



# Achieving Data Privacy Through Database Encryption

CoreGuard™ Information Protection Solution protects against theft of sensitive information from the IBM® DB2® Universal Database™ product by encrypting data at rest and tightly controlling data access.

June 2005



## Contents

Perspective .....	3
The Problem .....	4
The Solution.....	5
CoreGuard Access Control.....	6
MetaClear™ Data Encryption .....	8
DB2 Universal Database Server protected by CoreGuard.....	9
Achieving Data Privacy and Protection Objectives .....	9
Solution Applications .....	10
Summary.....	11

*“By 2005, enterprises that don’t encrypt stored sensitive data will spend 50 percent more than enterprises that do, due to failure to comply with regulatory or contractual data protection requirements (0.7 probability).”*

*“When and How to Use Enterprise Data Encryption”*  
Rich Mogull , Gartner  
March 2004

### Perspective

Pervasive network connectivity is providing organizations with the opportunity to leverage their digital assets, resulting in enhanced productivity and increased profitability. Customers, vendors, partners, and other constituents are able to access the information they require in real time, improving communications and reducing costs. At the same time, this increased accessibility to sensitive information increases its vulnerability to malicious activity and misuse from both within and outside of the organization. Perimeter and point security solutions alone have proven themselves inadequate for protecting valuable stored data from sophisticated internal and external threats.

Today, much of an organization’s sensitive data is stored ‘online’ in databases. Databases supply information for a variety of applications that support an organization’s day-to-day operations, as well as storing historical transactions. Particularly sensitive information may include customer lists, financial results, non-public personal information (NPI), credit card numbers, purchase and sales records, access codes, health records (EPHI), corporate intellectual property, and confidential government information.

A number of privacy regulations, such as the Gramm-Leach-Bliley Act, HIPAA, the Data Protection Act (EU), and California legislation SB 1386 and AB 1950, were created to make data owners responsible for protecting sensitive personal information from inappropriate use. Compliance guidelines and best practice standards generally include the use of specific technologies, including data encryption, user authentication and authorization controls, detailed auditing of access events, and protection from malware. Non-compliance with these regulations can result in significant sanctions and penalties, damage to corporate branding and reputation, and class-action lawsuits brought by the individuals whose personal information was compromised.

In addition, organizations such as financial institutions, public companies regulated by the Sarbanes-Oxley Act and government agencies are increasingly subject to external auditing of their IT security practices. Security auditors look for technical controls that prevent bypassing operational processes and procedures, particularly in regards to financial transaction processing. The controls need to include protection of not just the data itself, but also audit logs that provide irrefutable evidence of data accesses, configuration changes and other operations.

Managing the balance between information availability and privacy assurance can result in significant tension between security and IT operations organizations when security measures have a negative impact on business operations. Successfully resolving this tension requires the adoption of solutions distinctly designed to protect sensitive information without interrupting established processes and procedures.

**The Problem: Despite DBMS security, sensitive data remains vulnerable**

Much of the sensitive information that requires both a high level of accessibility and strong protection is stored in online databases. The increased frequency of malicious activity targeted at databases is widely acknowledged. In a recent survey, 20% of the respondents reported a direct security breach against a database over the past year<sup>1</sup>, with an untold number of security breaches going undetected or unreported.

Some applications such as the IBM® DB2® Universal Database™ product (DB2 UDB) offer sophisticated mechanisms that control access to data. However, all applications that store data on disk remain vulnerable to ‘side-door’ attacks, such as host system compromise via abuse of root privilege, or via the data management process, e.g., backup/disaster recovery. Cleartext (unencrypted) information is vulnerable to a number of security threats. Vulnerabilities of any application that stores data include:

