Prioritize Enterprisewide Encryption for Critical Datasets

Published: 28 June 2017 **ID:** G00331161

Analyst(s): Brian Lowans

Security and risk management leaders face a rapidly increasing volume and variety of complex data security challenges, both on-premises and in the cloud. Datasets must be prioritized for enterprisewide encryption to mitigate certain security threats and compliance requirements.

Key Challenges

- The dramatic year-over-year growth in the number of data breaches, coupled with increased data residency and compliance requirements, is creating new challenges on how to focus prioritization and protection across diverse infrastructures, on-premises and in the cloud.
- Despite the availability of the Key Management Interoperability Protocol, many vendors' encryption products are using proprietary interfaces that are incompatible with other vendor key managers.
- The wide variety of encryption products and vendors makes the selection process problematic, due to their siloed approaches to data types or storage environments, such as file servers, databases, NoSQL, applications, infrastructure as a service and SaaS clouds.
- The wide variety of access controls and segregation of duties enabled by encryption products makes use-case evaluations important.

Recommendations

Security and risk management leaders must develop application and data security strategies that:

- Prioritize expenditures on the protection of datasets with the highest business risk, consolidate storage, and delete or desensitize old datasets at their end of life.
- Encrypt only what is necessary to match the access controls for each data protection use case.
- Minimize the number of encryption products in use to simplify the orchestration of policies through enterprise key management.

Table of Contents

List of Figures

Strategic Planning Assumption

Through 2020, driven by the increasing risk of a data breach, more than 50% of enterprises will purchase enterprisewide encryption products, which is a significant increase from fewer than 20% today.

Introduction

Security and risk management leaders face increasingly complex decisions regarding the protection of stored data, due to a dramatic growth in data breaches (see data compiled from the Identity Theft Resource Center, depicted in Figure 1).

