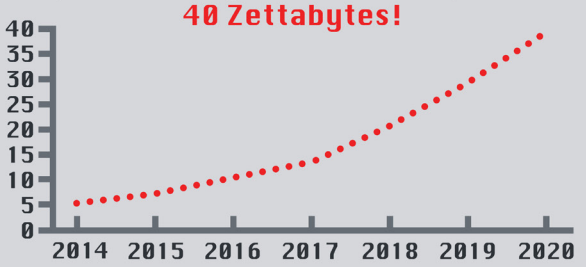

... That's roughly equivalent to **8.5 trillion DVDs.** Almost enough to DVDs to reach Saturn!

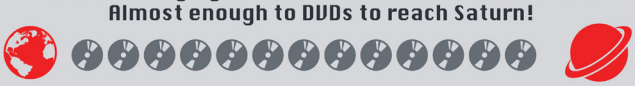

Expansion

Big data is expanding on 3 fronts at an increasing rate.

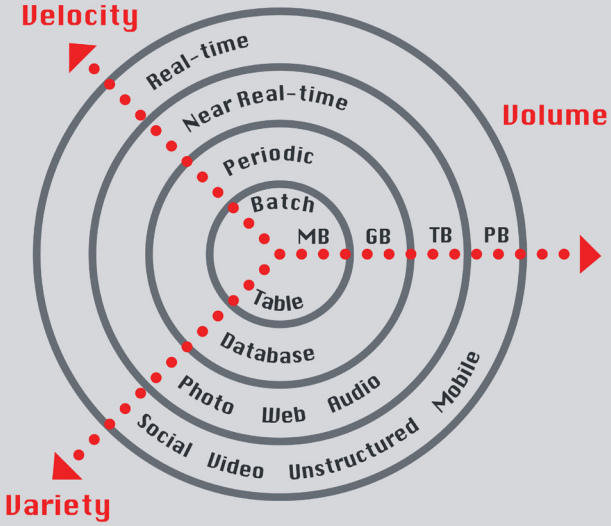

But...

80%

Up to 80% of online data is inactive.

THE SOLUTION

Enterprise Archiving

How does it work?

1. Data is generated and stored on an active server
2. Data stays on the server until no longer needed
3. Data is archived to the cloud with the help of Hadoop
4. When needed, archived data can be easily accessed using a simple search query.

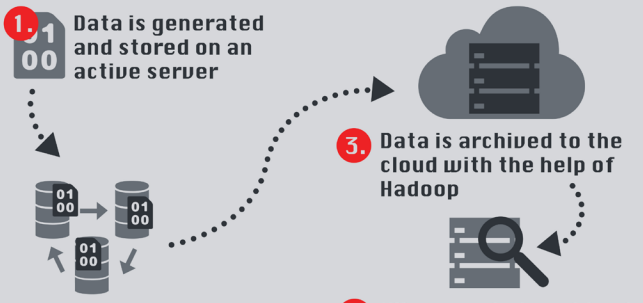

The benefits

Enterprise archiving...

- Improves application performance
- Allows faster backups and minimizes downtime
- Eliminates infrastructure, maintenance & support costs
- Reduces operational complexity

The cost of archiving on Hadoop

$1,080 vs **$60,022**

The average cost to store 1TB of data on Hadoop*

The cost of 1TB of data in a production tier

* Based on $30/month Amazon S3 Bucket pricing in December 2014, multiplied by three years – the average lifespan of a production tier.

The future of enterprise archiving

By 2016, 75% of enterprise archiving solutions will incorporate support for big data analytics.

By 2017, enterprise archiving will represent 25% of the information governance efforts in enterprises.

SOLIX
Empowering Data Management

This Infographic was created by Solix Technologies, Inc., a leading provider of Enterprise Data Management (EDM) solutions. We help companies achieve compliance and reduce data storage costs. For more information, visit www.solix.com

Source: Solix

From the Gartner Files:

Magic Quadrant for Structured Data Archiving and Application Retirement

Structured data archiving technologies help IT leaders retire legacy applications, reduce capital and operating expenses, and meet governance and compliance requirements. We evaluate vendors offering products and services that provide archiving for databases and data from enterprise applications.

