Data loss prevention as a privacy-enhancing technology

Received (in revised form): 28th April, 2020

William Stallings
Independent consultant, USA

Dr William Stallings is an independent consultant and author of *Information Privacy and Privacy by Design* (Pearson, 2020). Bill has written numerous textbooks on cybersecurity, cryptography and computer networking (williamstallings.com), including *Effective Cybersecurity: A Guide to Using Best Practices and Standards* (Pearson, 2019) and *Cryptography and Network Security, Principles and Practice, Eighth Edition* (Pearson, 2020). He has 12 times received the award for the Best Computer Science and Engineering Textbook of the Year from the Textbook and Academic Authors Association. Bill created and maintains the Computer Science Student Resources site at ComputerScienceStudent.com. This site provides documents and links on a variety of subjects of general interest to computer science students (and professionals). He is a member of the editorial board of *Cryptologia*, a scholarly journal devoted to all aspects of cryptology. Bill holds a PhD from the Massachusetts Institute of Technology (MIT) in Computer Science.

PO Box 2405, Brewster, MA 02631, USA
E-mail: wllmst@me.com

Abstract Data loss prevention (DLP) is a mature technology deployed by many enterprises to support information security requirements. Key characteristics of DLP also make it a powerful privacy-enhancing technology that can satisfy a wide range of information-privacy requirements. In essence, DLP is a set of integrated technologies that detect sensitive data and prevent unauthorised use, release or delivery to specific destinations or recipients, as well as their storage at prohibited locations. DLP works in real time to identify personally identifiable information and react to privacy risks based on data content and the dynamic context of the information environment. Thus, DLP provides technical enforcement of terms and conditions, or policies more generally, to prevent privacy leaks. This paper provides an overview of the main features and elements of DLP.

KEYWORDS: content analysis, context analysis, data at rest, data in motion, data in use, data leakage prevention, data loss prevention

INTRODUCTION

Data loss is the intentional or unintentional release of information to an untrusted environment. Data loss prevention (DLP), also known as data leakage prevention, refers to the use of integrated technologies that detect sensitive data and prevent their unauthorised use, release or delivery to specific destinations or recipients, as well as their storage at prohibited locations. DLP functions by applying a combination of contextual and content analysis methods and enforcing centrally managed data-protection policies. DLP controls are based on policy and include classifying sensitive data, discovering that data across an enterprise, enforcing controls, and reporting and auditing to ensure policy compliance. Sensitive information that is at risk of leakage or is actually leaked often includes shared

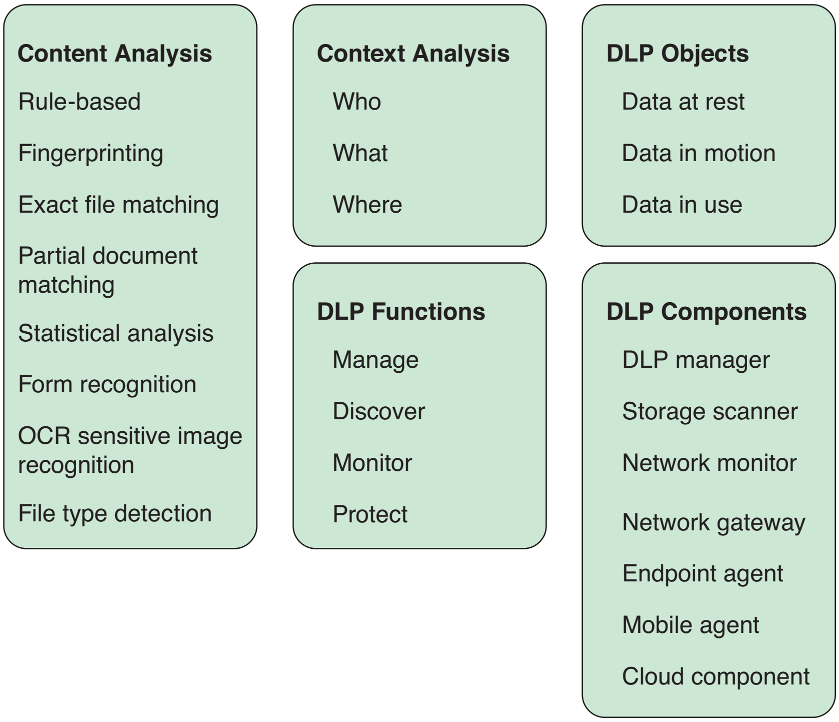

Figure 1: Elements of data loss prevention

Notes: DLP, data loss prevention; OCR, optical character recognition.

and unencrypted content, such as word processing documents, presentation files and spreadsheets, that could leave an organisation via many different points or channels (eg via e-mail, instant messaging, internet browsing or on portable storage devices).