Figure 1. The Number of U.S. Organizations Hacked Is Growing Dramatically

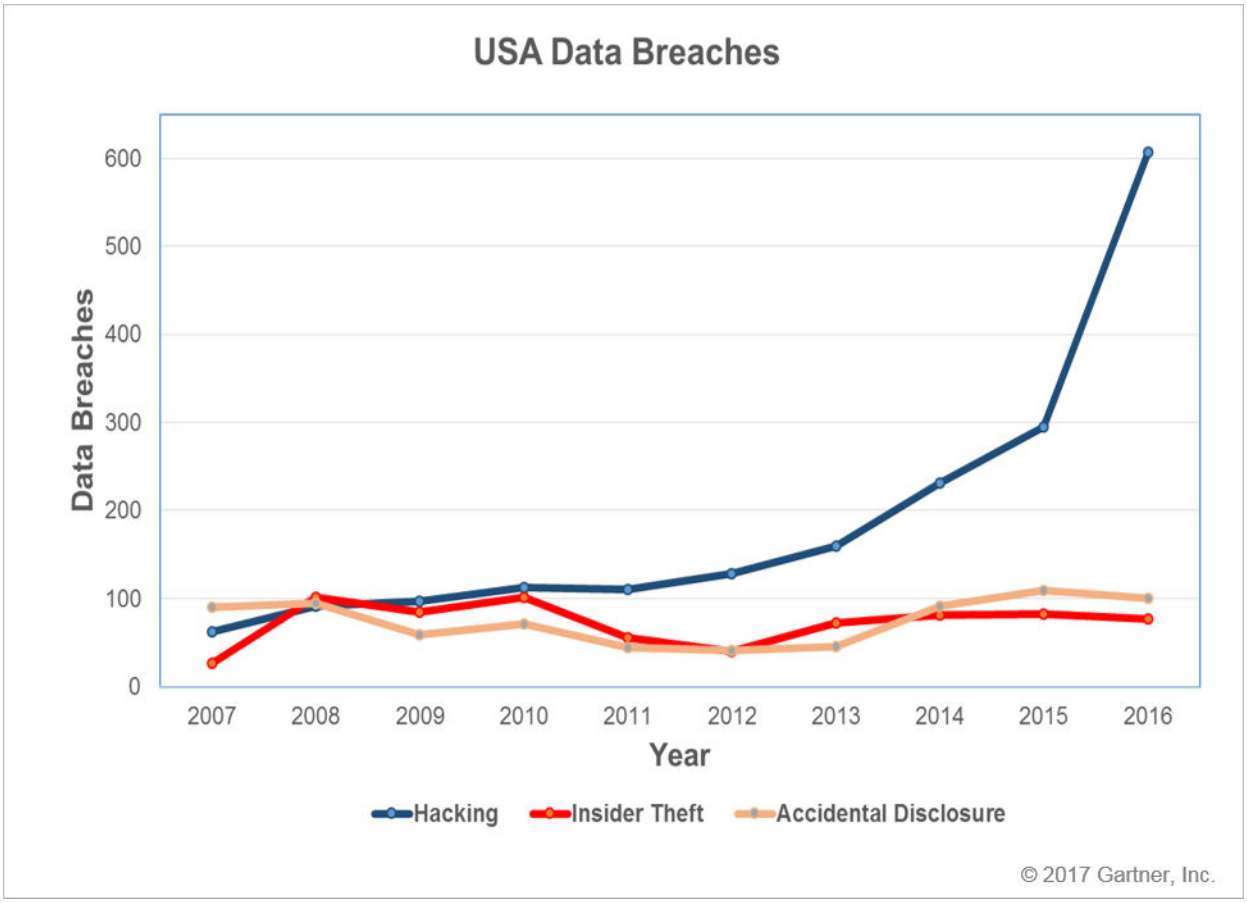

Source: Gartner (June 2017)

The business risks and financial liabilities involved in a data breach (see "Use Infonomics to Reset Data Security Budgets") vary with each dataset, and can include:

- Personally identifiable information (PII)
- Electronic protected health information (ePHI)
- Credit card information for the Payment Card Industry Data Security Standard (PCI DSS)
- Company financial information (CFI)
- Intellectual property (IP)

More than 100 national data privacy laws, including the upcoming general data protection regulation (GDPR) in Europe, require the protection of data when it moves across borders. In addition, internal and external auditors are applying increasing pressure for best-practice data protection levels, such as access controls and encryption (encryption is not always mandatory).

Data encryption is not a sufficient control by itself, because any external attack that compromises user identities, or actions by a malicious insider with appropriate access privileges, will be able to access sensitive data. Attackers want to access user accounts, escalate privileges, and access encryption keys or password stores. Application layer attacks using SQL injection (SQLi) are frequently used to access unprotected databases.

Evidence continues to show that the overwhelming focus of external attacks and insider abuse is on gaining access to sensitive data stored on-premises in file and database servers.[1] However, attacks on cloud services are beginning to occur, because more-sensitive data is used there, creating vulnerabilities that will be exploited and which encryption alone will not solve. Cloud data breaches are beginning to be reported.[2,3,4]

SRM leaders face a confusing array of encryption products from vendors, and these tools can be deployed at different levels, including applications, databases and storage media (see Figure 2). To add to the confusion, these products vary in terms of encryption capabilities, security controls, application interfaces and their ability to interface with key management products.

Figure 2. How Encryption Products Relate to Data Encryption Level/Storage Environment