Strategic Planning Assumptions

By 2017, archiving in support of big data analytics will surpass archiving for compliance as the primary use case for structured data archiving.

By 2016, 75% of structured data archiving applications will incorporate support for big data analytics.

Market Definition/Description

Structured data archiving is 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:

- 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, improved recoverability, lower risk of regulatory compliance violations, 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 (for example, 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 — 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.

FIGURE 1 Magic Quadrant for Structured Data Archiving and Application Retirement

Source: Gartner (June 2015)

Vendor Strengths and Cautions

Actifio

Actifio provides Actifio Copy Data Services (CDS) and Actifio Sky for the copy data management and retention of structured and unstructured data. CDS is a physical appliance, and Sky is a virtual one. Both work through copy data management, using a single "golden image" (copy) of production data, capturing it at block level and placing it in a separate storage system, from which it can generate multiple virtual copies for different use cases including archiving. The golden image is frequently updated on an incremental forever basis via change block tracking and — together with compression and global deduplication image size for long-term retention — is typically much smaller than the production environment. Currently, Actifio does not natively provide granular search capabilities, such as Boolean or faceted search, and cannot purge data selectively from the production source. CDS and Sky offer data masking through third-party software. Pricing is based on ingested data capacity.

Actifio is best-used for database archiving, when retention of older Microsoft SQL and Oracle environments is required, often in conjunction with test/development use cases. Looking ahead, Actifio plans to offer better search capabilities and support for larger-scale network-attached storage (NAS) environments.

Strengths

- Actifio delivers outstanding customer support.
- Actifio's technology can be applied to a number of different use cases in addition to archiving, including data protection and test/development.
- Actifio has flexible support for the cloud, with both public cloud and SaaS offerings through partners.

Cautions

- Actifio lacks fundamental capabilities like legal hold or search.
- The vendor is largely unproven in structured data archiving or application retirement, with a small number of customers in these areas.
- Actifio customers had the largest project costs of vendors surveyed in this year's Magic Quadrant, but the product's usage is rarely limited to structured data archiving; and the costs typically encompass other use cases, such as test/development and/or data protection.

Data Migration

Data Migration sells its JiVS product primarily for SAP archiving and retirement, but also, increasingly, for other business systems such as JD Edwards and Oracle E-Business Suite. JiVS is a platform approach, with support for capabilities like data masking and retention management through configuration. Application retirement is a key use case for JiVS, and the vendor is experienced with legacy systems like Baan and IBM Domino. Cloud is an area of focus for Data Migration, and is available with Data Migration providing its own cloud for midsize enterprises and value-added resellers (VARs), like T-Systems, offering branded cloud solutions for multinational enterprises. Pricing is flexible, based on three models: per application plus database volume; per defined scope of applications to retire; or per terabyte for retirement programs managed with JiVS. Data Migration is positioned for enterprises looking to retire multiple applications including SAP. Moving forward, the vendor looks to enhance its capabilities around preconfigured reports that will assist in reducing deployment time in JD Edwards, Oracle E-Business Suite, Baan and other legacy application environments.

Strengths

- Data Migration has improved its implementation and deployment times in large global enterprises.
- Data Migration supports a variety of application environments.
- Data Migration is experienced at delivering cloud deployment models for SAP and other application retirement use cases.

Cautions

- Data Migration is primarily focused on retirement scenarios and less on active archiving; for example, there isn't currently support for Hana.
- Data Migration lacks support for Hadoop.
- Support for North America is currently indirect via a partner.

