Best Practices for a Successful DLP Deployment

TITUS White Paper

www.re-soft.com info@re-soft.com 203 972 8462

At TITUS we work to help businesses better manage and secure valuable corporate information. Our focus is on building policy management solutions that make it easier for IT administrators to protect and manage corporate correspondence including email and documents.

Table of Contents

1.0 | Overview 4

2.0 | The Real Challenge: Inadvertent Data Loss 5

3.0 | Best Practices for User Driven DLP 7

Step 1: Identify the Top Data Loss Scenarios 7

Step 2: Create Simple Policy Rules 7

Step 3: Alert Users From Within the Application 8

Step 4: Allow Users to Remediate Policy Violations 9

Step 5: Focus On User Education 9

4.0 | User Driven DLP Examples 11

5.0 | The TITUS Advantage 13

1.0 | Overview

The one constant in business today is change: companies merge, employees join and leave the organization, new regulations and laws are passed. These changes put the organization's most valuable data at risk, presenting an enormous challenge for CIOs, CSOs, and IT departments who are responsible for protecting the information.

The reality is that there is only so much that the IT Department can do on its own. Data protection needs to be a joint effort, starting with the content owners and information workers who handle the organization's PII, financial data, and intellectual property on a daily basis.

In this white paper, we'll look at best practices for making users the first line of defense against data loss. By following these best practices, organizations can make users responsible for their own data, significantly reduce data loss incidents, and free up IT departments for more targeted security enforcement and education.

2.0 | The Real Challenge: Inadvertent Data Loss

Data leaks caused by insiders are a growing concern for CIOs and CSOs. The vast majority of these leaks are accidental, caused by users who are simply trying to do their work in a fast-paced, rapidly changing environment.

Most organizations place high expectations on the IT department to prevent these leaks. Yet there is only so much that IT can achieve on their own, due to the size and complexity of the problem. Two areas present particular challenges: 1) An ever-changing workforce with different skills and attitudes toward security; and 2) Huge volumes of data with diverse data protection needs.

A Diverse Workforce

End users have different skills and attitudes toward security. Some users are well-informed and conscientious, but may assume incorrectly that others will treat their data with the same care. Others may be stressed or distracted, causing them to make inadvertent mistakes. And still others may intentionally violate corporate policy, due to laziness, irresponsibility, or a genuine belief that their business objective outweighs the risks.

For the typical information worker, it can be difficult to see how their work fits into the organization's wider business strategy, especially in the areas of compliance and data protection. Information workers are experts in their own area of speciality, but may not understand how a seemingly innocent mistake could cost the organizations millions in fines and lost business.

To add to this complexity, workforces are constantly changing. The organization may acquire another company, or reorganize their business units. Employees will change roles, and new employees will join the company. Offices may open in other countries, and work may be outsourced to global service providers.

Diverse Data Protection Needs

Organizations also face a difficult challenge in protecting the huge amount of data generated by information workers each day. In just one year, an organization of 1,500 users will generate almost 70 million emails[1] - not to mention other data transferred through instant messaging, social media, and portable media.

Much of this data is subject to regulations. Organizations that fail to protect this data face significant fines, loss of customer trust, sanctions, or possible criminal charges. Examples of regulated data include:

- Personally identifiable information regulated under privacy protection legislation such as S.B. 1386 in California, MASS 201 in Massachusetts, PIPEDA in Canada, and DPA in the UK

[1] Osterman Research, Inc. *The Case for Outbound Content Management.* April 2010. <http://titus.com/resources/index.php?resourceid=101&tabno=4>.

