

Database Security for Cloud and Outsourced Environments

Ulf Mattsson, CTO Protegrity

Introduction

One of the biggest concerns about the cloud is the threat of data being stolen. The cloud is a high risk environment that decreases the database administrators' ability to control the flow of sensitive data. Because cloud introduces risk, exposure of database encryption keys becomes particularly vulnerable. Data security in today's business world is a classic Catch-22. We need to protect both data and the business processes that rely on that data, but in order to do so we need to move from a reactive, fear (or compliance) driven mode to a proactive data security plan.

It's not easy to secure any relational database, let alone one as enormous and feature-rich as Oracle. The product's massive and diverse deployments and legacy installations make it virtually impossible to identify and defend against every potential threat. What's more, the product's extensive feature set, while serving some extremely valuable business needs, means more room for trouble. The database's connectivity to Web apps brings open-source and third-party variables into the mix, making the end-user organization even more vulnerable. And end users are inconsistent at best when it comes to patching.

Many organizations will need to reach beyond a point solution for one database brand to address new threats to data across their IT environment. The nature of these breaches call for a different security approach, particularly in outsourced environments. Enterprises are currently on their own in deciding how to apply the principles of Payment Card Industry Data Security Standard (PCI DSS) data protection (e.g. segregation of regulated data) when reducing costs with virtualization or cloud computing, or reducing PCI exposure with tokenization and encryption. Tokenization eliminates keys by replacing sensitive data with random tokens to mitigate the chance that thieves can do anything with the data if they get it. The transparency inherent in random token also reduces remediation costs to applications, databases and other components where sensitive data lives.

Data Breaches

2010 Data Breach Investigations Report

The Verizon Business RISK team, in cooperation with the United States Secret Service (USSS), has been conducting an annual Data Breach Investigations Report. The purpose of the report is to study the common elements and characteristics that can be found in data breaches. In six years, the Verizon Business RISK team and USSS combined dataset now spans 900+ breaches and over 900 million compromised records. As in previous years, the 2010 Report showed that nearly all data was breached from servers and online applications, with 98% of all data breaches coming from servers originating from hacking and malware as the most dominant perpetrators. Financial services, hospitality, and retail comprised the "Big Three" industries, recorded as being 33%, 23%, and 15%, respectively, of all data breaches. Targeting of financial organizations is hardly shocking, as financial records represent the nearest approximation to actual cash for the criminal. An astounding 94% of all compromised records (note: records differ from breaches) in 2009 were attributed to financial services.

Searching to Protect the Business from Endpoint to Endpoint

Financial firms hold large volumes of sensitive consumer data for long periods of time, and because of this, fall under very stringent government regulation requirements that require them to submit remediation validation records if data is found to be vulnerable, as well as regular compliance reports proving that they are adequately securing the data they have access to. Despite being under such stringent compliance standards, 79% of financial firms whose data had been breached failed to meet PCI DSS compliance, the minimum security measure. Thus, organizations have been searching for a solution that protects the business from endpoint to endpoint, while efficiently meeting compliance.

Cloud Opens Up More Hacking Opportunities

In addition to the constantly evolving security threats that must be mitigated, enterprises are quickly adopting cloud computing practices that add a new element to the data security conundrum. According to Gartner forecasts, worldwide revenue from use of cloud services will increase nearly 17% this year to \$68.3 billion and will approach \$150 billion in 2014, a 20.5% compound annual growth rate over the next five years. While its growing popularity is undeniable, the cloud also has serious data security issues. In the cloud, data moves at a faster pace and frees up on-premise network bandwidth, which is what makes it attractive.

Unfortunately, those performing the data breaches recognize the cloud's vulnerabilities and are quickly capitalizing on them. At DEFCON 2010, one of the largest hacker conferences in the world, 100 attendees who have already hacked or have tried to hack the cloud participated in an in-depth survey. 96% of the participants believed that the cloud would open up more hacking opportunities for them. Given its rapid adoption rate, enterprises need a solution that will secure the cloud today and tomorrow.