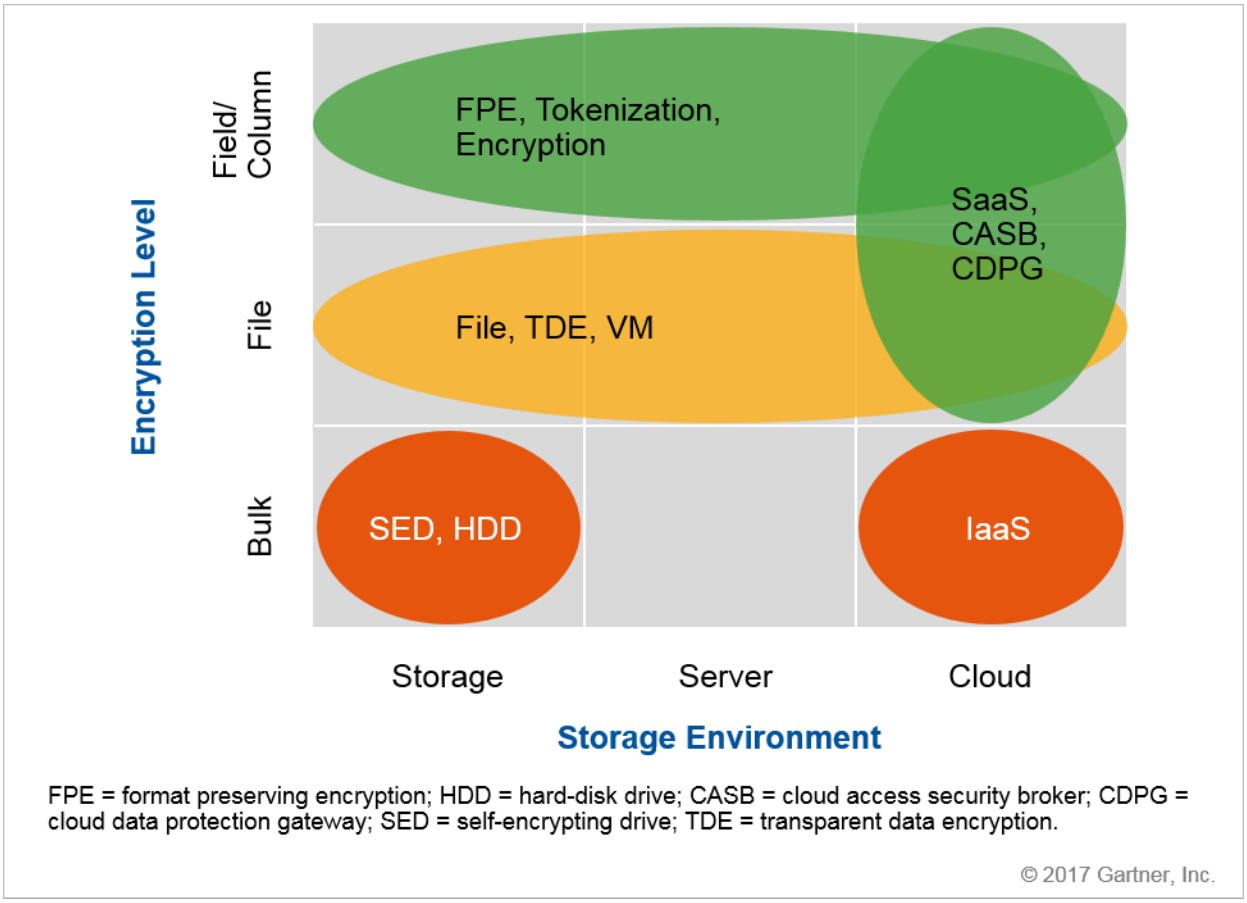

FPE = format preserving encryption; HDD = hard-disk drive; CASB = cloud access security broker; CDPG = cloud data protection gateway; SED = self-encrypting drive; TDE = transparent data encryption.

Source: Gartner (June 2017)

SRM leaders must understand the requirements and limitations of the different encryption products for protecting bulk storage media, files or databases, and, whether the data is stored on-premises, in an off-site data center or in the public cloud. Navigating the encryption and key management options can be daunting. The encryption options are described in the following sections as a guide through this selection minefield and can be used to develop a strategy to apply the appropriate products. However, any strategy must first come to grips with data hoarding and focusing on the classification of sensitive data. Security and risk management leaders will need to identify and prioritize the business risks resulting from the various security threats and compliance requirements applicable to each dataset and storage silo.

