G00263059

Market Guide for Data-Centric Audit and Protection

Published: 21 November 2014

Analyst(s): Brian Lowans, Earl Perkins

Organizations that have not developed data-centric security policies to coordinate management processes and security controls across data silos need to act. This market guide helps CISOs identify emerging data-centric audit and protection tools that can apply these policies.

Key Findings

- The exponential growth in data generation and usage is rendering current methods of data security governance obsolete, requiring significant changes in both architecture and solution approaches.
- Organizations lack coordination of data-centric security policies and management across their data silos, resulting in inconsistent data policy implementation and enforcement.
- Data cannot be constrained within storage silos but is constantly transposed by business processes across multiple structured and unstructured silos on-premises or in public clouds.
- Vendors identified as DCAP providers will develop capabilities across adjacent data silos and across big data platforms, either organically or via partnerships.

Recommendations

CISOs:

- Coordinate with key business stakeholders to establish strategic organizational data security governance and coordinated data security policy implementation.
- Identify the data security controls required to mitigate the risks and threats to each sensitive data type and storage silo. These controls must be coordinated through a single data security policy and by each silo's management team.
- Implement a DCAP strategy, and "shortlist" products that can apply the required data security controls to address all silos.

Strategic Planning Assumption

By 2018, data-centric audit and protection strategies will replace disparate siloed data security governance approaches in 25% of large enterprises, up from less than 5% today.

Market Definition

A DCAP product is characterized by the ability to centrally manage data security policies across unstructured, semistructured and structured repositories or silos. The policy will encompass security controls such as the ability to classify sensitive data and control access by centrally managing privileges, activity monitoring and data protection. Data protection techniques such as encryption, tokenization and masking can be used selectively to enhance the segregation of duties against both application users and highly privileged users. The ability to offer auditing and reporting can also support various compliance requirements.

Data classification and discovery are also core requirements for content-aware data loss prevention (DLP) tools (see "Magic Quadrant for Content-Aware Data Loss Prevention"). DLP is focused on protecting data in motion through the prevention of data leakage outside of the organization. However, DCAP is primarily focused on using data classification and discovery to help with the protection and activity monitoring of data at rest and in use within the organization.

DCAP products support several additional capabilities that include:

- Provide a single management console that enables the application of data security policy across multiple data repository formats (referred to here as data silos).
- Classify and discover sensitive data across relational database management systems (RDBMS) or data warehouses, unstructured data file formats, semistructured formats such as SharePoint, semistructured big data platforms such as Hadoop, and cloud-based file stores.
- Set, monitor and control privileges of unique user identities (including highly privileged users such as administrators and developers) with access to the data.
- Monitor user activity with customizable security alerts.
- Create auditable reports of data access and security events with customizable details that can address defined regulations or standard audit process requirements.
- Prevent specific data access by individual users and administrators. This may also be achieved through encryption, tokenization, masking or redaction.

Market Direction

The advent of big data platforms and cloud-based enterprise file-sharing services (EFSS) is driving organizations to review their strategy for data security. Traditional approaches are limited to data silos because the manner in which vendor products address policy is siloed, and thus the organizational data security policies themselves are siloed.[1] For example, the approach to

structured database security governance is frequently different from the approach taken for unstructured or semistructured data in an organization. Transposing data from one silo to another creates an interrelated data processing environment that lacks synchronization of security policy and leads to security chaos because organizations have not developed processes to deal with it. This leaves organizations open to internal malpractice, hacking, data breaches and financial liabilities.

Organizations must develop a comprehensive policy — based on data security governance principles — to apply appropriate security controls across all data silos. This often requires the purchase of more than one DCAP product to match the targeted silos and the development of management structures that coordinate and align data security policy and accountability across those silos (see "Big Data Needs a Data-Centric Security Focus").

