



The Security Division of EMC

White paper

Best Practices for Preventing Enterprise Data Loss



'Enterprise data loss' cost businesses nearly \$105 billion last year

More than 150 million records have been breached since January 2005, according to Privacy Rights Clearinghouse, a non-profit consumer information and advocacy organization.

Contents

Best Practices for Preventing Enterprise Data Loss.	page 1
Best Practice #1: Understand what data is most sensitive to the business.	page 2
Best Practice #2: Know exactly where the most sensitive data resides.	page 2
Best Practice #3: Understand the origin and nature of your risks.	page 3
Best Practice #4: Select the appropriate controls based on policy, risk, and where sensitive data resides.	page 4
Best Practice #5: Manage security centrally.	page 6
Best Practice #6: Audit security to constantly improve.	page 7
How Do I Get Started?	page 8
Conclusion	page 9

The numbers are staggering. More than 159 million electronic records have been breached since January 2005, according to Privacy Rights Clearinghouse (www.privacyrights.org), a non-profit consumer information and advocacy organization.

Not only is the number of data thefts and losses due to security breaches continuing to grow at an alarming rate, the resulting monetary impact of these losses is also skyrocketing. So-called 'enterprise data loss' cost businesses nearly \$105 billion last year, according to U.S. government estimates. Insider data breaches alone cost businesses an average of \$3.4 million per company each year, according to the Ponemon Institute. And industry analyst firm Gartner Group estimates that the cost of recovery can reach \$150 per breached record – a number that does not factor in the costs of regaining customer loyalty due to brand damage, potential fines, and legal representation.

This reality leaves companies – both large and small – in a vicious cycle. To optimize business processes, improve customer service, and enhance partner relationships, companies collect ever-greater volumes of data. And as this need grows, data becomes more broadly distributed across an increasing number of information systems throughout the enterprise. The result? The risk of compromising sensitive data rises because more users – both inside and outside the company – have access to sensitive data than ever before.

Adding to the problem, traditional security infrastructures were primarily designed to protect against external threats. Yet today, the more imminent threats to security today come from inside a company. The increasingly lucrative black market for information used for identity theft, financial fraud, and other criminal purposes has shifted the security focus to insiders with broad access to sensitive data – they know where the systems are, how they interact with each other, and what data resides on which systems.

This mismatch between current threats and traditional security infrastructures is leading to more data breaches, increased regulation, and higher operational costs. In turn, these slow down core business processes – the very opposite result intended by the information explosion.

So, how can companies protect themselves from a data loss catastrophe?

This paper outlines six best practices for corporations seeking to prevent enterprise data loss and thereby protect revenue, limit customer attrition, and meet government regulatory requirements.

Best Practices for Preventing Enterprise Data Loss.

Based on more than 25 years of experience in the security industry, RSA has developed a strong understanding of the best practices that help prevent enterprise data loss. By following these best practices, companies can not only improve their ability to secure sensitive customer data, but also better protect revenue, ensure customer loyalty, and meet government regulations.

These best practices are:

1. Understand what data is most sensitive to the business
2. Know where this sensitive data resides
3. Understand the risk model
4. Select the appropriate controls based on policy, risk, and where sensitive data resides
5. Manage security centrally
6. Audit security to constantly improve

Let's look at each of these best practices in more detail.

Best Practice #1: Understand what data is most sensitive to the business.

Not all data is of equal importance from a security perspective. The first step in preventing enterprise data loss is to determine which data is most sensitive – or at highest risk – to your business. Then, you can prioritize your efforts and define appropriate policies. But how do you know which data is most sensitive to your business?

To answer the question, you need to understand your business structure, examine the various departments and lines of business across your organization, and identify both the regulatory and non-regulatory security drivers for each department. For example, your finance department might need to comply with Sarbanes-Oxley and Gramm-Leach-Bliley Acts as well as SAS 70, while your retail operations group needs to focus on SB 1386 and PCI. Meanwhile, your international operations must comply with the European Union's Directive on Data Protection as well as a range of country-specific regulations, such as the Japan Privacy Act, Canada's PIPEDA, and the Australia Privacy Act, to name just a few.