Analysis

Prioritize Expenditures on the Protection of Datasets With the Highest Business Risk

Enterprises store a great deal of data unnecessarily, which requires additional management overheads. Data security governance should be applied to classify and prioritize the datasets that have the most significant business risks. Infonomics should be used to assess the financial asset and liability valuations of each dataset to understand the potential financial risks and the security budget impacts (see "Use Infonomics to Reset Data Security Budgets"). If data needs to be retained longer than the business value, then consider desensitizing with appropriate encryption and tokenization. The locations of these datasets must be identified in terms of storage silo and geography.

Deleting data that has no further business purpose or exporting the processing of data to a third party can greatly simplify operations and reduce overheads. For example, financial organizations regularly outsource the processing of PCI data, which has the added benefit of simplifying audits for compliance.

Consolidating the scope and location of regulated or sensitive data and encrypting only what's necessary, and only for the useful lifetime of that dataset with financial risks, will reduce the need for encryption deployments and provide operational cost savings. Risk assessments should be used to prioritize data protection policies. Protection can then be applied through products that employ data classification techniques, such as data loss prevention (DLP), data-centric audit and protection (DCAP; see "Market Guide for Data-Centric Audit and Protection"), encryption, tokenization, and data masking.

Developing a long-term data storage plan that regularly reviews the opportunity to delete or archive old data or, to stop protecting specific silos, is critical to control expenditures. This process can be optimized, if business units have incentives to manage and delete old data. Likewise, the focus of security products on storage silos with reduced scope and number of platforms can form part of an incentivized plan to simplify operations. Not all data is created equal, so prioritize expenditures throughout the datasets' life cycles. Big data and data lake environments may need special attention for desensitization, due to the collection of "dark data" and analytics (see "Securing the Big Data and Advanced Analytics Pipeline").

Encrypt Only What's Necessary

Encryption is essentially a form of access control to silos, files or particular structured fields. It associates access with membership — for example, to Active Directory (AD). Encryption should be used with identity and access management (IAM) controls. For example, any user-granted application, group privileges or entitlements by a data owner will require login credentials to see all the data in a particular silo. Encryption then provides the granular access to individual files or fields, but must be manually orchestrated among these security products. Encryption ensures that data is accessed only by authorized users, individuals within the groups or specific administrators.

Access can be role-based to encompass business-unit requirements, but should always be proactively managed to reflect data requirements, as well as personnel and role changes, with encryption permissions orchestrated across multiple silos. Some compliance regimes may require IAM, segregation of duties (SOD) or activity monitoring. Therefore, maintaining key management logs of encryption and decryption actions will be required. However, policies can only be orchestrated manually across disparate encryption and IAM products.

Due to the complexity of the encryption options, security and risk management leaders must segment the different encryption products in relation to some basic use cases:

- **Breach Mitigation —** Before embarking on an encryption project, enterprises should note that the encryption of centrally stored data is rarely mandatory. Regulations regarding PII and ePHI data, such as HIPAA, the USA State Privacy Laws and the upcoming GDPR, mention encryption, but only offer exemptions in the event of a data security breach.
- **Administrator SOD —** Transparent data encryption (TDE) offers protection of individual databases at rest, because the production database sits as clear text in memory while in use. Database administrators (DBAs) will still have access to production databases. File encryption ensures that files are visible only to authorized users on endpoints. These products support basic SOD to prevent access by system or IT administrators through AD membership.
- **Application User Access Control —** Protecting data at the field level with encryption, tokenization or data masking can prevent specific application users and administrators from having access to sensitive data fields when accessing databases. Security access policy is tied to AD membership and can be used to prevent access by system, IT and DBAs. Depending on how the product is deployed, this may enable access controls to be applied on the database server or the application server. The use of column encryption may require database schema and application changes to accommodate field size changes. FPE and tokenization will affect the ability to create database indexing, search and sorting, etc.
- **Data Residency —** Deploy data encryption when using SaaS or infrastructure as a service (IaaS) in public clouds to address data residency and compliance issues. However, the use of cloud-based storage and applications means data residency requirements are leading to encryption as a best practice (see "Simplify Operations and Compliance in the Cloud by Encrypting Sensitive Data" and "Choosing Between Cloud SaaS and CASB Encryption Is Problematic"). Some regulators offer guidance for cloud encryption related to financial data,[5,6] PII data[7] and PCI data.[8]

