Securing Sensitive Data

A Comprehensive Guide to Encryption Technology Approaches

Vormetric, Inc.
888.267.3732
408.433.6000
sales@vormetric.com
www.vormetric.com

Executive Summary

Enterprises can choose from a variety of encryption technologies to protect data at rest, each of which has its strengths and limitations. The optimal choice should secure multiple data types in scope, satisfy audit and compliance requirements, and minimize operating costs through centralized management of policies and encryption keys. This paper surveys the various encryption approaches for securing data at rest to understand what the technology does along with where each provides an optimal fit. After comparing the various encryption technologies, this paper explains why file level transparent data encryption provides the optimal balance of data type coverage, security, manageability and operational efficiency.

Securing the data

As data breaches and hacking attempts continue, enterprises are evaluating how to best secure data with encryption to meet compliance and governance requirements. There are a variety of different encryption approaches to consider. Different encryption technologies typically have a "design center" or job for which they are best suited. With the variety of different design centers and feature sets, one can quickly become lost in the minutiae of cryptographic technology.

Before stepping into the technology decision, consider the following fundamental questions that can help guide your decision-making process as you evaluate the different technology approaches.

- What type of information do you want to protect?
- What threats do you want to protect the information against?
- What application and infrastructure changes can you tolerate?
- What are your key management expectations?

The answers to these questions will guide your decision on the type of encryption technology that is optimal for your data security needs.

Technology Approaches

The various technology approaches can be divided into two major categories. Platform specific and those that are transparent to applications and databases:

1. Platform specific encryption refers to encryption that is designed for specific environments that must be specifically managed. Platform specific approaches include Application Encryption, Tokenization, Database Column-level Encryption and Format Preserving Encryption.

2. Transparent approaches include Native Database Transparent Data Encryption, File-level Transparent Data Encryption, and Storage-level Encryption.

What follow are descriptions of the various technology approaches to encrypting data and describes the strengths and challenges of each approach.

Application-level Encryption

Application-level encryption uses software development tools to encrypt data within the application before it is stored in a database or storage infrastructure. Application-level encryption protects data as it is processed within the application and keeps the data encrypted all the way to the storage environment where it will reside. Such efforts involve custom code development that brings encryption into the individual application data fields.

Application-level encryption can effectively "blind the Database Administrator (DBA)" as data arrives at the database already encrypted. Application-level encryption typically requires owning the application code and cannot be applied to commercial off-the-shelf applications, for example, SAP, PeopleSoft, or Oracle Financials.

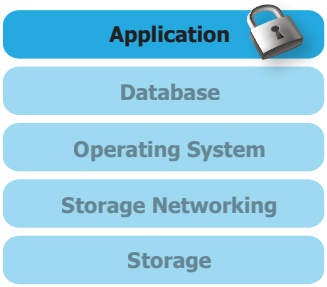

Application-level encryption can be operationally intrusive and take considerable time to deploy, especially if multiple applications are in scope; it is not transparent to the application environment and requires a development team that can program, maintain and support the application as business requirements change. This approach is costly in terms of time, money, and development staff to integrate encryption into the application and provide ongoing application development support. Without creating additional applications to support them, application-level encryption does not apply to unstructured data such as reports and spreadsheets. Operating system and programming language support are typically limited to what the application-level encryption vendor offers.

> **"Once again application encryption isn't a panacea – it can work well, but brings additional complexities and is very easy to screw up. Use with caution."**
>
> - **Securosis,** "Understanding and Selecting a Database Encryption or Tokenization Solution", 2010[1]