DataVard

DataVard offers OutBoard for SAP archiving. The vendor has been steadily growing its archiving installed base as an outgrowth of its SAP performance analysis capabilities, its primary area of focus. OutBoard includes the Erna module, which operates as a kind of housekeeper, deleting unneeded information and improving system performance. DataVard is also a co-developer of SAP NetWeaver Information Lifecycle Management (ILM) with a focus on database connectivity, and improving document and digital imaging processing. In addition to support for SAP Business Warehouse (BW) and ERP archiving, DataVard has a growing number of customers for SAP Hana, and offers a wide range of storage options including Sybase IQ, Apache Hadoop, any certified SAP database and file, and Amazon cloud-based storage.

DataVard is best-used in BW archiving scenarios and production SAP environments, where containing costs but making data readily available (such as in Hadoop-based data lakes) is required. Moving forward, DataVard looks to expand capabilities in areas such as application retirement, and grow its global presence.

Strengths

- DataVard provides a more open and modern approach to SAP with extensive infrastructure support.
- OutBoard is priced favorably in a market where SAP archiving tends to be prohibitively expensive.
- DataVard provides related services in SAP management, and has strong SAP domain expertise.

Cautions

- DataVard has no demonstrated SAP application retirement customers, despite being a co-development partner with SAP on the SAP NetWeaver ILM product.
- DataVard has a small number of archiving customers.
- The vendor is only SAP-focused, and not suitable for other application environments

DCSoftware

DCSoftware's Arctools focuses primarily on JD Edwards archiving through the availability of dozens of predefined purge modules for JD Edwards. Arctools supports AS/400 (aka IBM i), Oracle and SQL Server databases; however, the AS/400 product has a different codebase than the Oracle and SQL Server products. Its archiving method is straightforward, consisting of copying or moving data from a table in the live system and inserting into the same table in the archive. Arctools provides the ability to run reports from native applications for both the live system data and the archive data with little or no modification. Usually sold as an on-premises solution, Arctools does have a subscription model that can be purchased through a customer's service provider of choice.

Arctools is a solid choice for JD Edwards archiving, and has several hundred customers archiving from that platform. Arctools has recently introduced a product for Oracle E-Business Suite, but it is still in beta as of the publishing of this report.

Strengths

- Arctools has deep knowledge of JD Edwards' systems.
- The product is easy to use, even with limited knowledge of the system or the archive.
- Data integrity validation options before, during and after moving data to the archive are robust.

Cautions

- Arctools does not have granular retention management capabilities.
- Arctools does not support data masking.
- Arctools is for active archiving use cases only, not for application retirement.

Delphix

Delivered as a software appliance, Delphix's primary use case is focused on testing and development through providing virtualized images of the database. Delphix Live Archive provides application archiving by creating a virtualized image of the entire stack that can be mounted at a future date when access is required. Much of Delphix's archiving business is driven through application migration to the cloud, as creating an archive image ensures data integrity if needed during the process and allows for retention of compliance data that does not need to be migrated to the cloud application. In 2014, Delphix added vFiles, which can archive files and also archive flat files created from an archived virtual machine (VM) image.

In the case of archiving, Delphix is well-suited for shorter-term archive scenarios where licenses for the original database and application are readily available. In the future, Delphix will consolidate multiple data warehouse schemas into a single dataset that can be queried as a single tablespace. Expect more cloud options (namely with CSC) in 2015, in addition to existing IBM, AWS and VMware cloud support.

Strengths

- Delphix's TimeFlow feature allows the archived VM to be rolled back to any point and time.
- Fully integrated data masking allows archive sets to be used for analysis and test/development.
- Delphix instances can be spun up in a matter of minutes with compression reaching as much as 20-to-1 and even 40-to-1.

Cautions

- Delphix requires the database and application licenses to be available once the virtualized image is mounted in order to access or report on the data.
- Delphix is not suited as a long-term archive. If used as a long-term archive, the Delphix VM instance needs to be converted as a flat file for later access.
- Delphix operates at the block level versus the database schema level, therefore, archive can only be per database and not at more granular level, such as at a table level.