DLP was developed to address security concerns, but it is also a vital technology for information privacy.[1] DLP addresses the protection of sensitive data, and this includes PII.

Figure 1 depicts the key elements of a DLP capability. These are examined in this paper. The paper begins with a model showing DLP in the context of information security/privacy governance. This is followed by a look at specific functions and components of DLP. Next, a brief introduction to analytic techniques used to support DLP is presented. Finally, the paper presents two examples of DLP in practice.

DLP MODEL

Figure 2 depicts a conceptual model that illustrates all the elements that contribute to a holistic view of data loss prevention. There are three levels to the model. Data governance ensures that DLP aligns with corporate goals and requirements and drives the development of DLP controls. An organisation needs to develop a classification scheme for the specific types of personal data it holds in order to customise DLP controls for its needs. In this model, part of governance is the identification of what sensitive data the organisation stores and processes and where in the IT (information technology) architecture (servers, data centres, cloud, workstations) data are held.

Data governance

Central to data governance, from both a security and privacy perspective, is to assess the risk associated with security or privacy

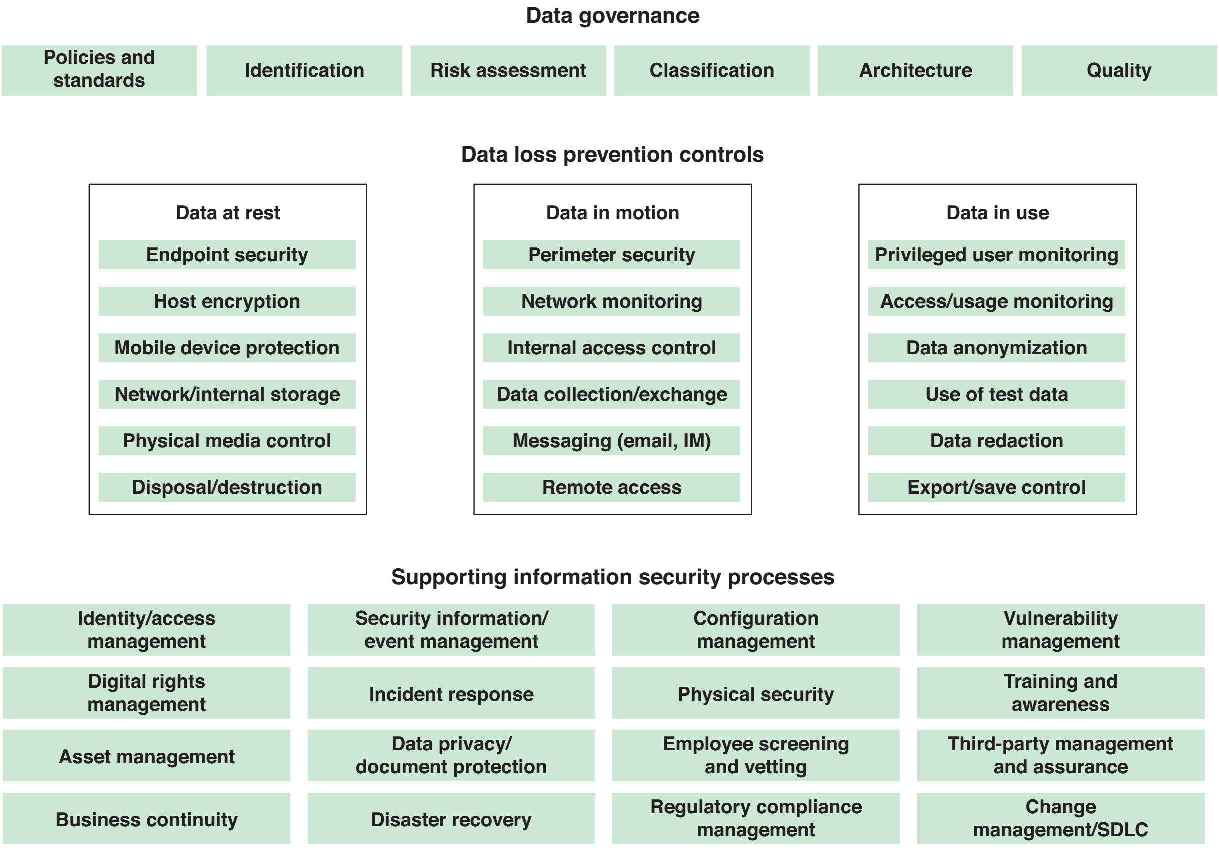

Figure 2: Data loss prevention architecture

Note: SDLC, software development life cycle.

breaches of different classes or categories of data. Examples of factors to consider include:

- Is the data protected by regulations (eg GDPR)?
- What is the relative value of internal data (eg board papers vs corporate customer lists)?
- What is the direct impact to customers and business partners of a specific type of breach (eg disclosure)?
- What is the impact on corporate reputation?
- Is there a potential loss of competitive advantage in the market?

Addressing these types of questions helps privacy and security managers prioritise DLP activities so that the highest risk data is protected first.

Data loss prevention controls

The next level of the model, DLP controls, applies to the two high-level ways data are stored within an organisation:

- Structured repositories: These are databases, such as relational databases, that are typically supported and controlled by the IT organisation.
- Unstructured repositories: This type of data is generally end-user driven and stored in less controlled repositories, such as network shares, SharePoint sites and workstations.

Key to effective design of DLP controls is to develop an understanding of the places and times at which data are vulnerable. A useful way of managing DLP is to categorise data into three states: data at rest, data in motion

Table 1: Data states

DLP state	Definition	DLP objective	Examples
Data at rest	Data in a stable storage system that is not often updated, if at all. Organisations often store such data in encrypted form.	Locate and catalog sensitive information stored throughout the enterprise.	Corporate files stored on servers or cloud storage; relatively static databases; and reference or stable data on desktops and laptops.