Tokenization

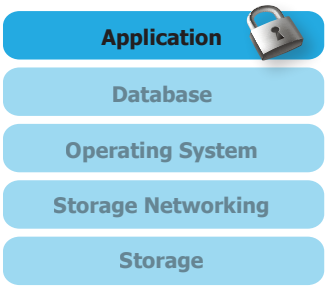

Tokenization is the process of replacing sensitive data with unique identification symbols (a "token") that acts as a proxy for the original information. The original data is kept in a master database that can be hardened, encrypted and keeps track of which token matches which original piece of data. The tokenization approach seeks to minimize the amount of sensitive data a business needs to keep on hand and typically is applied to a single field or column (e.g. credit card numbers, social security numbers). By swapping out all the sensitive numbers, most of the data security can be focused on the one system that's easier to control. Tokenization can be applied inside an enterprise environment or via a service provider. A typical use case is for a merchant managing cardholder data to contract with a service provider to handle the issuance of the token value and the provider bears the responsibility for keeping the cardholder data secured.

Tokenization can require some custom coding in the application ecosystem to accommodate the token. Tokenization does not lend itself easily to multiple data types.

Tokenization lends itself to situations where there are a few, extremely sensitive pieces of information to protect, for example, protecting credit card numbers in order to comply with PCI-DSS regulations. However, it tends to break down when there are more than one or two pieces of information to protect, as each conversion from token to real data (or vice versa) requires a network roundtrip to the tokenization appliance.

Format Preserving Encryption

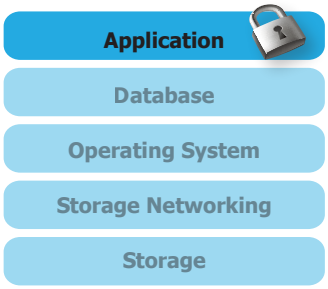

Format Preserving Encryption (FPE), also called Datatype Preserving Encryption (DPE), attempts to maintain the same data size and format as the original data being encrypted, but substitutes encrypted data without needing to change systems that use such data (e.g database structures, queries, applications, etc). FPE tries to avoid having to modify databases or datatypes to accommodate the larger encrypted data (cipher text occupies more space than clear text). While FPE might avoid having to modify application code, it can require enterprises to add and maintain FPE application code. FPE can become complex beyond one use case and poses security challenges since multiple trusted systems have access to decryption keys and decrypted data.[2] FPE does not apply to unstructured data such as spreadsheets, reports, or Extract-Transform-Load (ETL) files.

1 Securosis, "Understanding and Selecting a Database Encryption or Tokenization Solution", 2010
http://securosis.com/reports/Securosis_Understanding_DBEncryption.V_.1_.pdf

Column-level Encryption

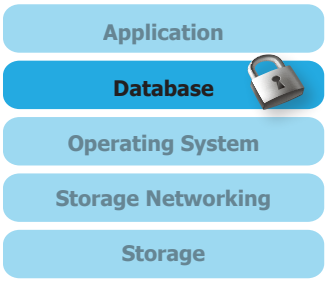

Database column-level encryption consists of database modules that utilize views, triggers, stored procedures, and external functions to encrypt structured data in a specific database column. Microsoft refers to this approach in SQL Server as "cell-level encryption". Column-level encryption was the initial step on the part of some database vendors to provide native encryption inside of the database. Most of the major database vendors now offer transparent tablespace-level encryption and many of these database vendors now position tablespace-level encryption as a superior approach than column-level.

Contrary to the impression that column-level encryption is more efficient, encryption inside of the database can be a performance burden for many reasons, for example, indexes on an encrypted column offer no advantage, so index range scans turn into poorer performing alternatives. A Microsoft Technical Article regarding Microsoft's cell-level encryption approach commented that, "performance for a very basic query (that selects and decrypts a single encrypted column) when using cell-level encryption tends to be around 20% worse [than Transparent Data Encryption]". [3]

Column-level encryption has the additional challenge that credentialed users, such as DBAs, with adequate privileges can view the encrypted data and manage the encryption.[4]

Column-level encryption does not apply to unstructured data outside of the database, cannot be extended beyond one vendor's database, and typically requires some key management solution on top of the encryption offering. Organizations needing to protect both structured and unstructured data and wanting to use column-level encryption would need to deploy an additional encryption solution to protect unstructured data.