<b>Privilege Theft</b>	The ability to illegally obtain ‘trusted’ root access privileges. <b>Methods</b> include: Password cracking, buffer overflow, local login and privilege escalation.
<b>Application Tampering</b>	The alteration or insertion of executable code for the purpose of running an unauthorized version of the application. <b>Methods</b> include: Trojans, worms, unauthorized patches and altered or corrupted code.
<b>Unintended Privilege</b>	The use of root access or DBA privileges to access and view information outside of the requirements of a user’s authorized role. <b>Methods</b> include: Copying of data, control and user files or tables, transmission of data, viewing of data, and deleting or altering data and access logs.
<b>Data or Log Tampering</b>	The alteration of data or audit logs. <b>Methods</b> include: Access to data at rest, including database or file content and audit logs, via (i) the operating system, (ii) in transit to or from storage, or (iii) local access to the storage system.
<b>Storage Media Theft</b>	The theft of storage media or hardware from a datacenter or in transit to or from a remote vaulting facility. <b>Methods</b> include: Physical theft.

---

<sup>1</sup> Evans Data Corporation, “Database Developer Survey, Volume 2 2002.” [http://www.evansdata.com/n2/surveys/database/database\\_02\\_2\\_xmp2.shtml](http://www.evansdata.com/n2/surveys/database/database_02_2_xmp2.shtml).

Critical data in audit logs, configuration files and resource libraries also require protection from tampering and unauthorized viewing to ensure system integrity and auditability. Implementation of strict security principles for segregation of duties may require that the DBA be blocked from accessing these files.

Controlling audit log access prevents the DBA from viewing and tampering with logs to cover any inappropriate accesses or changes to sensitive data. Audit log files, therefore, must be stored outside of the database table space where the DBA cannot access them. Storing the files in an operating system directory, however, leaves the files outside the sphere of DBMS security, and vulnerable to access by other unauthorized users if not properly protected.

In order to address these vulnerabilities and protect against the theft or unauthorized viewing of sensitive data, vigilant planning requires setting objectives for protecting sensitive data, including:

- **Authentication:** Ensure that all users requesting data access, including root users, are properly authenticated via approved authentication services (e.g., LDAP, Active Directory, etc.);
- **Authorization:** Protect data files in the application data directory or in any customized data file directories from access by unauthorized users;
- **Data Integrity:** Prevent unauthorized deletion of sensitive information and tampering with application files;
- **Confidentiality:** Protect copies of database files and exported data contents;
- **System Integrity:** Protect operating system image, binaries and configuration files, as well as prevent access logs from being altered or tampered; and
- **Auditability:** Audit, alert and logging of significant events, such as security policy violations, systems intrusions, and data management accesses for later forensic analysis.

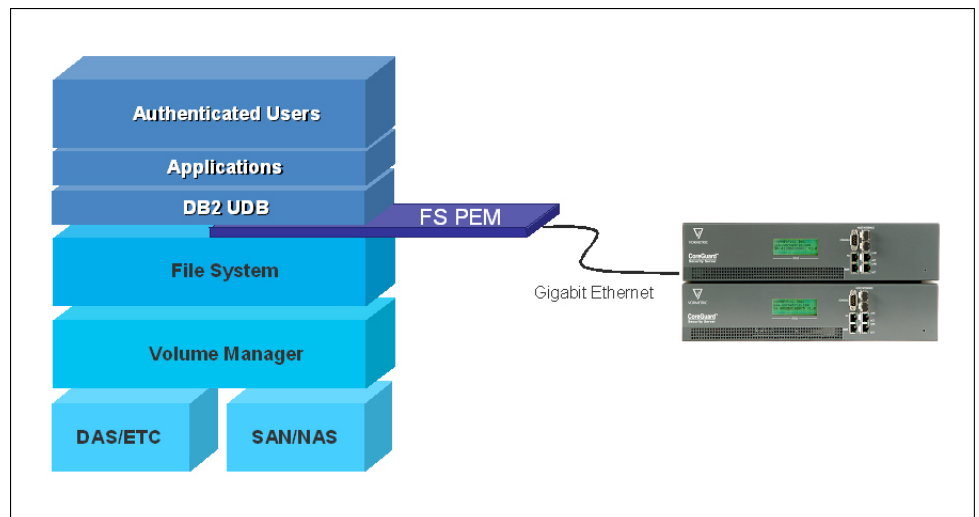
### **The Solution: CoreGuard protects sensitive information stored in DB2 UDB**