PCI DSS Compliance

PCI DSS allows Different Ways to Render Data Unreadable

There are three different ways to render data unreadable:

1. Two-way cryptography with associated key management processes
2. One-way transformations including truncation and one-way cryptographic hash functions
3. Index tokens and pads

Two-way encryption of sensitive data is one of the most effective means of preventing information disclosure and the resulting potential for fraud. Cryptographic technology is mature and well proven. The choice of encryption scheme and topology is critical in deploying a secure, effective and reasonable control.

Hash algorithms are one-way functions that turn a message into a fingerprint and are usually no more than a dozen bytes long. Truncation will discard part of the input field. These approaches are used to reduce the cost of securing data fields when data is not needed to do business and will never need the original data back again.

Tokenization is substituting sensitive data with replacement values that retain all the essential characteristics without compromising its security. A token can be thought of as a claim check that an authorized user or system can use to obtain sensitive data such as a credit card number. Using tokenization, all credit card numbers stored in business applications and databases are removed and placed in a highly secure, centralized encryption management server that can be protected and monitored utilizing robust encryption technology.

Cost Effective Approaches

Ponemon Institute summarizes respondents' rating on highly cost effective technologies with respect to achieving PCI DSS compliance goals in the following way. Leading technologies are Firewalls (82%), Anti-virus & anti-malware solutions (74%), Encoding (including Encryption) for data at rest (74%) and Encryption for data in motion (71%). Less cost effective approaches include Data loss prevention systems (43%), Code review (36%), Traffic intelligence systems (32%), Virtual privacy network (26%), Intrusion detection or prevention systems (22%), Database scanning and monitoring (18%) and ID & credentialing system (11%). More details can be found at <http://www.ponemon.org/local/upload/fckjail/generalcontent/18/file/PCI%20DSS%20Survey%20Key%20Findings%20FINAL4.pdf>.

Can I Outsource PCI Compliance?

A few Service Providers can help you by outsourcing PCI Compliance. Some of them are also active members of the PCI Security Standards Council. Physical security and policies and procedures should meet or surpass PCI requirements and be audited annually by a Qualified Security Assessor (QSA). The Payment Card Industry Data Security Standard (PCI

DSS) is an evolving set of security requirements designed for entities that store, process, or transmit cardholder data. These entities must maintain a secure Cardholder Data Environment (CDE). Compliance with PCI DSS is a sound business practice that also serves to keep sensitive data secure. As a business grows and conducts an increasing number of annual credit card transactions, it is subject to increasingly complex compliance requirements. PCI validation requirements are currently organized into four levels which are explained in detail on Visa's web site (http://usa.visa.com/merchants/risk_management/cisp_merchants.html).

Many Service Providers Can Only Deliver Basic Security Controls

Managed services providers offer a range of PCI solutions and services for businesses seeking to outsource compliance. These providers offer a deeper level of expertise and experience in PCI compliance than most businesses typically have in-house. Depending on their clientele, a service provider may have experience working with Level 3 or Level 4 validation requirements only, or they may have a stronger range of experience from serving clients who are subject to Level 1 or Level 2 audits. Expertise and security credentials will also vary among providers. While most service providers can deliver basic security controls, many lack the specific knowledge or advanced expertise required to architect and properly manage a complex PCI DSS compliant solution.

PCI solution offerings will vary depending on the service provider. Some just offer the tools necessary to meet a few aspects of compliance rather than a complete fully managed solution that achieves and maintains PCI compliance. These differences are not always readily apparent, but often come to light upon close examination of the provider's service agreement. If the service provider fails to clearly define which PCI requirements are being met and who is responsible to meet them (the provider and/or the client), a business may falsely believe they are compliant simply because they purchased a PCI toolbox. Although automated tools are one step toward meeting compliance, several PCI DSS requirements demand careful vigilance and analysis to continuously interpret logged data. One example of a service provider for PCI data can be found at <http://www.datapipe.com/solutions-compliance-pci-dss.htm>.