The DCAP market is in the earliest stages of development, and only a few of the vendors selected are beginning to address all, or only a few, of the siloes. Therefore, the market is currently characterized with most of the vendors focused on specific data silos such as databases, unstructured file stores and cloud. Current Gartner research refers to four major solution sets: database audit and protection (DAP), data access governance (DAG), cloud access security broker (CASB), and data protection (DP), including encryption, tokenization and data masking. In the near term, all of these vendors will develop capabilities across adjacent silos and across big data platforms, either organically or via partnerships. A number of vendors have also emerged to support cloud-based EFSS, but they lack support for on-premises data silos. Some DP vendors have recently emerged with tools that operate across multiple silos, with capabilities for privilege management, activity monitoring, security alerting and audit.

Vendor Categories

While no individual products fully meet the requirements of a DCAP product, there are currently four main categories of products that in aggregate comprise DCAP capabilities across the data silos:

- **DAP —** These products have developed over several years to cover implementation of data security policy, data classification and discovery, access privilege management, activity monitoring, audit and data protection. Focused on RDBMS and data warehouses, a few vendors are beginning to offer support for Hadoop and unstructured file shares.
- **DAG —** Typically, these products are focused on implementation of data security policy, data classification and discovery, activity monitoring and audit of unstructured data within file shares such as SharePoint and various file stores. These vendors are closely tied to identity and access management vendors. DCAP candidates are also beginning to develop capabilities for Hadoop.
- **CASB —** The ability of organizations to protect data within cloud storage environments such as salesforce.com, ServiceNow, Box and Dropbox has only started to become a reality within the last 18 months through a few recent startups. Data classification and discovery, access controls and activity monitoring are still developing.

- **DP —** These products have developed capabilities to protect data using encryption, tokenization or masking across multiple data silos (RDBMS, data warehouses, unstructured, big data and some cloud-based EFSSs). In addition, some activity monitoring and audit capabilities are being developed. Whereas DAP and DAG products may offer monitoring of access to all data within files or databases, these DP products are typically focused upon sensitive data types only.

Future Market Direction

The DCAP market will be very different by the end of 2018, with repositioned and more comprehensive product offerings. Several vendors will offer full DCAP products, covering all of the discussed data silos, while many of the other vendors will continue their organic development of capabilities in terms of breadth and depth across multiple silos. Compliance requirements and the advent of big data are forcing CISOs to apply their strategy across silos. This is pushing vendors to innovate through cross-siloed product offerings, which in turn are changing the dynamics and attractiveness of separate market segments into one larger market. This market will see new entrants and consolidation through mergers, acquisitions and some failures.

- Most of the vendors will offer centralized management platforms that can directly control data security policies across multiple data silos.
- Competition between vendors will intensify through support for Hadoop and other big data platforms and integration with policy enforcement functionality across multiple platforms.
- DP vendors will, through a necessity to compete, continue to diversify product offerings with deeper activity monitoring and privilege management capabilities. Likewise, more vendors will either directly integrate their own data protection functionality within the management console or will develop functionality where it is lacking.
- Products will be driven to address multinational data residency and compliance issues through the application and management of on-premises data protection and privilege management to control access.
- The cloud will become a new battleground for product differentiation, with potential for acquisitions, mergers and new entrants.

Innovation and product diversification will intensify, especially as DAP and DAG product capabilities converge toward Hadoop. Mergers and acquisitions across these silos will inevitably develop as a means to plug product gaps and as new entrants emerge.

Market Analysis

The four market segments that contribute to DCAP have evolved over vastly different timescales with different security focal points and business drivers. Convergence of product capabilities toward Hadoop, from vendors previously focused on database and unstructured files, has created a surge in interest toward a much larger market opportunity through the combined segments. Convergence of these adjacent markets is driving current organic developments but has also attracted new

entrants from the DP market. Imperva has acquired a cloud-based CASB vendor, which will open a new market direction and will also create exposure to existing CASB vendors that have been developing data protection and activity monitoring functions.

Most vendors have developed a common need to classify and discover data, manage and monitor access, provide auditable reports and provide some form of protection (see Figure 1). However, these capabilities are not created equal, and care should always be taken to address product requirements based on data security governance principles and the controls available through implementation of data security policies of the selected product or products.