- Personal health information regulated under HIPAA and HITECH
- Financial data regulated under SOX, GLBA, BASEL II, PCI DSS and SEC disclosure rules

Other types of data are not regulated, but are important and sensitive because they represent the organization's intellectual property and proprietary information. When disclosed to unauthorized recipients, the consequences can be lower revenues, loss of competitive advantage, and damage to the company's reputation. This type of data tends to be widely distributed throughout the organization, and may be difficult to identify. Examples include:

- Company strategy documents
- Project plans and designs
- Revenue and pricing information
- Sales forecasts
- Financial, legal, and competitive analysis

The diversity of the organization's workforce and data makes it challenging for IT departments to protect sensitive information while enabling business productivity. It is a problem that cannot be solved by IT alone. As we will see in the next section, truly protecting sensitive data requires a joint effort, starting with the content owners and information workers who handle the organization's PII, financial data, and intellectual property on a daily basis. This joint effort also needs to include the IT department who implement policies and enforce control over data, as well as the compliance/risk officers who understand the legal and business ramifications of the regulations that are pertinent to the business.

3.0 | Best Practices for User Driven DLP

To address the challenge of insider security threats, many organizations are deploying data loss prevention (DLP) solutions. Unfortunately, these deployments can quickly become multi-year projects as IT administrators attempt to translate the business process and policies into automated rules for every data loss scenario. IT administrators quickly find that it is impossible to accurately identify every type of sensitive document, just as it is impossible to predict the behavior of every type of user.

As well, business executives typically have significant concerns about putting in place systems that interrupt user workflow or business agility. As a result, the majority of businesses that have deployed DLP systems run them in "watch-mode", where they are monitoring for leaks, and never realize the true benefits of actually preventing data loss.

For a successful DLP deployment, it is it is crucial to make users the first line of defense. A user driven DLP strategy will start with the user, alerting them to potential data loss at the desktop while they work, and giving them the tools to fix policy violations before they happen. This user driven approach makes users responsible for their own data, and educates them on corporate policy while they work. It also significantly reduces data loss incidents, and frees up IT departments for more targeted security enforcement and education.

Let's consider the 5 essential steps to implementing a user driven data loss prevention strategy:

Step 1: Identify the Top Data Loss Scenarios

Many DLP deployments fail because they try to address too many problems at once. It is important to start small, focusing on the data loss scenarios that have the highest impact to the organization. These may be data loss incidents that are relatively minor but occur several times per day. Or they may be unusual, high impact breaches that cost the organizations millions of dollars in fines or lost business.

Here are some scenarios to consider:

- **Are there common, innocent errors that lead to data loss?** For example, do users frequently send email to the wrong recipients when their email program auto-completes email addresses? Do users fail to notice confidential information in long email threads?
- **Do you need to control the flow of internal email?** For example, an investment bank needs to prevent analysts from leaking corporate inside information to brokers who advise clients about buying shares. Or an executive needs to restrict certain email discussions to other executives.
- **Are there certain types of documents that present a greater risk?** For example, a retail organization may be concerned about spreadsheets leaving the organization with customer PII. Or an investment bank may have confidential information stored in pitch book presentations, which could lead to insider trading violations if the information was disclosed externally.

Step 2: Create Simple Policy Rules

Once the most important data loss scenarios have been identified, the next step is to define rules to prevent data breaches. Many organizations are tempted to implement complicated content scanners to

automatically determine the sensitivity of an email or document. However, in many cases, simple policy rules may cover many data loss scenarios and lead to faster deployment. Here are some questions to consider:

- **Do your sensitive documents have certain characteristics that you could use for a policy rule?** For example, do your confidential documents include a header or footer to indicate their sensitivity? Do they follow a naming convention, such as using a project code in the name of a project plan document? Do they contain any unique document metadata?
- **Do your users have certain attributes that you can use to control access?** For example, can you use Microsoft Active Directory to determine a user's role, group, or department? Are there external email addresses, such as competitors, that are known to be high risk?
- **Is certain content known to be sensitive?** For example, are there patterns you can search for to prevent PII disclosure, such as Social Security Numbers or credit card numbers? Are there key phrases you could search for in content, such as "Internal Only" in an email or document?

Step 3: Alert Users From Within the Application

Now that you have a list of common data loss scenarios and policy rules, the next step is to focus on the user as part of the solution. By installing DLP technology on the user's desktop, you can alert users to policy violations before a data breach occurs.