Many Solutions Fail to Meet Compliance Requirements

Achieving compliance in-house requires a significant level of expertise, and maintaining compliance quickly becomes resource intensive. As a result, outsourcing PCI compliance solutions to managed service providers has become a popular business trend in recent years. An increasing number of businesses outside the payment card industry are also deploying PCI DSS solutions across the enterprise to meet a range of industry requirements. However, not all PCI DSS service providers deliver a comprehensive compliance solution. In fact, many solutions fail to meet a majority of compliance requirements. For example, some third-party service providers offer services that meet basic security requirements such as ASV scans and antivirus, but their compliance offerings lack advanced security controls such as Web Application Firewalls (WAF), log management, and two-factor authentication.

A comprehensive solution is more than just hardware or software solutions. Achieving and maintaining PCI DSS compliance requires specialized skills and experience that only a few providers can deliver. The right provider will also offer the resources and expertise necessary to accommodate company growth and scale a PCI DSS solution. When selecting a certified service provider look for one with a deep understanding of the specific requirements of PCI DSS, demonstrated expertise in secure network architecture (including proper network segmentation to reduce the number of system components considered in-scope), and security service design and implementation.

Each Merchant is Liable for Security

It is essential that the service provider deliver a transparent service agreement which clearly defines both the client's and the provider's responsibilities. This ensures that all aspects of the standard are addressed and there are no gaps in responsibility. The organization or merchant that is storing, processing, or transmitting the cardholder data is ultimately liable for any gaps in compliance, which makes it imperative to clearly delineate who is responsible for each aspect of meeting the PCI DSS standard. Companies often underestimate the time, resources, and efforts required to continuously and rigorously maintain compliance in-house. Security systems require significant capital investments in hardware and software and are costly to implement, maintain, and monitor. Outsourcing to the right provider enables businesses to achieve and maintain compliance while controlling costs.

Other companies turn to outsourcing because their current business volume has outgrown existing compliance resources. As illustrated in the chart above, merchants are subject to increasingly complex compliance requirements as the annual number of credit card transactions processed grows. Meeting those requirements becomes extremely resource intensive which makes outsourcing an attractive option.

Oracle Security and Cloud

Oracle Investing in New Security Technologies

Oracle has been working hard and investing in new technologies to address these concerns, particularly in light of the bad press about its lack of responsiveness to the database's vulnerabilities. The result: a security makeover of sorts. The vendor now offers an extensive range of optional security tools that reaches far beyond the basic security items in the Oracle Database Vault. The company's Transparent Database Encryption, for instance, automatically encrypts information as it's stored in the database; it's designed so you can retrofit it easily into your existing Oracle setup. Likewise, integration of the Secerno firewall technology is designed to detect database attacks in real time; Oracle and Secerno together can now block SQL injection attacks and other types of queries that appear dubious, including remote code exploits. This may not be a cost effective approach for PCI DSS and you'll have to develop your own rules, policies and reports for each of the Oracle security tools.

The Default Oracle Configuration is Insecure

Not every Oracle customer uses version 11 of the database. In fact, a majority of customers use older versions. Large percentages of Oracle customers still use versions 8 and 9. The database is functionally stable, so customers are not in a hurry to make the investment in upgrading. But these versions were designed and built before most people had ever heard of buffer overflow attacks or remote exploits. Many of the known security threats have been addressed with patches -- provided they have been back-ported -- but these older versions lack some of the advancements in password management, encryption, separation of DBA roles, and auditing. Similarly, legacy applications -- control systems, homegrown applications, mainframe connectors, SAP R3, and so on -- that use older versions of Oracle don't have security built in. They have interlocking dependencies between the application and databases, and rely on external security services to detect and protect against threats.

The sheer number of features and options provided by Oracle creates a larger threat surface, with far more targets of opportunity for attackers. Oracle comes standard with many features that a lot of businesses rarely use. And just about every Oracle package has been compromised at one time or another. Because Oracle serves so many different use cases, there is no such thing as a secure default configuration. The default Oracle configuration is insecure, and users must take the time to remove features they don't need, and to verify that user, platform, and application security measures are in place before the database goes into production.

Oracle's Road to Cloud

There are many ways to implement cloud computing today. Some are more successful than others. The ones that have been successful have applied a discipline to the methods utilized in their creation. Oracle has found that by employing a framework and standard architectural development process focusing on the needs of the business has achieved the greatest success. Following are the components of that framework and development process.