> **"...keys are only as secure as your least trustworthy [Database Administrator] DBA. In any platform that stores the keys in a database table, they are accessible to database administrators. Despite vendor advertisements to the contrary, there are attacks that can sniff key operations within the database, or even scan through database memory structures where the keys reside during use."**
>
> - **Securosis,** "Understanding and Selecting a Database Encryption or Tokenization Solution", 2010[5]

Native Database Transparent Data Encryption

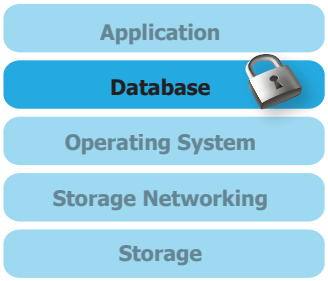

Transparent Data Encryption (TDE) provided by a database vendor encrypts tablespaces or columns within the database via functionality within the database engine. Credentialed users (DBAs) within the database have access to the unencrypted data (DBA is not "blinded"), so database and third party tools such as Database Activity Monitoring (DAM) are needed to monitor database usage.

2 Securosis, see previous citation
3 "Database Encryption in SQL Server 2008 Enterprise Edition", SQL Server Technical Article, February 2008, http://msdn.microsoft.com/en-us/library/cc278098(v=sql.100).aspx
4 Vormetric, Inc., "Debunking The Myths of Column-level Encryption", August 2011
5 Securosis, "Understanding and Selecting a Database Encryption or Tokenization Solution", 2010 http://securosis.com/reports/Securosis_Understanding_DBEncryption.V_.1_.pdf

Internal database TDE is limited to structured data within the database and cannot be used to protect unstructured data outside of the database (reports, spreadsheet extracts, trace files, Extract-Transform-Load (ETL) data, etc). Native database TDE applies to a single vendor's database and related data, so a heterogeneous environment with multiple vendor databases would require deploying the cost and management overhead of multiple database TDE solutions. TDE typically applies to the most recent database releases and is not available for many older database versions. Internal database TDE suffers from mediocre key management and typically requires the addition of a key management solution when deployments grow beyond a handful of databases.

File-level Transparent Data Encryption

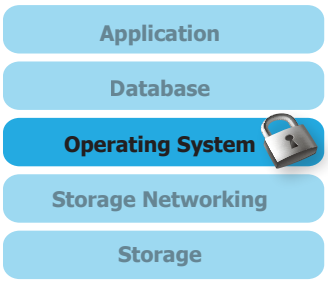

File-level encryption involves encrypting data at the file level within the operating system. This approach to encryption can be used to encrypt both structured data inside the database along with unstructured data outside of the database. File-level encryption is transparent to applications and databases, and can optionally provide access controls that enable separation of duties between different roles (DBA, operating system administrator, application developer, etc). File-level encryption is extensible across a variety of operating system environments and provides an extensible approach that secures multiple data types, databases and applications. Similar to internal database TDE, this approach leaves secured data visible to DBAs and is consequently augmented with database activity monitoring (DAM) tools.

Native File-System Encryption (Microsoft EFS, etc.)

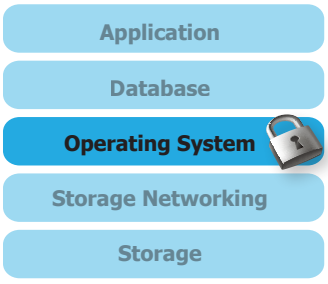

File-system level encryption is provided by some operating systems such as Microsoft's Encrypting File System (EFS). EFS is a component of the NTFS file system on Windows 2000, Windows XP Professional, and Windows Server 2003.[6] EFS is only as strong as the local administrator account or account under which it operates. Compromise of the passwords associated with these accounts, either due to malware intrusion or weak password selection, and EFS-protected files can be at the equivalent risk to their unencrypted neighbors.

