

Data-Centric Security vs. Database-Level Security

Contrasting Voltage SecureData to solutions such as Oracle Advanced Security Transparent Data Encryption

Introduction

This document provides a high-level overview of Voltage SecureData™ versus database-level security solutions such as Oracle’s Advanced Security Transparent Data Encryption (TDE) and similar “transparent” encryption solutions.

Database-level encryption had its origins in the 1990s and early 2000s in response to very basic risks which largely revolved around the theft of servers, backup tapes and other physical-layer assets. As noted in Verizon’s 2014, Data Breach Investigations Report (DBIR)¹, threats today are far more advanced and dangerous. Attackers and malware are capable of accessing systems directly by exploiting vulnerabilities: attacking at the SQL layer, and mimicking authorized users. In addition, attacks to servers and applications are prominent in the largest breaches, yielding millions of personal data records. Clearly, database-level encryption has not kept pace; breaches have happened despite the best intentions of enterprises that have followed compliance checklists and relied on infrastructure-centric security strategies.

Contrasting Data-Centric Security and Infrastructure Security

Encryption of data and encryption of the system or infrastructure that stores data are two very different security strategies with very different results in managing risk. The former protects data against new threats and streamlines compliance; the latter does not, offering only limited breach risk mitigation and compliance to privacy regulations. Unfortunately, they are often confused as they sound comparable in application scope. It is therefore critical for organizations to understand the differences – especially against the backdrop of the continuous data breaches, insider risks, and new threats to enterprise data. Figure 1 illustrates the inherent security gaps in traditional IT infrastructure solutions.

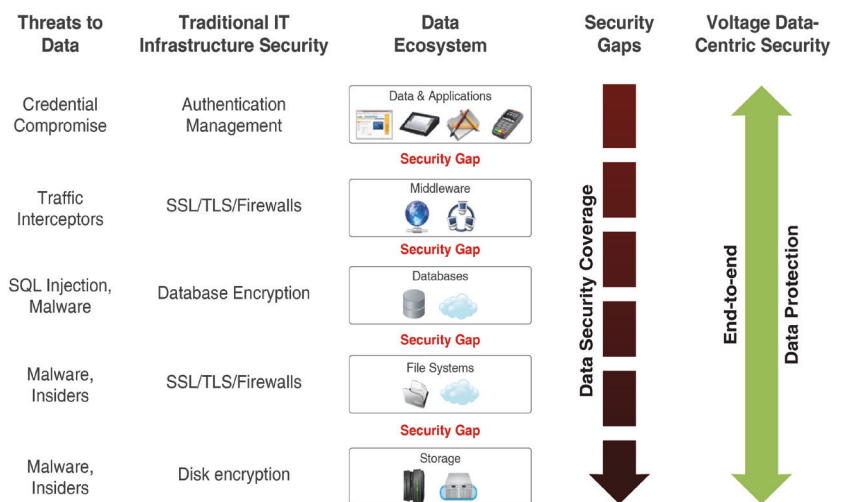


Figure 1. Traditional Infrastructure Security with Security Gaps vs. Data-centric Security

¹ <http://www.verizonenterprise.com/DBIR/4>

While database-level encryption solutions provide rudimentary data-at-rest protection for situations when the database is not in use, protection does not stay with the data when active. With an infrastructure based approach to data security, data is vulnerable to attack as it is read from storage, processed in applications, moved, and consumed by business users. This unnecessarily exposes data to advanced threats and insider risk across the data lifecycle. Malware and advanced threats have compromised data even with database encryption in place at the best prepared enterprises using traditional security.

Verizon’s 2014 DBIR report covers 63,000 individual incidents and highlights the shifting attack and risk profile to data across its lifecycle. The report also notes that among assets attacked by advanced malware and external attackers, “Databases and file servers, both repositories of so much valuable information, are also targeted regularly.” This trend is illustrated in the figure below which shows the significant and continuous rise in malware and external threats to enterprises, with relative decline in physical attacks that database-level encryption was designed for.




Figure 2. Verizon Data Breach Report 2014 – Shifting threat landscape: Hacking and Malware on the rise

Beyond the Database – Big Data and Hadoop Risks

Any organization handling sensitive data in databases today will at some point look at the capabilities Big Data offers for insight and competitive advantage based on existing data assets. Organizations looking to extend data to Hadoop platforms must carefully consider the need to protect data in analytic processes while avoiding data exposure. At the same time, retaining analytic and logic value of the sensitive information is a critical requirement in any solution. Database-level encryption does not address this need, and only introduces unnecessary overhead. As Gartner notes in their March 26, 2014 Research Note titled: “Big Data Needs a Data-Centric Security Focus”², *“CISOs should not treat big data security in isolation but require policies that encompass all data silos to avoid security chaos. New data-centric audit and protection solutions and management approaches are required.”*

A data-centric approach, using industry-standard Format-Preserving Encryption (FPE) or Secure Stateless Tokenization (SST) technology, avoids such complications while providing true end-to-end data protection over the data’s life across any platform or process. If an attacker steals data from a database, data warehouse, Hadoop cluster, or data process where the data itself is secured with Voltage SecureData’s data-centric approach, the attacker gets nothing of value – just useless randomized data. Data breach is thereby neutralized and organizations deploying the Voltage SecureData solution benefit from increased use of data that is safe, compliant and simple.