The best place to activate data loss protection is within the application where the user violated corporate policy. Integrating DLP into the user's application enables the organization to:

- **Prevent security incidents before they happen.** Users are informed immediately of policy violations, before they actually cause a data breach.
- **Integrate alerts into the user's workflow.** Because alerts appear within the application, users are guaranteed to see the warnings so that they can take action. This is in contrast to less effective methods, such as generating messages at the operating system level where users are less likely to notice them, or quarantining suspicious content for IT administrators to sift through.
- **Provide targeted security education that helps prevent policy violations in the future.** Users receive alerts that educate them on why the policy violation occurred, how to fix the problem, and how to prevent violations in the future.
- **Provide an effective way to deal with false positives.** Organizations can configure their alerts to simply warn the user rather than blocking the user from completing a task such as sending an email. This is useful in situations where content has been incorrectly identified as sensitive (i.e. a false positive). Users can also be forced to enter a justification if they override the warning, which is then logged and reported.

Alerts should be based on the data's content and how the information is being used. For example with email, the DLP policy should check the recipients, message content, and any attached documents. The alert should be triggered if certain criteria are met, such as an investment analyst sending IPO information to external recipients.

Step 4: Allow Users to Remediate Policy Violations

A key component of a successful DLP strategy is to give users the tools to fix problems themselves. Most security violations are inadvertent, such as users accidentally attaching the wrong document to an email, selecting the wrong recipient, or overlooking a confidential section in a long email thread. Users should be trusted to correct these sorts of policy violations, without having to involve IT and interrupt business productivity.

There are several ways in which users can self-remediate policy violations:

1. **Edit content:** A DLP solution should highlight security violations and give the user the opportunity to edit the content from within the original application.
2. **Redact content:** Uses should be able to redact (i.e. black it out) content that violates corporate policy. This is useful in situations where users need to preserve the original structure of the content while removing the sensitive sections.
3. **Remove recipients:** Users should be alerted to unauthorized recipients, and given the opportunity to remove these recipients before sending the information. For example, a user may not realize that an email distribution list contains external contractors. Or the email auto-complete feature may fill in the wrong recipient name. These unauthorized recipients should be identified in any alerts that are presented to the user.
4. **Classify and mark:** Some organizations have a data classification policy that requires users to identify confidential or sensitive information. Users in these organizations should be alerted when their content appears to be sensitive, and given the tools to classify the data so that it can be protected. Users may also wish to add protective markings, such as headers and footers, to visually indicate the information's sensitivity.
5. **Cancel the action:** One remediation option is to cancel the action that was triggering the policy violation. For example, a well-timed alert may convince a user not to try to forward documents to their home address to work at home.
6. **Override the alert:** In some cases, a warning may actually be a false positive, or there may be legitimate business reason for going against the standard corporate policy. When this happens, users should be given the ability to override the alert so that business productivity is not impacted. Users can be forced to provide a justification for overriding certain alerts, and the justification can be logged and reported for auditing and educational purposes.

An important benefit of self-remediation is that it puts accountability on the user. Users are responsible for their data, and are expected to correct any policy violations themselves. This helps to create a culture where security is seen as everyone's responsibility, not just that of the IT department.

Step 5: Focus On User Education

In many cases, data breaches happen because the user was not aware of the impact of their actions. Users may be unfamiliar with corporate policy, or they may not understand the larger compliance picture in regulated environments such as financial organizations. Sometimes it's simply a matter of underestimating the consequences of going against policy – such as sending confidential documents to a home email address to complete a project after regular business hours.

Education is one of the most effective ways to prevent data loss. A user driven DLP strategy provides two ways to educate users on security policy: 1) Through real-time targeted alerts within the application; 2) Through the monitoring of user activity, which helps to identify areas where more education is required.

Real-time Targeted Alerts

As discussed in the previous sections, a user driven DLP strategy will provide the user with policy alerts before they cause a data breach. This kind of education is very powerful because it occurs in real-time as the user is violating corporate policy. This causes users to stop and consider the consequences of their actions. It also allows the organization to identify exactly where the problem exists, and to give users the tools to remediate the problem themselves.