Once the regulatory and corporate compliance universe is understood, you can prioritize your data by grouping information into various 'classes'. For example, you might create three classes of information from the most restricted and sensitive (e.g., data relating to the company's unannounced financial results) to the least sensitive (e.g., data pertaining to vendor shipping rates).

The next step is to determine the data categories, elements, and owners for each class of information. For California SB 1386 compliance, for example, non-public personal information – government identification numbers and citizenship status, for example – are critical pieces of the compliance puzzle. You might classify this information as 'restricted'. Then, you should determine which elements of the information are most critical and which department or business unit within the company owns this data.

Finally, after you have classified your data, you must then define the policies – the rules for 'appropriate handling' of the data – including which employees and applications are authorized to access this data and how, when, and from where they are allowed to access it. For example, you might allow all employees in R&D to access the information pertaining to the company's products, but only certain employees to view the data about new, unreleased products – and only during specific hours and from within the corporate firewall.

Best Practice #2: Know exactly where the most sensitive data resides.

At first glance, the answer to the question, "Where does my company's most sensitive data reside?" seems to be obvious. You would probably answer, "In databases, of course!" But databases are really just the tip of the iceberg, especially with today's mobile, highly collaborative environments. If data is stored in a database, then it is also stored on a disk, which is likely backed up to other disks or tape media. Additionally, your data is probably accessed through a variety of applications and from a wide array of devices, transformed on desktops, laptops, and wireless hand-helds, e-mailed to other users, and then stored on yet other file servers or collaboration portals.

The truth is that the answer to the question is not so obvious – yet it is critical to preventing enterprise data loss. Most companies, however, do not take the time to conduct thorough data discovery, leaving them with three choices, none of which are viable: One, they secure all their data, which is only possible with an unlimited budget. Two, they secure none of their data, which is only acceptable if the company is willing to accept all the risk that accompanies this strategy. Or three, and most commonly, they secure some of their data in a haphazard manner – lulling themselves into a false sense that data is secure and ignoring significant risks, which is simply dangerous.

To prevent enterprise data loss and strike a balance between cost and risk, you must go beyond simply determining which databases house your critical data. Rather, you should undertake a complete data discovery process, which requires you to answer some basic questions about your infrastructure, including:

- Do you have sensitive data in databases? If so, in which database tables? In which columns or fields?
- Do you have sensitive data in file shares? If so, in which folders? In which files?
- Do you have high-risk data on laptops? If so, on whose laptops?

Next, you will also need to answer data type and usage questions, such as:

- Is your intellectual property unwittingly exposed through custom-built applications?
- Are your unannounced company financial reports illicitly finding their way onto laptops, PDAs, and USB drives?
- Is your customers' credit card information being transferred from databases to insecure file servers so that users can create spreadsheets and reports?
- Are back-up tapes containing consumer information guaranteed to arrive at their final location without interruption or tampering?

Through the data discovery process, your company can create a map of its critical and sensitive data, which serves as a foundation for your security policy and control strategy. And, to be effective, you must embrace data discovery as a continuous process, not a one-time event, as neither your organization's data nor your use of it is static.

Best Practice #3: Understand the origin and nature of your risks.

In addition to knowing where your important data resides and how it is being used, you need to understand your risks. How could your data be compromised or stolen? By whom? And how much risk would your company assume with exposure of this data?

The answers to these questions can be found both inside and outside of your organization. Lapses in business processes and innocent mistakes on the part of users are actually more common than a malicious attack from outside your organization. In fact, a recent study by Ponemon Institute shows that insider threats – from negligent or malicious employees, partners, and contractors and from process breakdowns – was the number one cause of data breaches in 2006. While the severity of each of these threats will vary by organization, defining that severity is essential to determining risk. The answers to all of these questions are critical in the development of your risk model.