Data in motion	Data moving through any kind of internal or external network.	Monitor and control the movement of sensitive information across enterprise networks.	E-mail attachments, web uploads/downloads, instant messages, transfers over local networks or to/from a cloud, and mobile device traffic
Data in use	Data that is in the process of being generated, updated, processed, erased or viewed through various endpoint interfaces.	Monitor and control the movement of sensitive information on end-user systems.	Data processing by office applications, viewing/editing PDF files, database functions, cloud applications, and mobile applications.

Notes: DLP, Data loss prevention; PDF, Portable document format.

and data in use. Table 1 defines these states and the DLP objectives corresponding to each state. Figure 2 suggests the types of controls that may be appropriate for the three states of data.

Data at rest presents significant risk for enterprises. A large enterprise may have millions of files and database records on drives and removable media. A particular set of data files or records may have a home location, but portions of that data may also migrate to other storage locations, and this situation, if not monitored and controlled, can quickly become unmanageable. One example of how data may be replicated and proliferated is file sharing. With networked computer systems, file sharing for collaborative projects is common, but this may mean that the owner or creator of a file has no idea of what happened to the file after sharing it. The same risk exists with the many web-based collaboration and document management platforms in common use.

The fundamental task of DLP for data at rest is to identify and log where specific types of information are stored throughout the enterprise. The DLP unit uses some sort of data discovery agent that performs the following actions:

- Seek out and identify specific file types, such as spreadsheets, word processing documents, e-mail files and database records. The automated search activity encompasses files servers, storage area networks, network attached storage, cloud storage and endpoint systems. The term endpoint system refers to systems with a fixed attachment to the corporate network, such as workstations and servers.
- Once files are found, the agent must be able to open each file to scan the content for specific types of information. There needs to be a policy and mechanism for encrypted data, such that an entire encrypted block is labelled as sensitive or the agent has capability to decrypt and determine sensitivity.
- The agent logs files that contain information of security relevance and may issue alerts if a security policy is violated. Security for the log file is paramount, because it would be a useful tool for an adversary.

Data in motion presents two types of risks. While sensitive data is in transit, it is subject to eavesdropping, and so policies must be in place to secure sensitive data, such as secure protocols, encryption and the use of VPNs. The other sort of risk is that a user may thoughtlessly or deliberately send data to a location that is not secure. For example, an employee may send sensitive data to a personal webmail account in order to work at home.