EMC

EMC InfoArchive is EMC's solution for structured data archiving and application retirement, which it offers stand-alone or with a number of complementary products. The product can be used for database table and data, file and collaborative application (for example, IBM Notes, SharePoint or EMC Documentum) archiving and decommissioning. Various content and platform connectors are available directly from EMC or via partners. Data is stored in a unified repository in an open format using XML, and all content stored can be accessed via granular full text searches, Java Database Connectivity (JDBC)-supported business applications or, in some cases, native applications via partners. The latest version of InfoArchive is tightly integrated with EMC's Isilon storage platform (including storage-level retention using Isilon SmartLock and support for Hadoop Distributed File System [HDFS]), as well as with its EMC ViPR and Elastic Cloud Storage (ECS) platforms (ensuring that storage provisioned by these applications will be supported by InfoArchive). In addition to on-premises deployments, InfoArchive can be run as a service using EMC's hosted Managed Services OnDemand offering. EMC also offers Documentum Archive Services for SAP with support for SAP ILM.

EMC has demonstrated an ability to support use cases where unstructured/structured requirements converge. Buyers should consider EMC at a more granular level if the organization has a need for a unified repository supporting multiple data types from multiple applications, and has strong retention management requirements.

Strengths

- Tight integration with EMC Centera, Isilon, ViPR and ECS gives EMC customers "one-stop shopping" for archival and decommissioning solutions offering strong compliance capabilities.
- InfoArchive delivers ease of access to a unified repository supporting multiple content types (including structured and unstructured data).
- Some vertically focused solutions are available (for example, EMC Clinical Archiving), and this, combined with EMC's strong professional services, contribute to high customer satisfaction in these areas.

Cautions

- InfoArchive's customer base remains relatively small, and some customers have faced lengthy deployment cycles.
- Archive Services for SAP is part of EMC Documentum and is not tightly integrated with InfoArchive.
- A reliance on partners in some cases for connector technology means that sales and deployment can be complex.

HP

HP offers HP Structured Data Manager for database archiving and application retirement. HP Structured Data Manager supports a broader variety of applications like Oracle E-Business systems and PeopleSoft, and underlying databases, including Microsoft SQL Server and IBM DB2. 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. Data masking is provided natively. On-premises, private, hybrid and public cloud (such as HP Cloud Services) implementations are supported. Apache Hadoop (HDFS) is supported as an archive target. The product is integrated with HP Records Manager for records management. SAP archiving is supported via SAP ArchiveLink. HP is continuing to make significant investments to regain market share in database archiving and application retirement.

HP Structured Data Manager is best-suited for enterprises seeking database archiving or application retirement in Oracle environments. HP has recently partnered with Attunity to extract SAP data in an open format and incorporate it with any other content type managed by HP Structured Data Manager.

Strengths

- Pricing is a positive for HP, with customers charged per source database instance (regardless of data volumes). This enables HP to scale down and enable prospects that either don't have a large data footprint or wish to start small.
- HP Structured Data Manager has good retention management capabilities and the ability to store archived data in HP Records Manager.
- HP Structured Data Manager is a feature-rich product with support for big data analytics through integration with HP Vertica, as well as data masking with multiple repository options.

Cautions

- Structured data archiving knowledge and skill sets are relatively narrow within HP.
- HP's customer base comprises primarily Oracle E-Business systems users. Buyers should test HP's domain expertise in other application environments.
- HP Structured Data Manager does not provide automated data life cycle management/tiering for storage.

IBM

IBM InfoSphere Optim has a significant lead in market share for database and application archiving/retirement. IBM utilizes highly compressed, immutable binary files (Optim File Format, CSV or XML) for archiving. Optim provides a significant assembly of application sources such as Oracle E-Business Suite, PeopleSoft, Siebel, JD Edwards, Baan and SAP, along with IBM DB2, Oracle and Microsoft SQL Server at the database level. In the spirit of supporting internal systems, Optim also supports direct archiving from PureData and ships with IBM HDFS. IBM has broad support for retention and compliance through policy management, legal hold and discovery.