In the case of Microsoft ETF, EFS is an NTFS-only offering and is specific to Microsoft environments. When files are copied either to network drives or to external media such as USB thumb-drives (that are customarily FAT32 formatted), they are no longer encrypted. The default EFS settings offer little practical protection but can, with effort, be tuned through group policy into something much stronger. This centralized management is only available in Vista or Windows 7 environments.

Storage Switch Encryption

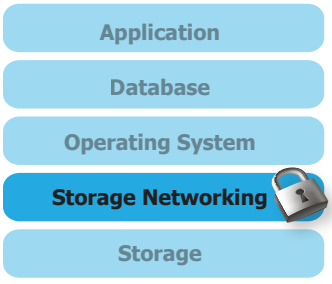

Storage switch encryption refers to using storage networking switches to encrypt data being written to storage devices. Storage switch encryption typically protects against the possibility of data loss from the physical theft of the storage media that data resides on or accidental loss of the storage media. While storage switch encryption is transparent to applications, databases and the servers accessing the data, it is typically storage infrastructure-specific; storage switch encryption applies to Storage Area Networking (SAN) but not to Network-attach Storage (NAS).

Storage switch encryption typically does not provide security functionality such as access control, separation of duties, or fine grained audit logging.

6 "An Overview of the Encrypting File System", http://technet.microsoft.com/en-us/library/cc700811.

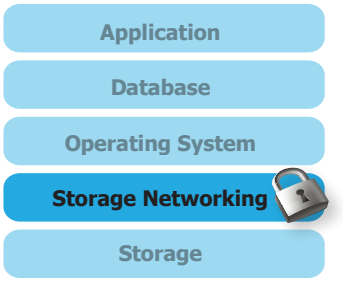

In-line Storage Encryption

In-line storage encryption refers to using in-line storage hardware technology to encrypt data being written to Network Attached-Storage (NAS) devices or a Storage Area Network (SAN). As with storage switch encryption , in-line storage encryption typically protects against the possibility of data loss from the physical theft of the storage media that data resides on or accidental loss of the storage media. While storage encryption is transparent to applications applications, databases and the servers accessing the data, it could require modifying the storage infrastructure and providing multiple paths to the data to avoid a single point of failure.

Storage encryption typically does not provide security functionality such as access auditing, access control, and separation of duties which limits its ability to control threats originating from the servers accessing the data.

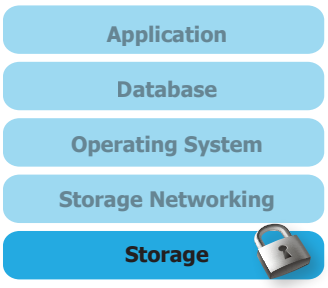

Disk-based Storage Encryption

Disk-based storage encryption refers to disk drives using the emerging Trusted Computing Group (TCG) Opal standard for self-encrypting disk drives (SEDs) in which the actual disk drive includes encryption functionality. The SED approach typically applies to mobile endpoints (laptops) and is not as typical in datacenter environments. This approach does protect against lost devices, but does not provide the separation of duties or auditing found in other technology approaches. It typically requires an additional solution for SED credential management.

Cloud Data Encryption

Encryption of data typically uses a variation or combination of some of the approaches mentioned previously. There are various delivery models for cloud computing: Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastructure-as-a-Service (IaaS). (See Vormetric white paper "Securing and Controlling Data in the Cloud".)

Typically the SaaS and PaaS provider handles data security as part of their overall service offering. In the context of infrastructure-as-a-Service (IaaS) where the user contributes to data security, data encryption can occur at the mounted storage volume layer, in which case there is little granular access control because the mounted storage volume is either accessible to all or not accessible. Other approaches can encrypt individual files and can provide granular access control.

An important variable in cloud data encryption lies in the location of encryption key management infrastructure. Encryption key management in the cloud encryption context can be placed in one of a number of locations including the enterprise datacenter, the cloud service provider, a cloud broker, or via a third party Software-as-a-Service (SaaS) key management portal. Enterprises need to weigh their risk tolerance and audit concerns to arrive at the optimal approach to encryption key management in the cloud.