Some of the more common risks have had their share of headlines over the last 24 months, including:

- **Lost or stolen media** – Back-up tapes or disk drives are frequently lost or stolen from data centers or en route to remote archival location. These tapes contain confidential customer or employee information and can be sold for top dollar on the black market, leading to criminal activity, such as identity theft.
- **Privileged user breach** – Privileged users, such as database or file server administrators, are found selling sensitive corporate data on the black market. For example, one of the nation's leading financial services companies discovered that a senior database administrator – responsible for managing data access privileges across the company – attempted to sell the personal information of more than two million customers on the black market for a substantial sum.
- **Unintentional distribution** – Sensitive data is sent out via public e-mail, exposed on public portals, or otherwise distributed to unauthorized users.


Data discovery must be embraced as a continuous process – not a one-time event.

Figure 1

“To what extent are the following trends related to data security challenges a concern to you and your business in upcoming years?”

Source: Encryption Key Management study of 199 IT decision makers responsible for security, conducted by Forrester Consulting and commissioned by RSA Security, April, 2007

Base: All respondents, n=199



- **Application hack** – A hacker (often an insider) breaches application authorization controls to access highly sensitive data through accounting applications, human resources applications, point-of-sale systems, and other critical business applications. For example, application developers often take production data to test new applications, violating numerous regulations and creating enormous risk.
- **Physical theft or loss** – A laptop or other portable device containing sensitive data is stolen or lost – along with sensitive customer data and corporate intellectual property stored on them.

According to a recent study conducted by Forrester Consulting commissioned by RSA, greater demands for mobile employee data access, collaboration, and partner data exchange present the biggest challenges to data security today and in the foreseeable future (see Figure 1).

Creating a risk model that takes into account all the potential ways your data might be compromised or stolen provides the context you need to implement an appropriate control strategy that outlines both the types of control mechanisms (i.e., how to secure the data) as well as the points of control (i.e., where to secure the data).

Best Practice #4: Select the appropriate controls based on policy, risk, and where sensitive data resides.

Once you understand your policies, where your sensitive data resides, and the risks at those locations in your infrastructure, you can develop an appropriate control strategy. That strategy will likely include both processes and technology.

The physical control strategy is comprised of two components: the control mechanisms (i.e., the types of controls), and control points (i.e., where in the infrastructure they are placed; at the storage, database, file server, application, network, or end point). A comprehensive control strategy will include a combination of controls from all three categories described below, implemented at various layers in the IT stack:

- **Access Controls** control both authentication (i.e., is the user who he or she claims to be?) and authorization (i.e., what can the user do once he or she gains access?). A wide range of products are covered in this category, including web access management, two-factor authentication, and knowledge-based authentication.

- **Data Controls** control the data itself. Data controls include products and technologies such as encryption, data loss prevention (DLP) and information rights management (IRM).
- **Audit Controls** provide the feedback mechanisms to ensure the policies and controls are in fact working as they should. Often called security information and event management (SIEM), audit control products provide the means to prove compliance as well as refine policies and controls.

Over the last several years, more companies are focusing on implementing data controls, especially encryption solutions and DLP systems (which are also sometimes referred to as "information leak protection" systems), due to the increasing number of data breaches and growing regulatory scrutiny of data privacy and integrity issues (see Figure 2).


Why? Because both encryption and DLP systems are highly effective in collaborative environments where data is mobile, shared, and transformed. These two types of data controls exemplify the notion of 'self-defending data'. That is, they enable your data to defend itself.

For example, even if an individual is able to circumvent your access controls and steal encrypted data, the data is useless to them. Likewise, if highly sensitive data subject to privacy regulations is transformed and e-mailed out of an organization, a DLP system can react and protect that data. In today's highly mobile collaborative environments, these controls are indispensable.