Oracle has seen cloud computing provide real cost savings and agility to its customers in all cloud service models. These companies all had a vision for their business that used cloud computing as an enabler. What is your company's vision for cloud computing? This question forms the foundation for a cloud computing strategy. Whether trying to create a private cloud to better serve internal customers or building a public cloud for an external customer base, there are several key considerations that Oracle has linked to successful cloud initiatives:

1. Understand the IT service portfolios, service-level requirements and service costs before building a private cloud service.
2. Develop a separate strategic plan for all services under consideration, as well as an analysis against external service offerings.

3. Build a private cloud service only after developing a complete business case analysis for doing so— it's all about return on investment, in terms of cost and business value.
4. Evaluate and constantly benchmark the solution against external cloud service offerings, and ensure that flexibility is designed in at the onset.

Oracle and Amazon Cloud

Amazon and Oracle have collaborated to offer their customers options and convenience when deploying enterprise applications on the cloud. Customers can not only build enterprise-grade solutions hosted by Amazon Web Services (AWS) using database and middleware software by Oracle, but they will also be able to launch entire enterprise software stacks from Oracle on EC2. In both cases, customers can benefit from the scalability, reliability, and cost-effectiveness of deploying on Amazon's cloud. Use Oracle Database 11g and Oracle Enterprise Linux to build enterprise-grade solutions in the cloud, leveraging the virtually unlimited compute power and storage of Amazon Web Services (AWS).

Hosting Oracle-based solutions in AWS enables you to use proven database and middleware offering within a proven cloud computing platform, providing greater reliability and performance than hosting solutions on your own hardware. Together, Amazon and Oracle provide businesses with a scalable, reliable, and cost-effective business application platform. Oracle customers can now license Oracle Database 11g, Oracle Fusion Middleware, and Oracle Enterprise Manager to run in the AWS cloud computing environment. Oracle customers can also use their existing software licenses on Amazon EC2 with no additional license fees. And for on-premise Oracle installations, AWS offers a dependable and secure off-site backup location that integrates seamlessly with Oracle RMAN tools.

New Data Security to Address the New Threats

Some organizations will need to reach beyond a point solution for one database brand to address new threats to data across their IT environment. The nature of these breaches call for a different security approach, particularly in outsourced environments. We need to understand how to deal with the threats conceptually before we jump into the more complex technical and operational issues that can confuse your choices.

Transparent Encryption won't protect sensitive content in the database if someone has access to it through legitimate credentials, but it will protect the information on storage and in archives, and provides a significant advantage as it is deployed independent of your business applications. If you need to protect things like credit card numbers where you need to restrict even an administrator's ability to see them, this option isn't for you. If you are only worried about lost media, stolen files, a compromised host platform, or insecure storage, then Transparent Encryption is a good option. By not having to muck around with the internal database structures and application logic, it often provides huge savings in time and investment over more involved techniques.

Encryption Approaches for Protecting Databases

Many newer database versions provide native database object encryption, external file/folder encryption or media encryption. Which option to choose depends on your performance requirements, threat model, existing architecture, and security requirements? Unless you have a high-performance system that exceeds the capabilities of file/folder encryption, we recommend you look there first. If you are managing heterogeneous databases, you will likely look at a third party product over native encryption. In both cases, it's very important to use external key management and not allow access by any local accounts.

Transparent Database Encryption for Cloud

The term Transparent Encryption is used by many vendors to describe the capability to encrypt data stored in the database without modification to the applications using that database. We've also added "External" to distinguish from external encryption at the file or media level. If you have a database then you already have access controls that protect that data from unwanted viewing through database communications. The database itself screens queries or applications to make sure that only appropriate users or groups are permitted to examine and use data.