Whichever use case is relevant, review the geographic location of encryption key management products and where the encryption/decryption process occurs. Installing strong encryption key management can still be made ineffective, if data is transmitted in the clear between locations or if data is decrypted in an environment that provides potential access to unauthorized staff or users. Deploying encryption key management in certain locations can cause noncompliance. Equally important is the need to develop backup key management storage, because the inadvertent deletion of keys will digitally shred the data (see "Develop an Encryption Key Management Strategy or Lose the Data").

Employ Enterprise Key Management

Enterprise key management (EKM) product capabilities have expanded with the advent of the Key Management Interoperability Protocol (KMIP) standard, which is sponsored by the Organization for the Advancement of Structured Information Standards (OASIS). EKM products can manage any cryptographic products that are compliant with KMIP. However, vendors offering key management products have not converted their cryptographic products to support KMIP; instead they use proprietary interfaces. This protectionist strategy retains control of cryptographic infrastructures; however, despite strong interest from clients, it remains a barrier to adoption. The negative effect of this is that, if multiple vendor products are deployed across different silos, then each product will require separate key management (and backup). These key management silos increase costs and management complexity.

In addition, because key management policies are not automatically orchestrated across silos, this will result in inconsistent user access privileges, which can lead to unintended data exposure and vulnerabilities. Therefore, the deployment of EKM with orchestrated policies across silos is critical, but it may require a single vendor to offer products across all silos.

Traditional "ticking the box" encryption aimed at storage media does not provide access control and will not prevent malicious activities, such as hacking. Therefore, the focus of enterprise solutions has moved toward file and field/column-level encryption products. The rapid emergence of products encrypting cloud applications or storage is being driven by data residency and compliance issues. The threat of hacking is also driving an increase in the adoption of on-premises and cloud encryption products.[9]

We summarize some encryption product options below that can help with the use cases described in the previous section (refer to Figure 2). Figure 3 provides a useful summary of the encryption options and access controls.