Figure 1. Summary of the Core DCAP Capabilities Offered by Vendors in Each Segment

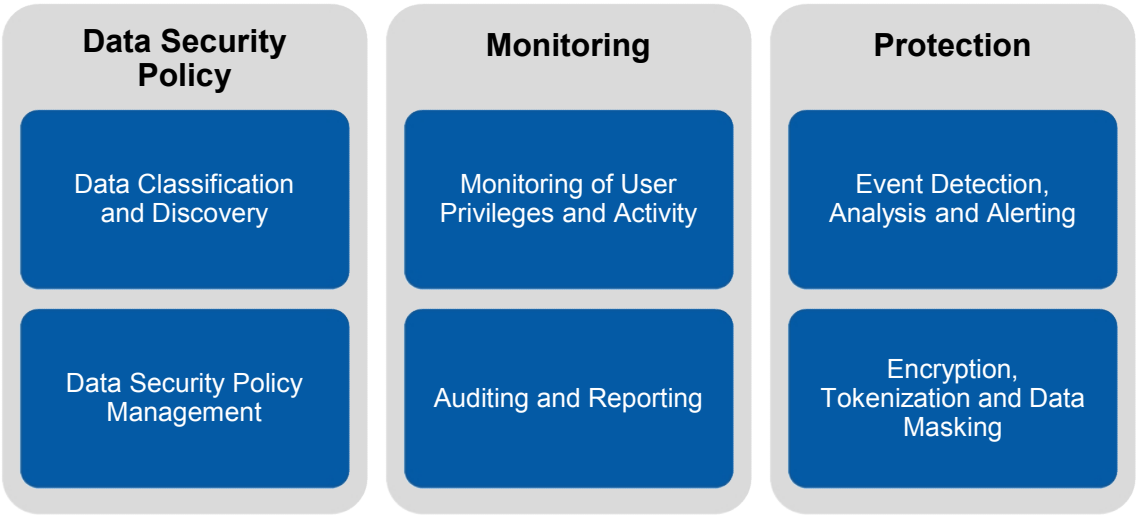

Source: Gartner (November 2014)

A vendor's ability to integrate these capabilities across multiple silos will vary between products and also in comparison with vendors in each market segment. Here is a summary of some key features to investigate:

Data Classification and Discovery **—** Many products come with built-in dictionaries or search algorithms tailored for use with some compliance regimes, but the search capabilities of different products will vary, for example, in terms of speed and false-positive performance. The ability to search within a specific RDBMS, file type, Hadoop or cloud EFSS will vary from vendor to vendor. If you are planning to use with RDBMS silos, note that some products may only search column/table metadata or within fields. Also check if data can be searched within a binary large object (BLOB) or character large object (CLOB) that may be stored within the database.

Data Security Policy Management **—** The ability to offer a single management console that controls policy across each silo is the desirable goal, and this will evolve as vendors encompass more silos. Most products will split this functionality, or separate vendor products will be required. In either case, separate software interfaces or separate management consoles will be required. Coordination of roles and responsibilities against the underpinning data security governance will be

important. The application of policy is typically based on user identities and business roles as authenticated through third-party solutions such as active directory (AD) or LDAP. Membership of groups can help define access to particular data within a silo, or even multiple groups when granting access to multiple silos. The ability to identify individual users at the application level can be a differentiator if applications use connection pooling to provide a more efficient group access account. This can sometimes be enabled by communication with the application through third-party interfaces such as Kerberos, but not all applications provide this capability. Other solutions may use application layer agents to gather identities.