The threat we want to address here is protecting data from physical loss or theft (including some forms of virtual theft) through means that are outside the scope of access controls. Keep in mind that even though the data is "in" a database, that

database maintains permanent records on disk drives, with data being archived to many different types of low cost, long term storage. There are many ways for data to be accessed without credentials being supplied at all. These are cases where the database engine is by-passed altogether -- for example, examination of data on backup tapes, disks, offline redo log files, transaction logs, or any other place data resides on storage media, as pointed out at <http://securosis.com/tag/database>.

Key Management Issues within Cloud

Existing cloud service providers may provide basic encryption key schemes to secure cloud based application development and services, or they may leave all such protective measures up to their customers. While cloud service providers are progressing towards supporting robust key management schemes, more work is needed to overcome barriers to adoption. Emerging standards should solve this problem in the near future, but work is still in progress, including <http://csrc.nist.gov/groups/SNS/cloud-computing> . There are several key management issues and challenges within Cloud Computing. Key stores must themselves be protected, just as any other sensitive data. They must be protected in storage, in transit, and in backup. Improper key storage could lead to the compromise of all encrypted data. Access to key stores must be limited to the entities that specifically need the individual keys. There should also be policies governing the key stores, which use separation of roles to help control access; an entity that uses a given key should not be the entity that stores that key. Loss of keys inevitably means loss of the data that those keys protect. While this is an effective way to destroy data, accidental loss of keys protecting missioncritical data would be devastating to a business, so secure backup and recovery solutions must be implemented. There are a number of standards and guidelines applicable to key management in the cloud. The OASIS Key Management Interoperability Protocol (KMIP) is an emerging standard for interoperable key management in the cloud. The IEEE 1619.3 standards cover storage encryption and key management, especially as they pertain to storage IaaS.

Data Tokenization

Encryption vs. Tokenization

End-to-end encryption can encrypt sensitive data fields throughout most of its lifecycle, from capture to disposal, providing the strongest protection of individual data fields. Therefore, end-to-end encryption, and its next of kin - tokenization - are very practical approaches to protect data between specific parts of a solution that are in high risk areas. While there is no silver bullet to the data security and compliance woes of large enterprise organizations, all eyes are on tokenization right now. Tokenization is different from encryption in that it is based on randomness, not on a mathematical formula. Encryption also requires compliance with key management, key rotation, selection of algorithm, etc. that are moot with tokens. Next generation tokenization offers a faster, more secure solution and uses less computing power than encryption. Also according to PCI DSS, encrypted data must be re-encrypted every year, while tokenized data can be left for a life time.

Security and the Data Lifecycle

The combined approaches of tokenization and encryption can be used to protect the whole data lifecycle in an enterprise. It also provides high quality production level data in test environments, virtualized servers and outsourced environments. In the development lifecycle there is a need to be able to perform high quality test scenarios on production quality test data by reversing the data hiding process. Key data fields that can be used to identify an individual or corporation need to be cleansed to depersonalize the information. In the early stages of implementation, cleansed data needs to be easily restored (for downstream systems and feeding systems). This requires two-way processing. The restoration process should be limited to situations for which there is no alternative to using production data. Authorization to use this process must be limited and controlled. In some situations, business rules must be maintained during any cleansing operation (addresses for processing, dates of birth for age processing, names for gender distinction). There should also be the ability to set parameters, or to select or identify fields to be scrambled, based on a combination of business rules.

There Are Two Forms of Tokenization Available

First generation tokenization solutions are based on the simple concept of a large and dynamic table of token/credit card pairs. While this is an obvious and reasonable approach, it has its disadvantages and issues with respect to performance, scalability, and availability. The core obstacle with the traditional tokenization approach is that the token lookup table is so large and dynamic that it's hard to manage. Next generation tokenization, on the other hand, addresses all of these issues through scalability with multiple, parallel instances, dramatically increased performance, availability, centralized or distributed deployment, elimination of token collisions, and support of PCI, PHI and PII data. When next generation tokenization is applied strategically to enterprise applications, confidential data management and PCI audit costs are reduced and the risk of a security breach is minimized. Security is immediately strengthened by the decrease of potential targets for would-be attackers, because authentic primary account numbers (PAN) is only required at authorization and settlement. Studies have shown annual audits average \$225K per year for larger payment card acceptors, and next generation tokenization reduces audit costs dramatically by eliminating the need for encryption keys.