The CoreGuard™ Information Protection System protects sensitive information and audit logs using transparent, high-speed encryption of stored information, context-aware access control, advanced host and application integrity protection, and detailed auditing of access events. The CoreGuard solution consists of Security Server appliances and software Policy Enforcement Modules (PEMs) that are installed on the protected hosts serving as access points to sensitive information, whether via network access or direct attached. The PEM acts as the enforcement point for data privacy protection policies, and the Security Server acts as the policy decision point. The CoreGuard architecture enables easy-to-manage, flexible mandatory access control and separation of policy management

## Data Privacy Through Database Encryption

from host management, enforcing the separation of duties between IT administration and security administration.

The FIPS 140-2 validated Security Server provides multiple services, functioning as a repository for data access policies, encryption keys, and cryptographic fingerprints of all executable files and resource libraries allowed to run on PEM-protected hosts. The CoreGuard PEM installed on the protected DB2 server intercepts all calls to protected file domains, shielding the applications and examining requests for critical context attributes. CoreGuard checks the authenticity of every application executed on the protected host—in the case depicted below, DB2 UDB itself—by validating its cryptographic fingerprint. This prevents the use of unauthorized root (system administrator) privileges to tamper with application binaries and the insertion of Trojans or ‘back-door’ access modules. The entire solution functions independently of DB2 UDB, applications, file systems, and storage architectures. Since the CoreGuard PEM requires a mounted file system for operation, independent of imbedded database file system services, support for raw devices is achieved through VERITAS Quick I/O.



*Fig. 1 – The CoreGuard Policy Enforcement Module protects the host system from attack by preventing access to the operating system by unauthorized processes and users.*

### CoreGuard Access Control

In order to access or view protected data, the request's context attributes must match the attributes of a predefined information protection policy. Without this permission in place, the response is a denial of access or, if chosen by the security administrator, a logging of the unauthorized data access attempt.

## Data Privacy Through Database Encryption

Information protection policies consist of elements that define the *who*, *what*, *where*, *when* and *how* of the access attempt.

**Who** – *Who* refers to the process User ID of the application requesting access to the database. This may be a DB2 instance running with access privileges or an authorized backup application requesting access to copy data. Without access permission, not even root users are permitted to access, view, write, or copy to the database, thus preventing abuse of root privileges. To access data, users must also be validated by an authentication service such as LDAP, Active Directory, Kerberos. Without this validation, unauthorized applications, process users, or even modified applications cannot access protected data stores. This protection, coupled with the native DB2 UDB front-end access controls, ensures comprehensive data privacy and protection.

**What** – *What* is the application program used to access sensitive data. This may be a DB2 UDB or a management utility program requiring access to the data files. Cryptographic fingerprints of each of the executables should be made by the authorized security administrator before the application is added to a data privacy protection policy. Cryptographic fingerprints ensure process authenticity, a guarantee that the application has not been tampered with prior to initiating a session.

**Where** – *Where* refers to the protected file domain where the data is located, which may correspond to a data location, such as a directory or file, or to a file type, defined by wildcard symbols in conjunction with particular file names or extensions. A strong AES 128- or 256-bit data encryption key can be created and assigned to a virtual file domain for data viewing control, if desired. A single encryption key can be assigned to as few or as many file domains as desired, such as databases or table spaces or tables or indexes, enabling the creation of secure silos of information.