Monitoring User Privileges and Activity — The access rules set out by the data security policy are a crucial guide for monitoring the privileges granted to all users with access to the data. This is important for checking for changes to AD membership or individual privileges to ensure they match requirements associated with business role, data type or geographic location. The ability to detect changes and create alerts for privilege escalation or changes to data is important to detect potential insider abuse or external hacking activities. However, not all products operate at the storage level, and they may not offer the ability to assess the privileges of highly privileged users such as database administrators, system administrators or developers. Monitoring application users and highly privileged users is important for compliance and is a critical analysis capability to detect insider misuse or hacking. Therefore, it may be important for a product to be able to intercept access by various administrators at the server level. Products need to demonstrate continuous operation during peak loading of servers or network communications congestion. Consideration must be given to the network architecture and demands required of products if intensive monitoring is required while infrastructure is highly loaded. This can lead to latency or, in extreme cases, even failure to monitor some activity.

Auditing and Reporting — As the data silo analysis requirements continue to grow, the demands on the reporting capabilities will grow also. Auditors in various regulatory environments will require an ability to produce insights into the activity of users on a historical basis, which can require up to one month of accessible data. Compliance will also require an audit trail of various monitoring capabilities, such as unusual user behaviors, changes to data, policy violations or changes to privileges.

Event Collection Analysis and Reporting — An ability to create security alerts based upon preselected monitoring criteria is critical, and these might encompass different levels of alert that range from policy violations to levels of suspicious behavior. Mechanisms for alerting include console displays and automatic messaging to key security or business staff. Other functionality may be enabled such as automatic blocking of a process or removal of privileges. Extreme responses might include shutting down all access in the event of very large data downloads. Future products may even correlate rules to detect unusual behaviors. Products vary in the ease of use of the management console interfaces to manage and report security alerts, and the granularity of reporting within the different data storage platforms. For example, in relation to databases, there may need to be a trade-off between the number of commands that can be inspected against the ability of software/hardware technology to process and communicate the results for analysis. This can happen if servers or network communications are already heavily loaded and the ability of local monitoring agents to process the large volume of commands is then constrained. Data access can be blocked based upon data content and group membership or privileges.

Data Protection — Some vendors offer separate data protection tools using encryption, tokenization or data masking, while others do not offer any tools and will require the purchase of separate vendor products. In either case, these protection products will not be integrated into a single management console and will require careful coordination with data security policies. The selection of these tools requires careful assessment of the threats and risks that each can offer. For example, implementing transparent database-level encryption can prevent access by system administrators, but DBAs would still have access. Applying dynamic data masking through an agent on the database server, and linked via AD, can be used to prevent access by DBAs. But, since the data is not protected when stored at rest, it may still be accessible by system administrators. Encrypting or tokenizing fields can protect the data elements in use and at rest, but care must be taken that this does not affect the operation of applications.

Representative Vendors

The vendors listed in this Market Guide do not imply an exhaustive list. This section is intended to provide more understanding of the market and its offerings.

The DCAP market is characterized by three primary sets of vendors that have focused on particular silos but may already be demonstrating or planning coverage of adjacent silos that include RDBMS, unstructured file stores (files), semistructured environments such as SharePoint and Hadoop, and cloud-based file stores (EFSS). Some innovative DP vendors are developing a centralized management console approach across multiple silos but lack comprehensive DCAP capabilities. A representative mapping is shown in Figure 2, but note that none of these vendors can currently be assessed as yet offering a complete DCAP solution.

Figure 2. Schematic Representation of the DCAP Market Showing How a Sample of Vendors Operates Across Different Data Silos