Figure 3. Encryption Options and Access Controls

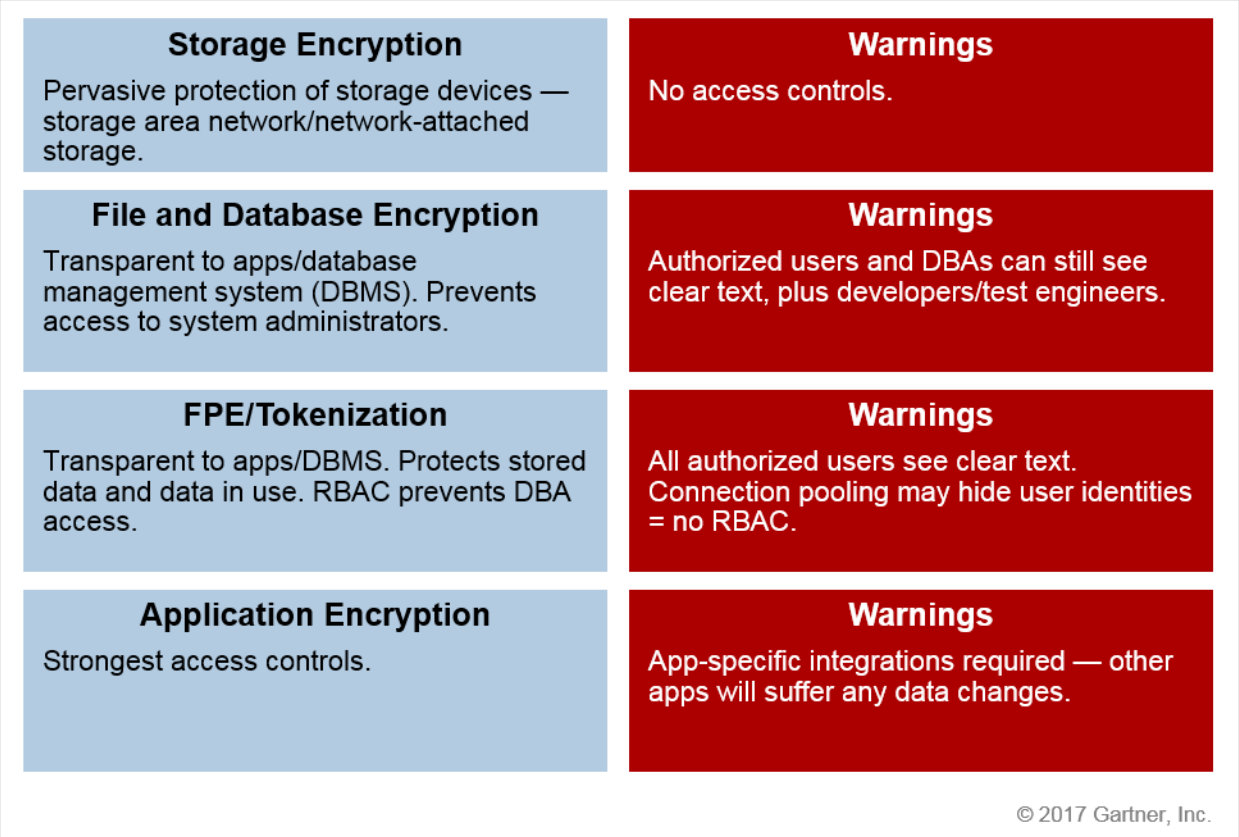

Source: Gartner (June 2017)

SEDs

SEDs are available as hard-disk and solid-state drives (SSDs) with built-in encryption capabilities to support bulk storage encryption deployments. The integrated encryption capabilities provide optimum performance that is unaffected by the number of SEDs installed. SEDs support third-party key management through OASIS KMIP and offer simplicity in retrofitting or new deployments. The higher cost of individual HDDs makes this a viable product for new installations; however, retrofitting a large system before its end of life may not be cost-effective, compared with alternative approaches.

Recommendation: Deploying SEDs is financially viable only for new storage installations.

Use Cases: Provides bulk encryption at rest, but does not provide any access control and has limited compliance benefit, because it protects only against loss or theft of the media from the data center. Most SED products are compliant with KMIP.

Sample Vendors: Hitachi, Micron Technology, Samsung, Seagate, Toshiba

HDD Controllers

HDD controllers manage the flow of data to the physical drive media typically included as part of storage product offerings. These offer bulk storage encryption and scale in proportion to the number of HDDs, with full bandwidth requirements. HDD controller encryption can reduce complexity by retrofitting environments and using existing HDD infrastructures. Most products are compliant with KMIP.

Recommendation: Use HDD controllers to retrofit existing or new installations, and typically offered as a license to existing storage solutions.

Use Cases: Provides bulk encryption at rest, but does not provide any access control and has limited compliance benefit, because it protects only against loss or theft of the media from the data center.

Sample Vendors: EMC, Fujitsu, HP, IBM, Oracle, Symantec

TDE and File Encryption

Vendors that offer the ability to target the protection of unstructured data typically offer both file encryption and TDE for database management systems (DBMSs) at rest. Vendors typically provide a centralized interface or the ability to encrypt multiple DBMS platforms and offer SOD through the management of access control lists (ACLs), typically linked through AD. Some vendors also provide APIs, which offers TDE to applications accessing DBMSs or unstructured files, such as SAP and Microsoft SharePoint. A few relational DBMS (RDBMS) vendors have native TDE encryption tools, such as MariaDB, Microsoft SQL Server, MongoDB, PostgreSQL, Oracle and Sybase, but the keys are accessible and managed by the DBA. However, most enterprises use multiple instances of RDBMSs from different vendors, making their selections unattractive.