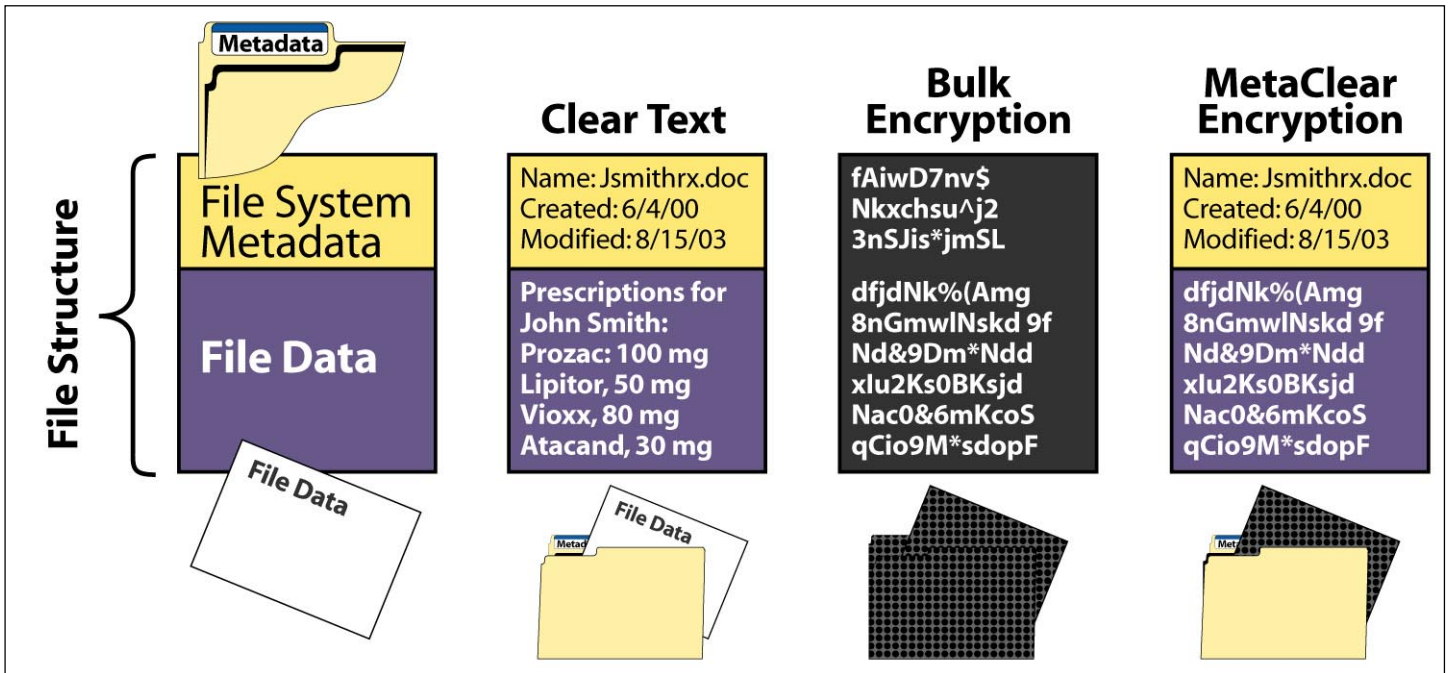
**When** – *When* defines the accessibility to protected data based on the time of day or the day of the week. A backup application, for example, may be limited to accessing data only during a predefined window of operation, limited by the time of day or the day of the week. Thus, access may be limited to a certain time window so that 'trap-door leave behinds' cannot be enabled, thus preventing the compromising of sensitive data.

**How** – *How* refers to CoreGuard's ability to separate **access** to data (for the purpose of managing the data) from the ability to **view** that same data (to use or manipulate actual file content). Access control is handled by the *who*, *what*, *where* and *when* discussed above. *How* determines viewing privileges. Thus, a storage administrator may be granted access to manage data with the file system metadata presented in cleartext, but with the file content remaining in ciphertext.

### MetaClear™ Data Encryption

Data privacy and protection is further enhanced by separating data access privileges from the ability to view that data. Encryption is increasingly viewed as a means of ensuring that only authorized users are able to view sensitive data, thereby preventing unauthorized viewing of data. (See Gartner “When and How to Use Enterprise Data Encryption,” Rich Mogull, 18 March 2004.)

Vormetric’s MetaClear data encryption technology works at the file system level, where it can distinguish between file content and file system metadata. By encrypting only the file content while leaving the file system metadata in the clear, data owners can control who can view data without disrupting management of the data. MetaClear encryption assures transparent operation for data management applications, while still protecting the file content. See the diagram below for details.