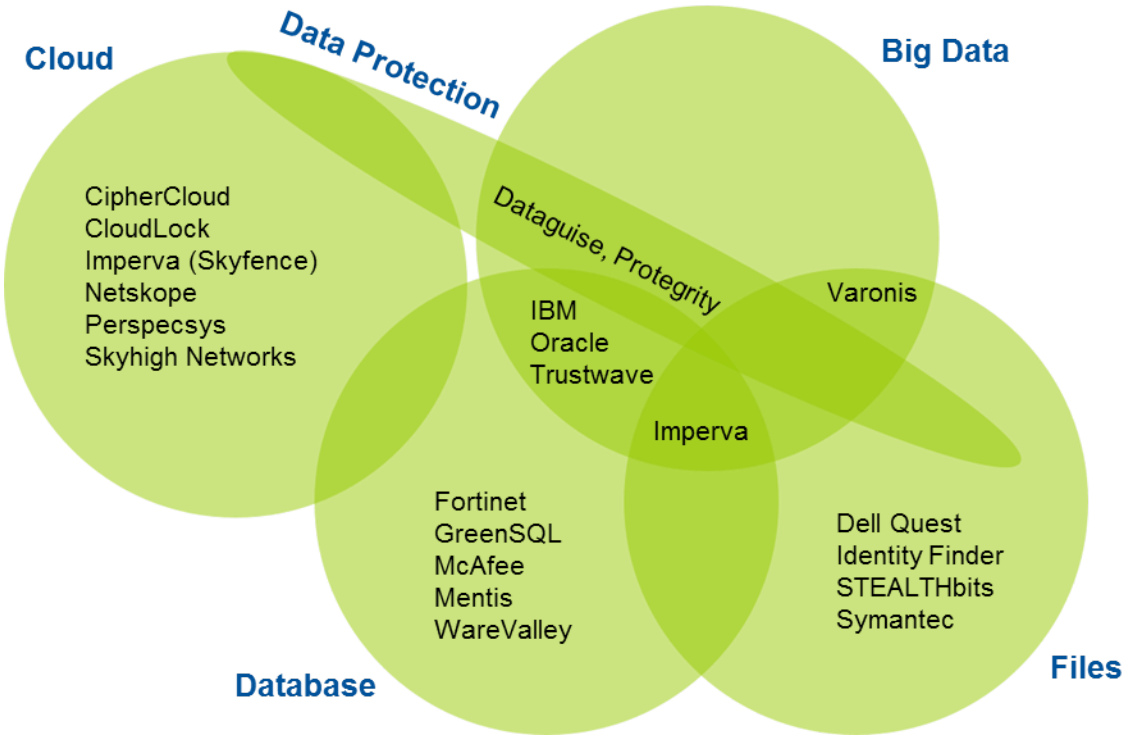

Detection tools may be applicable across multiple silos through a single management console but other functionality is limited.

Source: Gartner (November 2014)

Figure 3 shows a sample list of contenders to be identified as potential DCAP vendors. These vendors are categorized by the main capabilities outlined in this research (see the Market Analysis section), including:

- Capabilities for each of the data silos — RDBMS, file stores, big data and EFSS. Each vendor has a different coverage of these silos through on-premises products.
- Tools for data classification and discovery.
- Management of privileges such as read/write or access to classified data types. Privilege management of individual users may not always be possible at the application level. Some products focus on only classified data types.
- Activity monitoring of users and administrators — this may be limited to monitoring access to certain classified data types.
- Audit and reporting capabilities typically focused around specific compliance requirements.
- Data protection — this may be encryption, tokenization or data masking.

Table 1. Sample List of DCAP Vendors Against Capabilities for Four Data Silos

	Data Silos				DCAP Capabilities				
	Database	Files	Big Data	Cloud EFSS	Data Discovery	Privilege Management	Activity Monitoring	Audit	Data Protection
CipherCloud				Y	Y	Y	Y	Y	Y
CloudLock				Y	Y	Y	Y	Y	Y
Dataguise	Y	Y	Y		Y	Y	Y	Y	Y
Dell Quest		Y			Y	Y	Y	Y	
GreenSQL	Y				Y	Y	Y	Y	Y
Fortinet	Y				Y	Y	Y	Y	
IBM	Y		Y		Y	Y	Y	Y	Y
Identity Finder		Y			Y	Y	Y	Y	Y
Imperva	Y	Y	Y	Y	Y	Y	Y	Y	
Intel Security (McAfee)	Y				Y	Y	Y	Y	
Mentis Software	Y				Y	Y	Y	Y	Y
Netskope				Y	Y	Y	Y	Y	Y
Oracle	Y	Y	Y		Y	Y	Y	Y	Y
Perspecsys				Y	Y	Y	Y	Y	Y
Protegrity	Y	Y	Y	Y	Y	Y	Y	Y	Y