Tokenizing Sensitive Data Can Be Cost Effective

Tokenizing sensitive data including PAN and social security numbers can be a cost effective end-to-end solution that meets PCI compliance. Unfortunately, the PCI Security Standards Council (SSC) has not yet developed standards for tokenization, nor will they include tokenization in PCI DSS 2.0.

In attempt to fill this void, Visa published its "Best Practices for Tokenization" Version 1.0 on July 14 at http://usa.visa.com/download/merchants/tokenization_best_practices.pdf. Be careful, because this draft implies a "one-size-fits-all" architectural solution open enough for botched implementations including encryption pretending to be tokenization and home-grown tokenization that lack security requirements, where random-based tokenization is the only true end-to-end solution.

Next Generation Tokenization – Examples

Here are a few examples of this next generation data tokenization. A major retailer recently migrated from the traditional tokenization approach to a new Next Generation tokenization approach to meet the needs of operational performance. This retailer's Tokenization solution now enables them to perform 200,000 tokenizations per second on a single small commodity server which now enables them to meet their current and planned future needs for using tokenization on card holder data and PII/PHI data across the enterprise. In another example, a major retailer recently migrated from an encryption solution to this next generation tokenization approach to reduce cost. In a final example, a major US organization, encouraged by the scalability and small footprint of their new generation tokenization approach, are planning to put a hardened tokenization server at several hundred of their sites to meet their performance and availability requirements. The New Distributed Tokenization is fully distributed solution which does not require replication between servers and completely eliminates the severe issue of token value collisions.

Implementing a Tokenization Solution

Should a Company Build Their Own Tokenization Solution?

Developing all the capabilities to build a solution in-house can present significant challenges. To be implemented effectively, all applications that currently house payment data must be integrated with the centralized tokenization server. Developing either of these interfaces would require a great deal of expertise to ensure performance and availability. Writing an application that is capable of issuing and managing tokens in heterogeneous environments that can support multiple field length requirements can be complex and challenging. Furthermore, ongoing support of this application could be time consuming and difficult. Allocating a dedicated resource to this large undertaking and covering for responsibilities could present logistical, tactical, and budgetary challenges. For many organizations, locating the in-house expertise to develop such complex capabilities as key management, token management, policy controls, and heterogeneous application integration can be very difficult. Writing code that interfaces with multiple applications, while minimizing the performance impact on those applications, presents an array of challenges. The overhead of maintaining and enhancing a security product of this complexity can ultimately represent a huge resource investment and a distraction from an organization's core focus and expertise.

Security administrators looking to gain the benefits of centralization and tokenization, without having to develop and support their own tokenization server, should look at vendors that offer off-the-shelf solutions.

Should I Outsource Tokenization?

Standard tokenization is an available feature direct from gateway payment providers. But utilizing standard tokenization still requires that credit cards must first be handled and stored on the merchant's infrastructure prior to being tokenized. Newer integration options and Next Generation Tokenization handles the token exchange at the edge, before any data has entered the merchant's infrastructure. This newer approach plays a key role in helping to reduce and potentially remove ecommerce and online transaction activity from PCI scope. Other tokenization options for off-loading PCI transactions include re-directing the traffic to externally hosted and processed web pages, or inserting third party fields into existing ecommerce workflows. Both methods require merchants to outsource this business critical transaction to uncontrolled, non-customized, and un-reliable infrastructures. With newer approaches in integration of Tokenization, customers do not require workflow changes, externally hosted sites, or form fields.

Can I Outsource Risk and Liability?