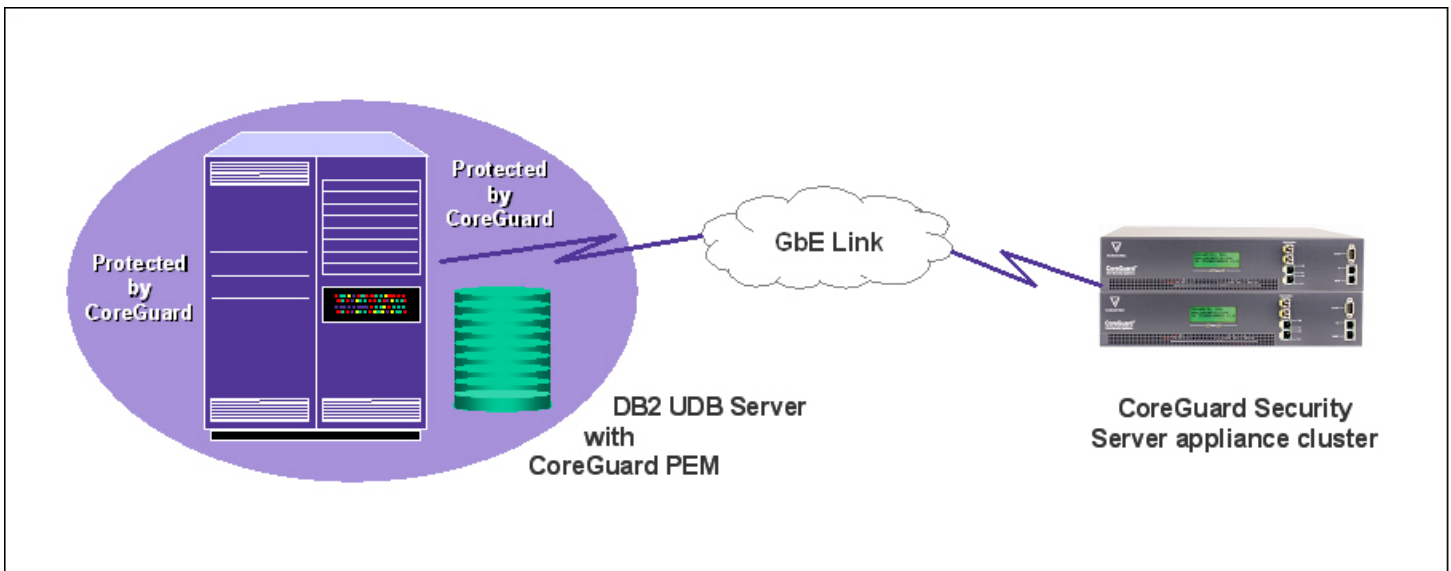


MetaClear data encryption performs cryptographic operations on a page-by-page basis. Thus, MetaClear encrypts and decrypts OLTP or DSS databases ‘on the fly’ in the same way that databases store and retrieve information. Moreover, because MetaClear data encryption operates at the file level, it can be invoked on selected DB2 databases or objects, leaving non-sensitive files ‘in the clear.’ This allows organizations to use strong encryption only on those DB2 objects that it deems sensitive, or only during a time window when the enterprise deems the

selected data files to be the most sensitive, such as information connected to a new product launch, a merger, or company financials prior to public release.

### DB2 Universal Database Server protected by CoreGuard

In the example below, the DB2 UDB database server is running on a host with direct-attached storage (DAS). The CoreGuard PEM intercepts database access calls between the application and the operating system and verifies the context attributes of the call match CoreGuard established data privacy and protection policies. Verified calls are granted permission to access sensitive data in the database. If data visibility is also permitted, the appropriate decryption operation is performed. Calls with context attributes that do not conform to a predefined policy are denied by the PEM, protecting the secure host from attacks by unauthorized users, including root users and unauthorized applications.