	Data Silos				DCAP Capabilities				
	Database	Files	Big Data	Cloud EFSS	Data Discovery	Privilege Management	Activity Monitoring	Audit	Data Protection
Skyhigh Networks				Y	Y	Y	Y	Y	Y
STEALTHbits		Y			Y	Y	Y	Y	
Symantec		Y			Y	Y	Y	Y	
Trustwave	Y		Y		Y	Y	Y	Y	
Varonis		Y	Y		Y	Y	Y	Y	Y
WareValley	Y				Y	Y	Y	Y	Y
Notes: Imperva achieved access to the EFSS market through the acquisition of Skyfence in February 2014.									

Source: Gartner (November 2014)

Vendors within each market category have developed products with different breadths and depths of capabilities. Organizations, therefore, need to select products that suit their data governance strategy. For example, if there are compliance requirements to provide full auditable assessments of users accessing a database, then DAP vendors will have the most comprehensive RDBMS abilities to inspect all SQL commands.

Market Recommendations

- Establish a data security governance strategy supported by DCAP that is approved by key business stakeholders, which should include business, IT, security, compliance and risk. The main outcomes will be a classification of sensitive data types, compliance requirements, risks and threats. These will be balanced against any necessary trade-offs to business access through, for example, IT systems, staff locations or customer needs.
- Establish and develop a data security policy that sets out the necessary security controls for each data type and storage location. This will require coordination and cooperation with the various security management teams that are deployed for each of the data silos to coordinate the controls that will be implemented by the DCAP products.
- Evaluate the data controls required and which silos need protection, and decide if any DCAP product options meet these requirements against each category. The early market trends have shown that some vendors already cross silos (see Figures 2 and 3).
- If you only have one RDBMS platform or file storage type such as SharePoint, you may be able to provide the required controls through native capabilities or extraction of log data to other network security tools such as security information and event management (see "Mitigate Breaches With Real-Time Discovery").
- Identify applications that extract data or transfer data to adjacent silos, given certain user privileges. Any new user access to the data must be understood within the new silo to verify if the privileges are appropriate, given the original data sensitivity. Therefore, identify how the controls need to be implemented, with appropriate DCAP product or products to address each silo.
- Given the shortlist of relevant DCAP products, analyze the internal architectural loading and network communications constraints to determine each product's ability to deliver the necessary monitoring and data protection.

Gartner Recommended Reading

Some documents may not be available as part of your current Gartner subscription.

"Competitive Landscape: Database Audit and Protection, Worldwide, 2013"

"Magic Quadrant for Data Masking Technology"

"Magic Quadrant for Content-Aware Data Loss Prevention"

"Apply the Nine Critical Capabilities of Database Audit and Protection"

"Enhancing Security and Compliance with Database Audit and Protection"

Evidence

[1] Over the past year Gartner has spoken to more than 300 clients to discuss requirements for data protection and activity monitoring. These discussions have highlighted the shortcomings of existing siloed products and their own lack of coordination of a data security policy across silos. Discussion with many vendors has also highlighted their desire to address unmet needs for data security governance and data security policy through a DCAP approach.

Note 1 Data Security Governance

A subset of information governance that deals specifically with protecting corporate data (in both structured database and unstructured file-based forms) through defined data policies and processes, and implemented via technologies that are drawn from solutions such as database audit and protection (DAP), data access governance (DAG) and data loss prevention (DLP), among others.

GARTNER HEADQUARTERS

Corporate Headquarters
56 Top Gallant Road
Stamford, CT 06902-7700
USA
+1 203 964 0096

Regional Headquarters
AUSTRALIA
BRAZIL
JAPAN
UNITED KINGDOM

For a complete list of worldwide locations,
visit http://www.gartner.com/technology/about.jsp