Maintaining The Optimal Balance: Vormetric Data Security

Vormetric Data Security enables enterprises to meet their compliance requirements and executive mandates with proven encryption and key management, superior manageability and extensibility compared to alternative approaches. Vormetric protects information by transparently encrypting data at the file and or volume level of the operating system. In addition to securing database information irrespective of the vendor, Vormetric can also secure unstructured data outside of the database without application or IT operational changes. Vormetric Data Security protects data in physical, virtual and cloud environments while avoiding any changes to the application, database or storage infrastructure. This approach allows enterprises to better secure their sensitive data and achieve operational efficiency without modifying existing applications or data.

The major benefits provided by Vormetric Data Security include:

Reduced Administrative and Operational Costs

- Protects structured and unstructured data accessed by Linux, UNIX and Windows systems in physical, virtual and cloud environments.
- Complete encryption solution includes integrated key management that avoids the cost of acquiring and managing HSMs and third-party key management software typically needed for TDE or column-level encryption.

Reduced Risk through a Unified Data Security Solution

- Controls privileged user access (System Administrators, etc) and allows them to perform tasks without exposing sensitive data.
- Single solution provides common policy framework for accessing both structured and unstructured data.

Rapid, Cost-Effective Deployment

- Vormetric Data Security is transparent to user operations, applications, databases and storage operations.
- High performance encryption maintains service level agreements.