Recommendation: Use this encryption when specific files, DBMSs or other common file types need to be encrypted. Protected files can be stored on-premises, in data centers or in the public cloud.

Use Cases: Provides targeted encryption to mitigate the need for breach notification requirements, even if the storage media is lost or stolen. When combined with access controls through applications and endpoints, this can prevent access by administrators or unauthorized users. TDE does not prevent access to DBAs.

Sample Vendors: eperi, Gemalto, HPE, IBM, KSign, Penta Security, Protegrity, PKWare, QuintessenceLabs, Security First, Townsend Security, Thales e-Security

Column Encryption, FPE and Tokenization

Protecting sensitive fields or columns within databases by using FPE[10] or tokenization has emerged as a best practice during the past few years. Tokenization is predominantly used to protect PCI data, due to its simplification of compensatory controls. Dynamic data masking may also be used, but it is outside the scope of this research (see "Market Guide for Data Masking"); however, some

vendors are now offering products that mask data on presentation to application users, in combination with FPE or tokenization. This concept means that the sensitive fields are replaced with randomized data that has the same format as the original data. Competition is growing, and several vendors have introduced new products into this market during the past few years.

Traditional approaches to using column-level encryption are still available, but are less commonly used. These products do not maintain the field structure, which may interfere with the database operation requiring schema changes or software changes to applications. A few RDBMS vendors have native column-level encryption tools, such as Microsoft SQL Server, PostgreSQL, Oracle and Sybase, but the keys are accessible and managed by the DBA. However, most enterprises use multiple instances of RDBMSs from different vendors, making their selections unattractive.

Regardless of the technology, these tools can affect database and application performance through problems that include key management standards, indexing, schema changes or connection pooling.

Recommendation: Use third-party encryption products to enable EKM of multiple-vendor RDBMSs.

Use Cases: Provides targeted encryption to mitigate the need for breach notification requirements, if the storage media is lost or stolen. This can prevent access by database, system or IT administrators, and provides access restrictions to database and application users.

Sample Vendors: Dataguise, eperi, DBSec, eGlobal Systems, Gemalto, HPE, IBM, KSign, Liaison Technologies, Netlib, Penta Security Systems, Protegrity, Thales e-Security

SaaS

Enterprises continue to expand their use of public-cloud-based SaaS. Protecting sensitive data to meet data residency and compliance requirements has led to the use of products that include searchable-encryption algorithms, FPE or tokenization offered by cloud data protection gateways (CDPGs) and CASBs. However, protecting the sensitive data may affect the processing ability of a cloud-based service. For example, FPE or tokenization can be used to enable indexing, searching and sorting, in combination with external tables.

Care is required when using searchable-encryption algorithms, because the vendor's particular implementation will result in weakened security with unquantified risk of cryptanalysis. However, all cryptographic implementations will result in lost functionality for numeric calculations performed in the cloud. Protection can be performed within an on-premises gateway appliance or a hosted appliance, which interfaces with SaaS applications as a reverse proxy. Alternatively, endpoint products can operate on the endpoints and use a forward proxy.

Several SaaS providers — e.g., Salesforce, ServiceNow, Google Docs, Microsoft Office 365, Box, Egnyte and Dropbox — are offering native encryption of structured data and files. However, this means that enterprises must trust the access provided to the SaaS, because key management will reside in that cloud. If not KMIP-compliant, then enterprises must orchestrate policies across

multiple independent key management products (see "Choosing Between Cloud SaaS and CASB Encryption Is Problematic").

Recommendation: The vendors in this market focus on SaaS data protection, and you need to ensure that they integrate with an on-premises EKM.

Use Cases: Primarily used to address data residency and compliance issues in the cloud. Access controls can prevent access by SaaS, local administrators or unauthorized users. They can be used to protect data fields or files while in use or when stored in a public cloud.