“CISOs should not treat big data security in isolation, but require policies that encompass all data silos to avoid security chaos. New data-centric audit and protection solutions and management approaches are required.”

*Gartner,
“Big Data Needs a Data-Centric Security Focus”,
Brian Lowans, Earl Perkins,
26 March 2014*

² Gartner “Big Data Needs a Data-Centric Security Focus”, Brian Lowans, Earl Perkins, 26 March 2014

Database Encryption vs. Data Security and Tokenization – PCI DSS

Leading organizations addressing new threats have turned to tokenization and data-centric encryption strategies to remove live data from systems, while still preserving data value during use, storage and transmission. The tangible benefit to organizations implementing these strategies is the reduction in compliance cost overhead by as much as 95%.

The use of tokenization technologies, such as Secure Stateless Tokenization, removes live data from systems to reduce PCI audit and assessment costs, and overall scope. Database-level encryption actually increases scope while not addressing new threats. For example, even with a database-level encryption solution, a breach like Target in 2013 would still be successful, giving the attackers access to live data. In such a case, data is read by the application into memory in the clear and malware steals it from memory, sending it upstream to compromised servers for upload to the attackers' systems.

Documented industry cases indicate that at least two major payment processors have been breached in this fashion, representing millions of credit cards, despite database-level transparent encryption being in place. In both cases, the attacks took place to data as it was read transparently from the database, decrypted transparently, and thus made vulnerable. Transparent to applications also means transparent to malware and insiders as well.

More significantly, database-level encryption does not reduce PCI scope and cost. Complying to the more than 250 controls of PCI DSS to protect data-at-rest using database-level encryption has proven cost prohibitive due to the complexity and many complications associated with key management. In contrast, Voltage's customers have repeatedly reduced their compliance costs by up to 95% which is a huge savings for large and small companies dealing with annual compliance audit costs.

Weaknesses of Database-Level Encryption

The top six critical areas where transparent database encryption creates pain, cost, or risk are as follows:

- **No True Separation of Duties.** Native database encryption provides no separation of duties—data is automatically decrypted when being read out of the database. As a consequence, a DBA or malware will have full access to sensitive data. Data being used in memory in the database is always clear during operation, and may also leak when paged to disk. Keys are also present with data during operations and not separated out. In some cases, including certain database implementations, encryption keys are stored in the database allowing data to be decrypted by anyone with administrator level database access or via malware vectors accessing database memory or wallets that are not protected.
- **Cannot Support Multiple Environments.** Native database encryption provides no capability to protect data in any other environment outside the database, limiting use for data protection across data lifecycles. Large enterprises have data stored in a variety of non-Oracle stores (e.g., MS SQL, DB2, Sybase, Teradata, IBM Mainframe, Hadoop, etc.) which are not addressed by the native Oracle solution.
- **No Application-level Protection.** Native database encryption provides no capability to protect data within applications in use – data is automatically decrypted when read out of the database, meaning that data is completely unencrypted within applications or as data moves and in memory. Major breaches have taken place from malware accessing data in this situation (such as Target in 2013). Insider access to data during use has also resulted in major breaches.
- **Limited and Manual Key Management.** Native database encryption provides no key management beyond the very basic and is often manual – database vendors often require the organization to build its own key management strategy, requiring dedicated resources to manually rotate, backup, restore, and manage keys. In addition, when data spans multiple systems such as batch files, applications, Hadoop systems, reporting tools and backup systems; additional key management point solutions with dedicated resources will be required. Re-keying operations require table moves, backups, and manual intervention and planning.
- **Introduces Operational Complexity.** Native database encryption requires manual password management for encrypting master keys in native database wallets. By only protecting data at rest, other encryption solutions are required to protect data in use, in transit, in BI tools, in sync operations, and in test and development. With so many security gaps, complexity and costs rise dramatically while the organization tries to mitigate each point of compromise and exploitation. This significantly increases operational overhead. Unfortunately, despite increasing resource investments, incomplete data protection will persist.
- **No opportunity for PCI Scope reduction or audit cost reduction.** Database encryption may meet basic compliance needs, but it cannot reduce PCI scope. With keys and data in the same place, all 250+ PCI controls will apply to the database system along with anything that uses it. This includes backup and operational systems, sync systems and operational staff. Tokenization outside the database is the only accepted method to achieve scope reduction, as recognized by PCI QSA's.