Some companies do not want to outsource secure handling of data since they cannot outsource risk and liability. Larger organizations may not be willing to move the risk from its environment into a potentially less secure hosted environment. Further, enterprises need to maintain certain information about transactions at the point of sales (POS), as well as on higher levels. In most retail systems, there are a plurality of applications that use or store card data, from the POS to the data warehouse, as well as sales audit, loss prevention, and finance. At the same time, the system needs to be adequately protected from attacks from data thieves. Merchants who gather card data via Web commerce, call centers and other channels, should ensure that the product or service they use can tokenize data through all channels. Not all offerings in the market work well or cost-effectively in a multi-channel environment, particularly if the token service is outsourced. Merchants need to ensure that their requirements reflect current and near-future channel needs. Another concern is that tokenization is new and unproven can pose an additional risk relative to mature encryption solutions. A risk management analysis will reveal whether the cost of deploying tokenization in house is worth the benefits. An outsourcing environment must be carefully reviewed from a security point and provide a reliable service to each globally connected endpoint. Many merchants continue to object to having anyone keep their card data other than themselves. Often, these are leading merchants that have made significant investments in data security and simply do not believe that any other company has more motivation (or better technology) than they do to protect their data.

How to Develop and Deploy a Risk-adjusted Data Security Plan

Protecting data according to risk enables organizations to determine their most significant security exposures, target their budgets towards addressing the most critical issues, strengthen their security and compliance profile, and achieve the right balance between business needs and security demands. Other issues that risk-adjusted security addresses are the unnecessary expenses, availability problems and system performance lags that result when data is over-protected. And cloud-based technologies, mobile devices and the distributed enterprise require a risk-mitigation approach to security, focused on securing mission critical data, rather than the now-unachievable 'protect all the data at all costs' model of years past.

1: Know Your Data

Begin by determining the risk profile of all relevant data collected and stored by the enterprise, and then classify that data according to its designated risk level. Data that is resalable for a profit -- typically financial, personally identifiable and confidential information -- is high risk data and requires the most rigorous protection; other data protection levels should be determined according to the value of the information to your organization and the anticipated cost of its exposure -- would your business be impacted? Would it be difficult to manage media coverage and public response to the breach? There are several models that a business can use to classify data. Larger enterprises will likely want to rely on policy-driven automated tools. Smaller businesses can use the simplest model: assign a numeric value for each class of data; high risk = 5, low risk = 1.

2: Find Your Data

Data flows through a company, into and out of numerous applications and systems. A complete understanding of the high risk data flow is essential to the risk-adjusted process. You can't protect data if you don't know where it is, and assigned risk levels will change depending on how data is being collected, used and stored. High risk data residing in places where many people have access is obviously data that needs the strongest possible protection.

Locate all of the places that data resides including applications, databases, files, and all the systems that connect these destinations such as data transfers across internal and external networks, etc. and determine where the highest-risks reside and who has or can gain access to data (see "Understand your Enemy" below).

Other areas to examine for data stores include your outsourcing partnerships as well as data that is being used for nonproduction purposes such as third-party marketing analysis or in test and engineering environments. It's not uncommon for organizations to invest in protecting production systems and data centers yet have live data sitting unprotected on the systems of application developers and other outsourced parties. If live production data is being used in a less controlled environment there has to be attention paid to regulatory compliance and security threats. Here, too, data de-identification technologies like Format-Controlling Encryption and tokenization can help.

Step 3: Understand Your Enemy

The next step is conducting an end-to-end risk analysis on the high risk data flow to identify the highest risk areas in the enterprise ecosystem and the points where data might be exposed to unauthorized users.

Currently web services, databases and data-in-transit are at high risk. The type of asset compromised most frequently is online data. Exploiting programming code vulnerabilities, subverting authorized user credentials and malware targeting the application layer and data (rather than the operating system) are the attack methods that are being utilized most frequently. These vectors change so keep an eye on security news sites to stay abreast of current threats.

Most data breaches are caused by external sources but breaches attributed to insiders, though fewer in number, typically have more impact than those caused by outsiders. Nearly three-quarters of the breaches examined in the Verizon Report were instigated by external sources. Unauthorized access via default credentials (usually third-party remote access) and SQL injection (against web applications) were the top types of hacking, access to a network was often followed by malware being planted on the system.

Step 4: Choose Your Defenses

Look for multi-tasking solutions that protect data according to its risk classification levels, supports business processes, and is able to be change with the environment so that you can easily add new defenses for future threats and integrate it with other systems as necessary.