IBM InfoSphere Optim is on most RFIs for medium to large organizational archiving and application retirement initiatives. Expect IBM to continue to refine the user interface in 2015. In addition, product initiative and marketing messages will be aimed more to the business users going forward. IBM is a market leader for Hadoop as an archiving platform — assume more offerings in this area going forward.

Strengths

- Optim is a proven archiving solution for large-scale environments.
- As part of IBM's Information Lifecycle Management (ILM) strategy, Optim offers strong business rule capabilities and policy management.
- Optim offers tight integration with the rest of the IBM Information Management Systems, including InfoSphere products such as Guardium, FileNet and Atlas.

Cautions

- Although IBM continually works on improvement of user interface and report viewing, many users find them difficult to work with and cite them as being antiquated.
- Pricing and licensing continue to be hurdles for customers, including pricing based on ingestion instead of final storage.
- IBM is focused on the technologies involving ingestion of data into the archive, and less on the business value of retrieving data and providing analytics.

Informatica

Informatica offers Informatica Data Archive, which supports a wide array of applications and underlying databases resident on distributed systems or mainframe platforms. The product also supports archiving from Apache Hadoop, IBM PureData System for Analytics and Teradata, as well as SaaS applications, such as Salesforce. Informatica includes a capability it calls Smart Partitioning, which allows administrators to relocate related inactive records across tables into a single database table space, and then effectively exclude "smart partitions" from database production operations, including queries, reports and nonproduction copies. This "archive in place" ensures that the database remains intact, retains data for compliance and improves performance.

Informatica Data Archive offers both native dynamic data masking and compliance management, which includes retention policies, legal hold and data disposal. In 2014, Informatica released its Application Retirement for Healthcare Solution, which includes a patient-centric portal for reporting, Audit Logging for HIPAA compliance and Accounts Receivable Burndown to manage legacy receivables as part of its effort to focus on that vertical. Informatica Data Archive also supports SAP archiving. Informatica has revamped its pricing, with the introduction of a Secure Edition aimed at application retirement that includes Data Vault, a repository component providing extreme compression, role-based security and open access. Other product packages are available, including one aimed at live archiving of production environments, and another focused on application retirement for vertical applications as well as historical analytics use cases. The product can also be deployed as a customer-managed service via public cloud storage, such as Amazon Simple Storage Service (S3). On 7 April 2015, Informatica announced plans to be acquired by Permira Funds and Canada Pension Plan Investment Board. The deal is pending shareholder and regulatory approvals, scheduled to close in 2Q15 or 3Q15. Informatica has told customers and partners that its commitment to delivering on its roadmap remains unchanged.

Strengths

- Informatica's deep data integration heritage allows it to leverage highly optimized connectivity for archiving use cases.
- Informatica has a strong vision with a targeted and focused roadmap, and has delivered on its commitments in the past. The company addresses a large number of use cases including retirement, compliance, performance optimization and analytics.
- Performance optimization is a strong use case for Informatica. Smart Partitioning is a popular choice for active database archiving.

Cautions

- Informatica's customers have cited ease of use of Informatica Data Archive as an area in need of improvement.
- With the exception of cloud service pricing, Informatica is priced for clients with a sizable data footprint or large number of applications. Its offerings do not scale down in price.
- Pending shareholder and regulatory approvals, Informatica will be acquired by private equity this year, so funding levels for Informatica Data Archive are unclear.

OpenText

OpenText offers OpenText Archiving for SAP Solutions, including OpenText Data Archiving for SAP Solutions and OpenText Document Access for SAP Solutions, for archiving SAP data and content. 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 relational database management system (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. Active archiving is a strength for OpenText, as documents like scanned images and unstructured content can be combined with archived data in support of business processes and workflow. Retention management includes support for records management and legal hold. Content can be classified during archiving or using OpenText Auto-Classification capabilities. 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.