Sample Vendors: Bitglass, CenterTools, Centraya, CipherCloud, Cisco (Cloudlock), CloudMask, Covata, eperi, nCrypted Cloud, Netskope, Ohanae, PKWare, Protegrity, Skyhigh Networks, Symantec, Thales, Vaultive

IaaS Storage

Many vendors are developing products to address the encryption and key management of data in public cloud environments or as a SaaS security offering. Typically, these are working at the bulk encryption level in the virtual machine (VM) or individual files accessed from memory. This means that the encryption keys will be available, and data will also be decrypted in memory. The products are maturing, but can operate at the infrastructure or platform as a service (PaaS) level. Transporting data from on-premises to the chosen cloud also requires encryption using protocols such as TLS; however, not all vendors offer this capability, and some require a third party to provide transport-layer encryption plus EKM (see "Enabling High-Risk Services in the Public Cloud With IaaS Encryption").

IaaS encryption products are a good choice to protect everything and to ensure digital deletion at the end of life, but because they don't provide access controls, they'll have limited ability to prevent data breaches.[11] File encryption adds granular access control; however, data will still appear in clear text in memory or accessed by applications. In all cases, the choices of EKM should be considered, as well as the choices of native cloud key management or even bring your own key reviewed as part of a risk assessment.

Several CSPs with IaaS or PaaS capabilities now offer storage encryption and encryption key management, including Alibaba, Amazon, Google, IBM and Microsoft.

Recommendation: A thorough risk assessment is required to review access to data and encryption keys.

Use cases: Primarily used to address data residency and compliance issues in the cloud.

Sample vendors: Cloudera, EMC Cloudlink, eperi, Gemalto, HPE, HyTrust, Protegrity, QuintessenceLabs, Security First, Thales e-Security, WinMagic

Multiple Encryption Scenarios

There is an increasing need to combine encryption products to match multiple use cases or storage environments. The problem is the inevitable conflict between the stated goals of deploying a single EKM product and selecting best-of-breed encryption products for each individual use case. In this scenario, a compromise is needed (see "Develop an Encryption Key Management Strategy or Lose the Data"). Security and risk management leaders must:

- Evaluate the requirements for each data protection use case and compose a list of native and third-party encryption vendors.
- Assess the use-case vendor options and develop RFI/RFP as normal.
- Establish which encryption products require independent EKM.
- Prepare for a cost/management efficiency analysis that will enable you to:
  - Select a best-of-breed strategy, with different vendors' encryption products for each use case; this may require multiple, independent EKM products.
  - Minimize the number of encryption vendors for all storage requirements to minimize the number of EKM products required.

Recommendation: Ideally, focus on one vendor's encryption product, when multiple storage silos need to be protected. Although this will not enable a best-of-breed approach for individual silos, it will provide consistent orchestration of policies through EKM and reduce costs.

Gartner Recommended Reading

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

"Market Trends: Database Security, Worldwide, 2017"

"Better Safe Than Sorry: Preparing for Crypto-Agility"

"How to Make Cloud IaaS Workloads More Secure Than Your Own Data Center"

"Application and Data Security Primer for 2017"

"Focus on Five High-Priority Changes to Tackle the EU GDPR"

Evidence

[1] Verizon's 2017 Data Breach Investigations Report

[2] May 31, 2017 Security Incident (Updated June 8, 2017)

[3] Security Notice Update

[4] CloudPets' Data Breach Underlines Need for Secure Cloud Apps

[5] Cloud Computing: Information Security, Outsourcing Technology Services

[6] U.K. Information Commissioner's Office: "Guidance on the Use of Cloud Computing"

[7] European Commission

[8] Special Interest Groups: Cloud Selected as 2017 SIG Topic

[9] Encryption purchasing trends based on evidence from Gartner client inquiries

[10] Recommendation for Block Cipher: Modes of Operation: Methods for Format-Preserving Encryption

[11] D. Goodin, "Virtual Machine Used to Steal Crypto Keys From Other VM on the Same Server," Ars Technica, 6 November 2012.

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