High risk data is best secured using end-to-end encryption or tokenization of individual data fields. Tokenization removes sensitive data from the information flow at the earliest possible point in the process, replacing it with a token that acts as an alias for the protected data. By associating original data with an alias, high-risk data can systematically be removed and protected from malicious hackers over its lifecycle under a fully auditable and controllable process. This practical protection method is perfectly suited for securing high risk data like payment card information and social security numbers.

Newer solutions provide targeted protection for data in use and doesn't interfere with business processes. For example, Data Format Controlling Encryption retains the original format, on a character-by-character basis, of encrypted data, putting an end to the data re-formatting and database schema changes required by other encryption techniques. It's especially well-suited to protect data that's being used for testing or development in a less-controlled environment. Partial encryption can then be applied to provide the ability to encrypt selected parts of a sensitive data field based on policy rules. Policy-Based Masking provides the ability to mask selected parts of a sensitive data field. Implemented at the database level rather than application level, policy-based Data Masking provides a consistent level of security across the enterprise without interfering with business operations and greatly simplifies data security management chores.

Step 5: Deployment

Risk-Adjusted data protection enables enterprises to stage their security roll-out. Focus your initial efforts on hardening the areas that handle critical data and are a high-risk target for attacks. Then continue to work your way down the risk-prioritized list, securing less critical data and systems with appropriate levels of protection.

Security is an ongoing process not a series of events. The level of protection required by data may change according to how it is being collected, transmitted, used and stored. Reevaluate risk levels annually and on an as-needed basis if business processes change.

Step 6: Crunch the Numbers

Risk-adjusted data security plans are cost effective. Among the typical benefits of a risk-adjusted plan is the elimination of the all too common and costly triage security model which is ineffective whether you're triaging based on compliance needs or the security threat of the moment. Replacing triage with a well thought-out logical plan that takes into account long-range costs and benefits enables enterprises to target their budgets toward addressing the most critical issues. By switching focus to a holistic view rather than the all too common security silo methodology, an enterprise will also naturally move away from deploying a series of point solutions at each protection point, which results in redundant costs, invariably leaves holes in the process, and introduces complexity that will ultimately cause significant and costly rework. Additionally, an understanding of where data resides usually results in a project to reduce the number of places where sensitive data is stored. Once the number of protection points has been reduced, a project to encrypt the remaining sensitive data with a comprehensive data protection solution provides the best protection while also giving the business the flexibility it needs.

Conclusion

Not too long ago, many security experts believed that the best way to defend data was to apply the strongest possible technological protections to all of the data, all of the time. While that plan may work perfectly in theory, in the real world of business this model creates unacceptable costs, performance and availability problems. What works from both IT and management standpoints? Risk-adjusted data security. Protecting data according to risk enables organizations to determine their most significant security exposures, target their budgets towards addressing the most critical issues, strengthen their security and compliance profile, and achieve the right balance between business needs and security demands.

A holistic solution for data security should be based on centralized data security management that protects sensitive information throughout the entire flow of data across the enterprise, from acquisition to deletion. While no technology can guarantee 100% security, tokenization and encryption are proven to dramatically reduce the risk of credit card data loss and identity theft. Next generation tokenization in particular has the potential to help businesses protect sensitive data in the cloud in a much more efficient and scalable manner, allowing them to lower the costs associated with compliance in ways never before imagined. By switching focus to a holistic view rather than the all too common security silo methodology, an enterprise will also naturally move away from deploying a series of point solutions at each protection point, which results in redundant costs, invariably leaves holes in the process, and introduces complexity that will ultimately cause significant and costly rework. Additionally, an understanding of where data resides usually results in a project to reduce the number of places where sensitive data is stored.

Many organizations will need to reach beyond a point solution for one database brand to address new threats to data across their IT environment. The nature of these breaches call for a different security approach, particularly in outsourced environments. Enterprises are currently on their own in deciding how to apply the principles of PCI data protection (e.g. segregation of regulated data) when reducing costs with virtualization or cloud computing, or reducing PCI exposure with tokenization and encryption. Tokenization eliminates keys by replacing sensitive data with random tokens to mitigate the chance that thieves can do anything with the data if they get it.