Figure 2

“What are your top data security initiatives over the next 12 months?”

Source: Encryption Key Management study of 199 IT decision makers responsible for security, conducted by Forrester Consulting and commissioned by RSA Security, April, 2007
Base: All respondents, n=199



Data controls include products and technologies such as encryption, data loss prevention and information rights management

Where to Encrypt?

A variety of different encryption solutions are available today, enabling the encryption of data at virtually any level in the IT stack: storage, database and file server, application, end-point, and network layers. Encryption addresses different risks at each layer. For example, organizations commonly encrypt the storage media (e.g., tapes) and end-points (e.g., laptop computers). While encrypting at these levels is non-invasive to applications that run above them, it only addresses the physical theft or loss of the media itself – a small spectrum of the risk model.

An increasing number of organizations are implementing encryption at both the database and file server levels because it is still relatively non-invasive to applications and provides protection against a broader range of threats, including privileged user breaches (e.g., a DBA compromising data) and unintentional distribution (e.g., a developer using production data to test a new application).

Encrypting data at the application layer gives organizations much broader protection, securing data throughout its lifecycle as it moves from the user and end point to the application, then to the database, and finally to the underlying storage and backup infrastructure. However, this added level of protection does not come without a cost. Application encryption requires that organizations add calls to encryption systems from within the application code. While advances in these solutions have greatly simplified this process, it is still invasive. Applications that process highly sensitive or highly regulated data, such as point-of-sale systems, are prime candidates for this type of control.

The right answer to the question, "Where should I encrypt my data?" will almost always be "At a combination of these control points." And the exact mix will depend on the nature of data and infrastructure processing that data.

Best Practice #5: Manage security centrally.

More than any other factor, the management of control mechanisms has a greater impact on both the effectiveness of controls and their total cost of ownership.

Organizations often make the mistake of managing each control mechanism separately, which results in policy misalignment, high management costs, and a lack of business process continuity.

To avoid these problems, companies must manage their security control mechanisms – including both policies and keys – centrally. Centralizing the administration of security policies ensures that control points consistently enforce security rules and makes proactive monitoring of activity that could result in a security violation easier to automate. In addition, centralization helps ensure that users consistently follow appropriate usage rules for sensitive data to avoid unintentional leakage.

The second piece of centralized security management involves encryption keys. With centralized key management, encryption controls can be effectively and consistently implemented across all control mechanisms, protecting the organization from data breaches due to human error, lost keys, or incompatible and conflicting encryption policies. For example, while all encryption products come with some form of management console, these consoles often lack higher-level, policy-based capabilities. And because each encryption product comes with a separate console, it can be nearly impossible to align the configuration and operation of these systems with the business' underlying security policies. Furthermore, different employees, often non-security staff, are tasked with managing the various encryption products, increasing management costs. Finally, when encrypted data needs to be shared between applications, groups, or infrastructures, the lack of centralized management for key sharing means either the data needs to be decrypted before sending it from one point to another and re-encrypted on the other end, causing increased overhead and vulnerabilities. Without centralized management of both security policies and encryption keys, processes can be irrevocably broken, leading to business disruption.

"Without centralized management of both security policies and encryption keys, processes can be irrevocably broken, leading to business disruption."

Companies that do not manage security centrally encounter three significant problems:

1. **Misaligned policies.** Managing these mechanisms individually makes it difficult – if not impossible – to ensure that the organization's security policies are uniformly and consistently implemented across all control mechanisms.
2. **High management costs.** The cost of ownership is multiplied several-fold when managing these mechanisms individually, because many employees (often non-security personnel) must manage a myriad of management consoles from different vendors. Typically, these consoles are rudimentary tools packaged with the control mechanism and lack any significant policy-based management capabilities, thereby forcing users to interact with systems at a highly technical, low-level manner. This incurs significant management, training, and other overhead expenses.
3. **Lack of business process continuity.** Most business processes rely on the sharing of data across applications, users, infrastructure, and sometimes organizations. The security mechanisms in place should facilitate, not hinder, these processes, but the lack of centralized management often creates barriers between these components and can even break processes completely.