Strengths

- OpenText enjoys a strong relationship and has the largest joint customer base with SAP, including numerous large, referenceable customers.
- OpenText provides strong retention management capabilities.
- OpenText has wide global sales coverage and support with well-balanced sales penetration across the Americas, EMEA and Asia.

Cautions

- OpenText does not support Hadoop.
- Support for use cases like database or Oracle application archiving is limited.
- Data masking capabilities are limited to unstructured content.

PBS Software

PBS is a small software company specializing in SAP archiving. The company offers PBS ContentLink for SAP content-based archiving, PBS CBW NLS for near-line storage archiving, and PBS Nearline Analytics Infrastructure (NAI) for archiving and analysis of large quantities of information seen as complementary to SAP data. The solutions support SAP BW and 30 different SAP ERP and SAP IS-U modules, as well as vertical solutions. The NLS and NAI solutions support SAP IQ, with support for Actian Vector and IBM DB2 with BLU Acceleration planned for 2015. PBS ContentLink supports the archiving of structured and unstructured data via ArchiveLink/WebDAV, and the product is integrated with SAP ILM. Typical use cases include archiving to manage database growth, compliance, and "nearline" analytics and reporting. Users can access archived and nearline data with standard SAP transactions, queries and report tools. Any file formats are supported within the file-based archive. Numerous back-end storage devices providing write once, read many (WORM) capabilities are supported, as are Amazon, Google and OpenStack cloud APIs. While the product is available worldwide, the company's strongest presence is in Europe, with products and support available from multiple value-added resellers.

PBS Software is best-suited for enterprises looking to create an active archiving environment such as Hana for SAP that incorporates both data and unstructured documents.

Strengths

- The solutions are feature-rich, and the company's long focus on SAP has resulted in a strong set of products for managing aging SAP data in nearline and archive environments.
- Customers cite technical support as excellent, via both partners and directly from PBS.
- Ease of use is rated highly, as the products are tightly integrated into the SAP environment ensuring transparent access from both the online database, as well as the archive.

Cautions

- PBS is releasing support for additional nearline databases, as well as a new product (EDS for NAI) — prospects and customers should ensure they are comfortable with levels of support and feature/product maturity.
- Indexing large data volumes in PBS is time-consuming for SAP Archive Development Kit (ADK)-based indexing.
- PBS relies on partners for implementation and consulting services, most of whom are small but specialized and focused on SAP archiving and ILM.

Solix Technologies

Solix Enterprise Data Management Suite (EDMS) supports many applications and databases. It is mostly used in Oracle E-Business Suite environments, but also supports Siebel, JD Edwards and PeopleSoft. EDMS stores data in a flat file format, in relational databases and Hadoop. The vendor has solidified support for unstructured data in 2015. Solix EDMS is available via a broad variety of platforms including on-premises, appliance, private/hybrid cloud and SaaS. Solix has relationships with other vendors/products, such as EMC InfoArchive, Kronos' cloud/SaaS platform, and integrates with Sybase IQ and Teradata (Rainstor) for columnar compression/deduplication. EDMS supports Apache, Cloudera and Hortonworks Hadoop distributions as repositories for archived data. Solix introduced much-improved data retention policy capabilities and legal hold in 2014.

EDMS is most often sold and utilized in small to midsize environments. On the roadmap are enhancements for big data archiving and retirement, partition-based archiving, aggregating of various archived objects into an object workbench, concept search and predictive clustering, and an upgraded rollback feature.

Strengths

- Search and reporting has greatly improved in 2015, including support for Open Database Connectivity (ODBC), Java Database Connectivity (JDBC), and Java and Web services for reporting. In addition, support for Tableau and Datameer is solid for big data analytics.
- Solix is consistently noted as being cost-effective for archiving projects. ROI is frequently achieved in well less than a year.
- Solix is at the forefront of offering archiving capabilities for and supporting use cases involving big data, with wide support for various Hadoop distributions.