Ten Issues with Database-Level Encryption in Today's Threat Ecosystem

The following table provides the Top Ten specific areas of weakness across contemporary security threats and the compliance risks that enterprises face. It also provides contrasting views of the database-level and data-centric data protection strategies to address them.

THREAT OR COMPLIANCE CHALLENGE	DATA PROTECTION METHOD	EFFECTIVENESS	COMPLIANCE AND DATA RISK
1. DBA insider access to tables. Oracle user using TOAD or similar tools to access data directly, including system data.	TDE/Database Encryption Data-centric approach	Minimal. A DBA has full access. Data is decrypted as it is read from the database. Data is in the clear. System may not be protected, resulting in inadvertent exposures. High. DBA will not have live data access. DBA will see only protected data. Key management is external and independent of the database. DBA's can still do their job without the compliance complexity of live data handling.	Insider Risk. Lack of separation of duties. Data protection relies on paper policies. Malware risk by mimicking DBA to steal data. Simplified compliance. Simple separation of duties for compliance including minimized audit scope.
2. Business risk management best practices and compliance controls require independent access between live data and protected data, especially separation of access to data and operations on data.	TDE/Database Encryption Data-centric approach	Minimal. If data is read from the databases, it will be presented to everyone who has access in clear form. TDE does not provide fine grained authorization to distinguish between clear data and encrypted data fields. DBA's grant privileges to users for table/column access which implies full access to live data in the column. In addition, permissions are stored in the database, not external – so permissions and group memberships propagate quickly, increasing inadvertent data exposure and complicating separation of duties requirements. Unique and High value. Separation is clearly defined and managed on a granular basis at a data level. Users can operate on protected data without decryption. Encrypted fields can be exported and used directly in applications and business processes due to the use of Voltage FPE and SST technology.	Granting access to data always implies access to clear data. This increases risk for each and every data access, regardless of the user's actual needs for access to live data. Fields cannot be exported or read in encrypted form, and cannot be used in encrypted form by applications. Compliance simplified. Only a small subset of overall users are granted access. Other users operate on protected data without decryption, reducing risk and exposure.
3. Malware accessing data directly in memory of the database or application. E.g. POS malware, application malware.	TDE/Database Encryption Data-centric approach	None. Data is decrypted into memory as it is read by the application. Live data in memory is stolen by the Trojan software and moved to attackers' staging servers. Unique and High value. Malware will steal protected data, resulting in a non-event. By reducing the exposure of live data to a minimum, risk is greatly reduced. Stolen data is useless to attackers.	Many data breach risks. Many contemporary breaches are due to this malware vector. Proven and effective in reducing risk. Neutralizes breaches when they occur.
4. SQL injection resulting in database compromise and issuance of unauthorized data.	TDE/Database Encryption Data-centric approach	Minimal. If the database presents data by an exploit at the SQL layer, data is decrypted before it is presented to the calling application, resulting in live data exposure. High. Data is persistently protected in the database and beyond. Attackers get nothing of value if data is stolen.	Major data breach risk. Common data breach entry point. Breach risk mitigated. Even when new exploits are used to attack systems.

THREAT OR COMPLIANCE CHALLENGE	DATA PROTECTION METHOD	EFFECTIVENESS	COMPLIANCE AND DATA RISK
5. Exploitation of zero day vulnerability, resulting in authorization bypass.	TDE/Database Encryption Data-centric approach	Minimal. Once database authentication is bypassed, an attacker has complete control over access to data and permissions. High. Authentication and key management are centrally controlled, independent of the database. Database compromise does not yield data or access rights.	Major data breach risk. Requires: a) Continuous patching for every CVE update b) Will not prevent breaches and attacks taking place yielding live data. Breach risk mitigated.
6. Live data in non-production systems, accessed by developers.	TDE/Database Encryption Data-centric approach	None. Does not provide any protection. Vendors often require additional tools to de-identify data, adding cost, and in many cases the methods of protection have no validation or third party assessment to prove security of de-identified data, resulting in increased risk. High. Voltage technology can cover data protection and de-identification in a single platform – for live data, non-production and analytic use cases.	Breach risk, cost impact. Does not protect developer access to production data. Point solutions increase cost and complexity. Breach risk mitigated. Enables outsourcing, test and developer use of data while maintaining data value.
7. Data going to and from cloud applications.	TDE/Database Encryption Data-centric approach	None. Creates security and compliance gap, resulting in exposure. Does not protect data. Requires additional technology solutions to protect data. Keys and data in a database in the cloud are at risk of VM snapshots, hypervisor access, and unintended legal discovery searches. High. Enables protection of data in the enterprise, protection persists with data going to cloud SaaS , PaaS applications, enabling secure cloud adoption. Key management can be completely separated from the data and cloud.	Breach risk, data residency risk, Data compliance risk from lack of control over live data in cloud. Breach risk mitigated. Enables cloud adoption for production, test and analytics while remaining compliant and secure.
8. Data attack outside the database – applications, data arriving to the database, presentation layers.	TDE/Database Encryption Data-centric approach	None. Data is in the clear outside the database and easily compromised by intent or accident on advanced threats. High. Data remains protected in use, in motion, and in storage. A small subset of apps can access live data which can be easily managed. Reduced exposure of data streamlines compliance and reduces risk.	Breach risk, data residency risk, Data compliance risk – in use and in motion and beyond the database. Data remains protected. Risk is mitigated. Standards-recognized data-centric technologies provide strong encryption and tokenization of fields and objects.