Best Practice #6: Audit security to constantly improve.

As with any corporate process, a security program should have a feedback mechanism that enables the organization to assess its compliance with policy, and provide feedback on the effectiveness of data controls. Business is not static – neither are the security mechanisms that protect it. You need real-time tracking and correlation of security events in order to respond quickly to change.

SIEM systems enable you to analyze and report on security logs and real-time events throughout your enterprise. To enable proper auditing of your data security infrastructure, you need an SIEM system that automatically collects, manages and analyzes the event logs produced by each of the security systems, networking devices, operating systems, applications, and storage platforms deployed throughout your enterprise. These logs monitor your systems and keep a record of security events, information access, and user activities both in real time and for forensic analysis.

By correlating events in your data control systems – such as encryption and loss prevention – in real time, you can quickly respond to incidents as they occur, remediating any potential losses. Such proactive log management is the foundation for a comprehensive auditing strategy. An SIEM system enables you to regularly review your security infrastructure for:

- Incident investigation and forensics,
- Incident response and remediation,
- Compliance to regulations and standards,
- Evidence for legal cases and
- Auditing and enforcing data security policy.

By establishing auditing best practices and implementing an effective SIEM system, you can reduce the cost and increase the efficiency of compliance, risk management, and forensics. Equally important, auditing provides an opportunity for continuous improvement. Security should always be viewed as a process rather than an event.

How Do I Get Started?

Now that you have an understanding of industry best practices to prevent enterprise data loss, how do you begin implementing them? The process may seem daunting, but RSA, the leader in data security solutions, can help.

RSA provides a comprehensive solution to help companies through all stages of the data security process, from risk assessment to control strategy to implementation. The RSA solution combines services and products designed around the best practice framework described in this paper. Addressing security as a process can greatly increase the effectiveness and efficiency of your security program.

RSA's products and services help you address the best practices outlined in this paper:

Best Practices #1 and 2: Determine what data is most sensitive and where it resides.

Once you have classified your information, you can pinpoint all instances of the data across the network (e.g., in file systems, on desktops, and on PDAs) and when crossing network boundaries (e.g. when sent in an email).

Data discovery and classification is the first step toward securing your data, but the fact that sensitive data exists in different forms (e.g., database records, email messages, and unstructured files) and different contexts (e.g., at-rest in data center storage, in-motion through the network, and in-use on laptops, mobile devices, and portable storage) complicates the process.

Tablus Content Sentinel™ enables you to perform enterprise-wide classification and discovery so you can rapidly identify where the sensitive data resides in your infrastructure and identify your data risk areas. It helps you manage your data according to specific governance and compliance requirements based on corporate, government, and industry regulations. RSA Professional Services are also available to assist you with defining your data classification policy and using Tablus tools effectively.

Best Practice #3: Understand your risks.

RSA Professional Services provides risk assessment services to help you accurately assess the relevant risks and threats to your information, and identify the relevant policies, procedures, and controls to address those risks.

Best Practice #4: Define your control strategy.

RSA Professional Services can also help you develop an appropriate control strategy to address your risks. In addition, RSA provides a range of products to enforce your security policy and control access, usage, and distribution of your data throughout your infrastructure:

- **Tablus Content Alarm NW™ and Tablus Content Alarm DT™** detect sensitive data in motion across your network and in use on your laptops and desktops to help remediate incidents of violated policies. The product automatically monitors and blocks transmissions containing sensitive content to minimize required intervention and maintain compliance with regulations and corporate policy. For example, by automatically routing emails containing sensitive content to an encryption server to secure messages and attachments on-the-fly in accordance with content protection policies, Tablus Content Alarm NW and Tablus Content Alarm DT enforcement of your enterprise data policies.
- **RSA® Database Security Manager** manages encryption across multiple types of databases, including various versions of Oracle, SQL Server, Sybase, and DB/2, from a single, centralized console. It can selectively encrypt content down the column/field level. RSA Database Security Manager enforces separation of duties between database administrators and your security personnel to prevent those users from performing intentional or unintentional operations that compromise data. By selectively encrypting only the data objects you care about, RSA Database Security Manager ensures transparent operation and the uninterrupted flow of business processes and information.
- **RSA® File Security Manager** manages encryption on both Windows and Linux file servers, transparently enforcing security for both the users and administrators of those file servers, RSA File Security Manager enables encryption in the file share, folder, and file levels, protecting against unauthorized use and distribution of sensitive data.
- **Storage Encryption offerings from Cisco and EMC integrated with RSA® Key Manager** enable seamless encryption for all storage media, including backup tapes and disks.

One of these solutions is Cisco Storage Media Encryption (SME), a heterogeneous, high-performance encryption solution integrated directly into the storage network fabric that works across tapes, virtual tapes, and disks.

RSA is also working with EMC to integrate our key management and encryption technologies directly into EMC storage product offerings to provide more transparent storage encryption solutions for customers.

Best Practice #5: Manage security centrally.

RSA Key Manager provides centralized provisioning and key lifecycle management for encryption keys and other security objects throughout the enterprise. These reduce the complexity in the deployment and ongoing management of encryption controls. RSA Key Manager also can be easily integrated into specialized applications, such as retail point-of-sale (POS) terminals and financial accounting systems. Used in conjunction with RSA Database Security Manager, RSA File Security Manager, and Cisco SME, RSA Key Manager is a robust, easy-to-use policy-driven solution for enterprise-wide encryption management.

Tablus Content Sentinel, Tablus Content Alarm NW, and Tablus Content Alarm DT also provide centralized management for data discovery, classification, reporting, auditing, and leak prevention capabilities.

Best Practice #6: Audit security to constantly improve.

The **RSA enVision** platform gives you the power to gather and use log data to understand your security, compliance, or operational status in real time or over any period of time. It provides efficient collection, analysis, and management of all data from any IP device in computing environments of any size, without filtering and without the need to deploy agents. All data is correlated in real time to produce dashboards and reports on the health and effectiveness of your security infrastructure.

Conclusion

Protecting your company from enterprise data loss is a strategic imperative. The risks – both monetary and image-related – are too great to ignore and simply hope your company does not become a victim. But laying out a data security strategy and then implementing it is no easy feat.

RSA, the leader in information-centric security solutions, has developed the set of six (6) best practices contained in this paper based on significant, long-term experience with thousands of companies. By following these best practices, you can not only improve your ability to secure sensitive customer data, but also to protect revenue, limit customer attrition, and meet government regulations.

Call RSA today to protect your company from potentially crippling enterprise data loss.

The risks – both monetary and image-related – are too great to ignore.

About RSA

RSA, The Security Division of EMC, is the expert in information-centric security, enabling the protection of information throughout its lifecycle. RSA enables customers to cost-effectively secure critical information assets and online identities wherever they live and at every step of the way, and manage security information and events to ease the burden of compliance.

RSA offers industry-leading solutions in identity assurance & access control, encryption & key management, compliance & security information management and fraud protection. These solutions bring trust to millions of user identities, the transactions that they perform, and the data that is generated. For more information, please visit www.RSA.com and www.EMC.com.

RSA and RSA Security are registered trademarks or trademarks of RSA Security Inc. in the United States and/or other countries. EMC is a registered trademark of EMC Corporation. All other products or services mentioned are trademarks of their respective owners. ©2007 RSA Security Inc. All rights reserved.

PEDL WP 0907



RSA Security Inc.
RSA Security Ireland Limited
www.rsa.com

The Security Division of EMC