*Fig. 2 – CoreGuard protects DB2 host and database files from attack by unauthorized users and applications*

### Achieving Data Privacy and Protection Objectives

By implementing properly configured policies, CoreGuard enables organizations to achieve their data privacy and protection objectives. The table below provides an example of how specific database access combinations can be defined to limit access and visibility to the least privilege required.

Policies can be granular, if so desired, allowing organizations to match access privileges to existing paper security policies for data access by groups or departments. Security administrators are able to assign the proper level of

## Data Privacy Through Database Encryption

privilege and access to each process user, including IT administrators, DBAs, security administrators and the DB2 UDB itself, so that each user can only access the applications and data required to perform his or her duties. Implementing the above policies allows the organization to:

- Enable users to access and view data from DB2 UDB as required, with no changes to business processes or applications.
- Permit system administrators to access DB2 directories so that they can perform their functions without the ability to view sensitive data.
- Deny unauthorized users who hold root privilege the ability to access or view sensitive data.

By using the CoreGuard System, an organization is able to protect sensitive personal data and valuable digital assets, and enhance systems auditability to facilitate compliance with regulatory requirements.

	<b>Applications</b>	<b>File Access Privileges</b>	<b>Accessible File Domains</b>	<b>Data Viewability?</b>
<b>DB2 user group</b>	DB2 UDB	read/write	DB2 data directory	Yes
<b>System admin users</b>	System tools	read/copy	DB2 data directory	No
<b>Security admin users</b>	System tools	read	DB2 audit logs in OS directories	Yes
<b>Root &amp; other unauthorized users</b>	None	None	None	No

### Solution Applications

There are a number of security and compliance applications that can be implemented in a DB2 UDB environment:

- Protect sensitive data stored in a DB2 UDB database from internal and external attacks through the host or directly on the storage systems.
- Safeguard sensitive personal data from unauthorized viewing, providing verifiable and auditable compliance with regulatory requirements for personal data protection.
- Protect data access, administrative and diagnostic log files from unauthorized modification or viewing, which enhances the DB2 UDB database auditability.

- Record detailed CoreGuard audit logs for a 'full-context' record of all access events, providing forensic evidence, and demonstrate compliance with security policies.
- Prevent confidential information viewing by IT administrators, improving the confidentiality of data and mitigating the risk of outsourced data management and storage services.
- Define information protection policies that conform to business-level security policies and roles within the organization.
- Establish 'trusted silos' of information in connection with the consolidation of storage operations.
- Enforce a common operating environment for host servers and contractor workstations.

### Summary

The growing need to make sensitive information more broadly available to individuals within and outside of an organization greatly increases the opportunity for malicious activity and misuse. Perimeter and point security solutions cannot adequately protect sensitive data at rest from these threats, nor do they satisfy today's increasingly stringent audit requirements.

The CoreGuard™ Information Protection System protects sensitive information stored in IBM DB2 UDB. Using CoreGuard, organizations with sensitive data stored in IBM DB2 UDB databases can deploy centrally manageable CoreGuard policies that control who, what, where, when and how data is accessed. CoreGuard protects data using transparent, high-speed encryption of stored information, context-aware access control, advanced host and application integrity protection, and detailed auditing of access events.

\* \* \*

Copyright 2005 Vormetric, Inc. and IBM Corporation. All rights reserved.

Vormetric, Inc. and IBM believe the information in this publication is accurate as of its publication date. The furnishing of this document does not imply giving license to any IBM or Vormetric, Inc. patents.

The information in this document is subject to change without notice and must not be construed as a commitment on the part of IBM or Vormetric, Inc. Neither IBM nor Vormetric, Inc. assumes responsibility for any errors that may appear in this document. No part of this documentation may be reproduced without the express prior written permission of the copyright owners. The software described in this document is furnished under a license and may be used or copied only in accordance with the terms of such a license.

#### Trademark Information

Vormetric, CoreGuard, MetaClear are trademarks or registered trademarks of Vormetric, Inc. in the U.S.A. and certain other countries. IBM, DB2, DB2 Universal Database are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. Other company, product, or service names may be trademarks or service marks of others.