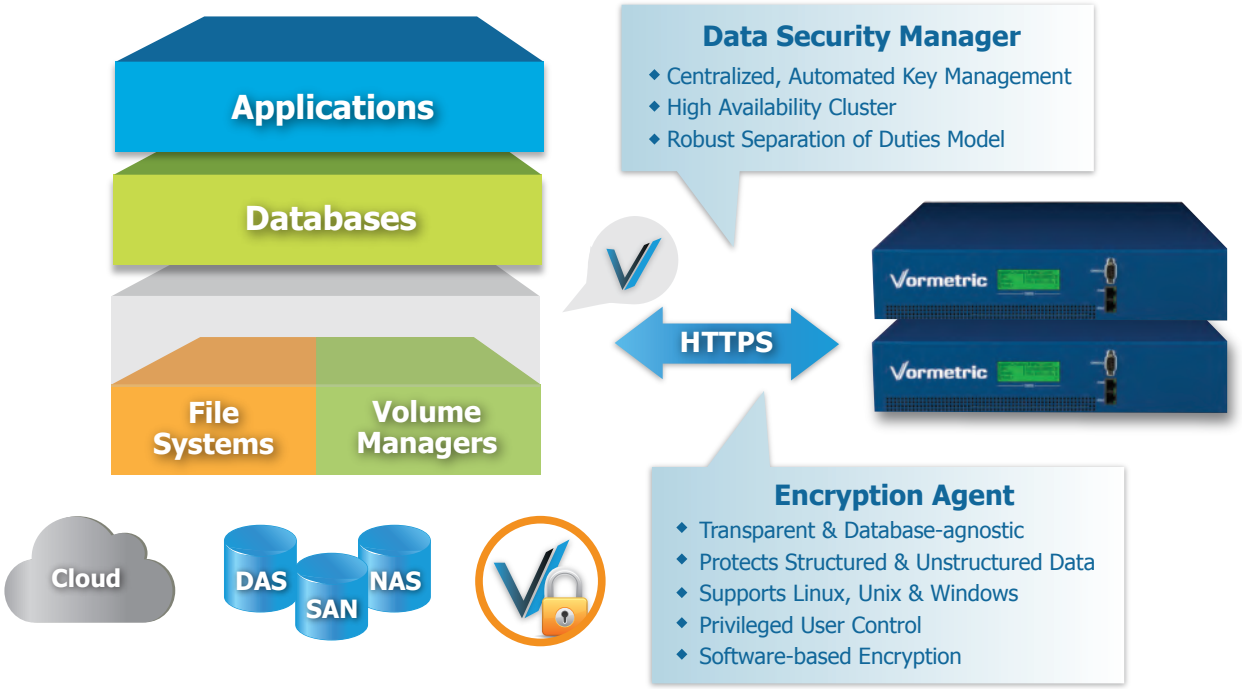

Extensible Solution for Structured and Unstructured Data

Vormetric Data Security can secure both structured and unstructured data to satisfy rigorous audit and compliance requirements and provide comprehensive protection for sensitive data. While data at rest inside of the database can catch the attention of auditors, the data inserted into the database and extracted from the database can be of equal importance to the auditor validating security.

Vormetric can protect sensitive data residing in reports, spreadsheet extracts, archives, Extract-Transform-Load (ETL) data or pdf files. Hackers and rogue employees can use such data stored outside of the database to obtain sensitive information.

Vormetric can evolve as your enterprise's data security requirements evolve in ways that are not possible with alternative encryption approaches.

Universal Solution for All Databases

Vormetric Data Security minimizes administrative overhead with a single key and policy management console providing a secure, easy method of administering encryption keys. This enables organizations to establish consistent and common best practices for managing the protection of both structured and unstructured data.

Operational Efficiency through Encryption Key & Policy Management

Vormetric Data Security provides robust key management along with granular and configurable auditing and reporting of access requests to protected data, as well as changes to policies and keys. The system's audit management reduces audit scope, integrates with existing Security Information & Event Management (SIEM) solutions, and aids compliance with industry and regulatory practices regarding the handling and protection of private and confidential information.

Performance

Benchmarking from a variety of customers has demonstrated the Vormetric Data Security solution provides superior performance relative to native database TDE and other encryption approaches. Vormetric performs encryption and decryption operations at the optimal location of file system or volume manager and consequently minimizes performance overhead. Vormetric's extensive OS and file system expertise and understanding of how to optimize the usage of the AES encryption algorithm provides for optimal system performance while minimizing the encryption CPU requirement.

Future-proofed Transparent Encryption

IT infrastructure and security is changing at a rapid pace with a steady flow of new applications and evolving compliance mandates. Enterprises are embracing virtualization and considering what applications to deploy in the cloud. To maximize their return on IT investments, enterprises need data security solutions that can evolve as their requirements change. The solution for protecting a database today might expand to include different vendor databases or "big data" in the future. Vormetric supports physical, virtual and cloud environments across the dominant operating systems to protect sensitive data wherever it resides.

Conclusion

The data within enterprise's IT infrastructure is the lifeblood that permits efficient business operations. An optimal solution to securing sensitive information needs to provide operational efficiency and robust security. A careful evaluation of encryption technology alternatives can ensure that your business selects an approach that minimizes operational costs and maximizes business flexibility as business requirements evolve.

Vormetric enables operational efficiency with a single solution for both structured and unstructured data combined with strong key management. Vormetric helps you to avoid the security limitations posed by other encryption approaches by protecting both structured and unstructured data, minimizing administrative burdens such as managing encryption keys in complex, heterogeneous environments, and providing extensibility as enterprise data security needs evolve.

With Vormetric Data Security, enterprises can transparently protect their data today with the extensibility to evolve to meet tomorrow's needs to encrypt multiple data types, both structured and unstructured, across multiple platforms and use cases. Vormetric enables your business to minimize operational costs of securing data while providing the flexibility to evolve as your data protection requirements evolve.

About Vormetric

Vormetric is the leader in enterprise encryption and key management for physical, virtual and cloud environments. The Vormetric Data Security product line provides a single, manageable and scalable solution to manage any key and encrypt any file, any database, any application, anywhere it resides— without sacrificing application performance and avoiding key management complexity. For more information, please call: (888) 267-3732 or visit: www.vormetric.com.