THREAT OR COMPLIANCE CHALLENGE	DATA PROTECTION METHOD	EFFECTIVENESS	COMPLIANCE AND DATA RISK
<p>9. Data moving to Hadoop for analytics – Risk to leakage from data-scientists or BI tools.</p>	<p>TDE/Database Encryption</p> <p>Data-centric approach</p>	<p>None. Requires additional third party solutions. Data is exposed on export and in use. Additionally, high risk Hadoop systems which have limited data security controls themselves are exposed.</p> <p>High. End-to-end protection with FPE and SST. Data analytics can operate on protected data, reducing decryption requirements and live data exposure without complex controls in production.</p>	<p>Breach risk, cost impact. Does not protect data outside data at rest in the database. Data is exposed in Hadoop, BI tools, and other applications.</p> <p>Data in Hadoop and BI tools stays secure. In most cases, all operations can take place on protected data, completely mitigating exposure risk.</p>
<p>10. Performance Impact.</p>	<p>TDE/Database Encryption</p> <p>Data-centric approach</p>	<p>On all reads and writes. Data is encrypted and decrypted on every access irrespective of whether the full field is needed or not. Even with optimizations that turn off integrity checking, performance is impacted on all operations. In addition, the data field size expands dramatically to handle meta data.</p> <p>Zero impact for most cases. The Database operates at full speed on protected data. Data does not require decryption. Encryption takes place outside the database at data capture for end-to-end field or object protection. Field sizes do not have to change.</p>	<p>Impact to applications and overhead for all data accesses, even for partial non sensitive fields such as last 4 digits of an SSN. This creates more exposure and risk by design.</p> <p>Minimal impact, data does not need to be decrypted to be used in most cases, reducing risk. Partial non-sensitive fields can be used directly without decrypting the whole field.</p>

Conclusion

Not all encryption techniques are the same. In a new world of advanced threats, and with data moving across systems far beyond the data store or database, new methods of protection are required for compliance and risk reduction. Data-centric security, as provided by Voltage Security, provides new powerful methods to protect data across its full lifecycle. The benefits include dramatic risk reduction and streamlined compliance, while enabling more access to data to grow the business without fear of breach risk across enterprise databases, the data warehouse, big data Hadoop systems and the cloud.

ABOUT VOLTAGE SECURITY

Voltage Security®, Inc. is the leading data protection provider, delivering secure, scalable, and proven data-centric encryption and key management solutions, enabling our customers to effectively combat new and emerging security threats. Leveraging breakthrough encryption technologies, our powerful data protection solutions allow any company to seamlessly secure all types of sensitive corporate and customer information, wherever it resides, while efficiently meeting regulatory compliance and privacy requirements.

For more information, please visit www.voltage.com.

Voltage Security, Inc., Voltage Identity-Based Encryption (IBE), Voltage Format-Preserving Encryption (FPE), Voltage Page-Integrated Encryption (PIE), Voltage Identity-Based Symmetric Encryption, Voltage SecureMail, Voltage SecureMail Mobile Edition, Voltage SecureMail Application Edition, Voltage SecureMail eDiscovery Compliance Tool, Voltage SecureMail Archive Connector, Voltage SecureMail Statement Generator Service, Voltage SecureMail Cloud, Voltage SecureData, Voltage SecureData Enterprise, Voltage SecureData Payments, Voltage Secure Stateless Tokenization (SST), Voltage SecureFile, Voltage SecureData Web, and Voltage Cloud Services are registered trademarks of Voltage Security or are trademarks and service marks of Voltage Security, Inc. All other trademarks are property of their respective owners.