Cautions

- Solix is still small enough that it can often be spread thin from a personnel standpoint
- SAP archiving is a work in progress with a major release slated for year-end 2015.
- As a platform, Solix requires ample out-of-the-box configuration during the setup phase.

Teradata (RainStor)

Teradata acquired RainStor in December 2014, expanding beyond its Integrated Big Data Platform Appliance 1xxx class systems for structured data archiving into heterogeneous storage solutions (shared storage, WORM for compliance and Hadoop-based solutions). Teradata (RainStor) provides a much needed "back end" online/queryable archival option for traditional RDBMSs, including Oracle, Sybase IQ, SQL Server, IBM DB2, IBM PureData and, of course, Teradata systems, which, to this point, had mainly relied on archival to tape. Compression of 20X to 40X is not uncommon with RainStor. It supports ANSI SQL-92 and elements of SQL:2003 for running queries. When combined with WORM storage, it provides immutability of the data. Record-level tagging, legal hold, logical and physical purge, and source schema change pattern recognition are all provided with RainStor.

Teradata RainStor is particularly well-suited for archival of large volumes of data and objects, and for RDBMSs including data warehouses. Going forward, Teradata looks to integrate much of RainStor's capabilities into its native environment, while at the same time continuing to build upon RainStor's support of heterogeneous data sources.

Strengths

- RainStor's compression of structured data greatly reduces the storage footprint; data does not have to be decompressed to run SQL queries or MapReduce.
- RainStor supports a wide ecosystem of storage platforms, like EMC Isilon, data warehouses and Hadoop distributions.
- RainStor provides encryption on data at rest and in transit.

Cautions

- RainStor's administration tools are command-line-driven and need to have a graphical user interface that can centralize management.
- Non-Teradata customers who use RainStor should pay attention to the roadmap as it develops to ensure Teradata continues to support their non-Teradata environments.
- The total cost of ownership for Teradata RainStor needs to be carefully evaluated.

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.

Added

Actifio, DataVard and DCSoftware

Dropped

dataglobal, RSD and ZL Technologies

Inclusion and Exclusion Criteria

- 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 roadmap 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 2015)

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) 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 roadmap.

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 2015)

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 five Leaders in this Magic Quadrant: Delphix, 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. There are two Challengers in this Magic Quadrant: OpenText and PBS Software.

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 Teradata (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 Actifio, Data Migration, DataVard and DCSoftware.

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; 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 $263 million — with additional related product revenue at $27 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.

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 cost 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.

The 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. This will put pressure on structured data archiving vendors to innovate further in this area; for example, some vendors are beginning to provide integration and support for technologies like Hive to allow for SQL-like queries into the archived data.

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 a warning, 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.

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. This trend is accelerating, with a steady stream of Gartner clients inquiring about SAP retirement as extended maintenance will be expiring soon for some SAP applications, such SAP CRM, Supply Chain Management (SCM) and Supplier Relationship Management (SRM) 2007. 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 inquiries show that clients are open to alternatives to this approach, 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. A new use case emerging for SAP archiving is the need to extract data from SAP Hana, as organizations look to manage older, less frequently accessed data in lower-cost environments like Hadoop rather than with in-memory databases like Hana. 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:

- **CommVault** provides Simpana, a single-platform approach to backup and archiving, and supports data and document archiving for SAP modules. CommVault has stated its roadmap intentions for supporting additional structured data archiving capabilities in future product releases.
- **Gimmal**, 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. Gimmal 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 ILM helps organizations comply with audit and compliance requirements, and consolidates SAP instances.
- **ZL Technologies** provides ZL Unified Archive, for archiving of all content types (structured and unstructured). The product has historically been deployed for email and instant message archiving, and, over the past few years, ZL's focus has expanded to include files, popular social media platforms, collaborative applications and some RDBMSs such as Oracle, Microsoft SQL Server and IBM DB2.

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.

Source: Gartner RAS Core Research Note G00266173, Garth Landers, Alan Dayley, Sheila Childs, 16 June 2015

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.