Organizations should expect to see a drop in policy violations as users become more security conscious. Users will become more familiar with corporate policy, and will also become more aware that their actions might be monitored. This will drive a change in user behavior, contributing to an overall reduction in security incidents.

User Monitoring and Education

A user driven DLP strategy will include the option to monitor user activity through logging and reporting. Monitoring allows the IT department to easily identify areas of concern, such as employees generating high numbers of alerts and overrides. This allows IT to focus on more targeted security enforcement and education, such as following up with the employee or their manager, or providing further education and training to specific groups within the organization.

Monitoring also helps to fine-tune the organization's security policies. If many users are overriding a specific alert, it may be that the policy is too restrictive, or the rule needs to be fine-tuned to avoid false positives. Monitoring user activity helps the IT department assess the effectiveness of the security rules and identify areas where the rules and alerts can be improved over time to balance data protection with business workflow.

4.0 | User Driven DLP Examples

By following the best practices outlined in this white paper, users will become the first line of defense against data loss. In partnership with the IT department and compliance/risk officers, users can dramatically reduce the risk of insider data leaks – without compromising business productivity.

Here are some examples of a user driven DLP solution in action:

Prevent Careless Disclosures That Result In Lost Business

An analyst has put together an analysis of an upcoming initial public offering that his firm is underwriting. He intends to send the analysis to some internal recipients, but accidentally selects a personal distribution list that includes several outside investors. To send the email would violate securities regulations, and would result in his company being removed as an underwriter – a loss of $50 million in revenue.

Before the email leaves the desktop, the DLP solution checks both the content and the recipients. The user is alerted to the fact that the email contains IPO information and is going to external recipients. After getting over the initial shock of having almost cost the company $50 million, he removes the external recipients and sends the email to the internal recipients. The alert is logged, as is the user's action.

Remind Honest Employees about Corporate Policy

An employee needs to finish a presentation for work, and decides to forward the document to her home email address so that she can work on it at night. The presentation, while not highly sensitive, does contain some financial analysis and strategy recommendations.

Before the email leaves the desktop, the user receives an alert that she is responsible for any documents that are sent to external addresses. This prompt reminds her about the company's security policies, and she decides not to send the email after all.

Protect Sensitive Data in Documents

An analyst has been working on valuation estimates for an upcoming merger. He needs to get feedback from another analyst, so he attaches his spreadsheet to an email, and attempts to send it. However, the analyst has inadvertently selected a broker with a similar name to the analyst. Sending this email would violate insider trading rules.

Before the email leaves the desktop, the analyst receives an alert that the recipient (the broker) is not authorized to receive the email. The DLP solution has checked the content against the recipients, and determined that the recipient does not have the correct Microsoft Active Directory attributes. The analyst realizes his mistake, removes the broker from the recipient list, and sends the email to the correct recipient.

Protect Personally Identifiable Information

An employee receives an email inquiring about a particular customer account. The email contains a long thread of replies and somewhere in the thread it happens to contain the customer's credit card number and expiry date.

The employee needs to forward the email outside the organization to an external supplier in order to get more information. When he clicks Send on the email, the DLP solution immediately alerts him that the email contains personal identifiable information which violates corporate policy.

The employee is prevented from sending the email until the policy violation is rectified. The DLP solution provides him with the option to either remove the credit card information or redact it. He chooses to redact the information, causing the credit number and expiry to be deleted and blacked out, which then allows the email to be sent and the policy warning to be logged.

Manage False Positives without Business Interruption

A financial analyst needs to send some information to her client for a meeting later that day. Before she can send the information, the DLP solution alerts her that the information looks like it could be a confidential pitch book. However, the information is all publicly available, and there is no risk of insider trading. The analyst overrides the warning, is asked to provide a justification, and sends the information. The action is logged along with the justification, and the IT department follows up by investigating how they could fine-tune the DLP rule to avoid false positives in the future.