Data in motion solutions operate in one of the two modes:

- Passive monitoring: Observes a copy of data packets as they move across a network link. Packets or sequences of packets containing information of interest can be logged and security violations can trigger an alert.
- Active monitoring: Interposes a relay or gateway type of device on a network line to analyse and forward data packets. The active monitor can log and issue alerts but can also be configured to block data flows that violate a security policy.

To inspect the information being sent across the network, the DLP solution must be able to: monitor the network traffic, recognise the correct data streams to capture, assemble the collected packets, reconstruct the files carried in the data stream and then perform the same analysis that is done on the data at rest to determine whether any portion of the file contents is restricted by its rule set.

The advantages of passive monitoring are that it is nonintrusive and does not slow down the flow of traffic. This approach, when combined with instant alerts and fast incident response, can be quite effective. The advantage of active monitoring is that it can enforce a DLP policy proactively. This comes at the cost of increased network and computational demands.

The types of risks involved with data in use involve some sort of misuse of sensitive data. For example, a disgruntled employee copying files containing personal information to portable devices, or a user printing sensitive data to equipment in common areas that can be accessed by others.

Data-in-use solutions generally involve installing DLP agent software on endpoint systems. The agent can monitor, report, block or quarantine the use of particular kinds of data files and/or the contents of the file itself. The agent can also maintain an inventory of files on the hard drives and removable media that is plugged into the endpoint. The agent can allow or disallow certain types of removable media, such as requiring that the removable device support encryption.

Supporting information security processes

A number of information security processes support the implementation of the DLP controls. The lowest portion of Figure 2 depicts the most important of these. In designing a DLP capability, the organisation needs to identify the other information security processes in order to establish and monitor multiple layers of defence.

For example, effective logical access controls may be in place, but if physical controls fail and sensitive hard copy information is removed from corporate facilities, data loss still occurs. Similarly, if changes to the IT infrastructure are not carefully controlled, existing DLP controls can become ineffective.

DLP FUNCTIONS

As indicated in Figure 1, there are four principal DLP functions: manage, discover, monitor and protect.[2]

The manage function is typically hosted on a central management platform. The platform enables an authorised administrator to define, deploy and enforce DLP policies by communicating with DLP agents on other platforms. The manage function includes a data loss reporting capability and management of incident response and remediation.

The purpose of the discover function is to locate the organisation's sensitive information, which may reside in an in-house data centre, the cloud, networked storage systems, mobile devices, network devices and endpoint devices. Discovery involves communication between the central management server and software agents. Discovery is an ongoing function because an organisation may generate or import new sensitive data over time. The discover function makes use of content analysis to identify sensitive data. This function creates a master inventory of where sensitive data is stored.

A key element of the discover function is locating sensitive data that is unprotected or not fully protected. The discovery of vulnerable data could trigger a relocation of the data to a secure platform or a signal to the data owner to secure the data.

The monitor function is used to understand how the organisation's sensitive information is being used, including what data is being handled, and by whom. With respect to network traffic, the monitor captures and analyses traffic. The function can be configured to monitor traffic on specific protocols, such as simple mail transfer protocol (SMTP), file transfer protocol (FTP), hypertext transfer protocol (HTTP), and various instant messaging (IM) protocols. As described earlier, there is place for both passive and active monitoring in DLP solutions.

Monitoring can also be performed by DLP agents on endpoints, to detect the downloading of data to local drives, copying to removable media devices and printing or faxing electronically.

The protect function is intended to stop sensitive information from being leaked or stolen by enforcing data loss policies and educating employees. The protect function extends to endpoint, network and storage systems. For data at rest, this function can invoke automatic encryption, relocation or removal of sensitive data. For data in use, this function can restrict printing, saving, copying, accessing, moving and transferring of sensitive data to portable media. For data in motion, the protect function can stop data from being sent in violation of security or privacy policies, or encrypt for secure exchange over a secure protocol.

DLP COMPONENTS

A comprehensive DLP solution that covers all sensitive data in an organisation consists of a number of components, as listed in Figure 1. The components are the following:

- DLP manager: It has been discussed previously. This is the management platform where an administrator can build and deploy policies. The DLP manager controls DLP agents in various network and endpoint devices.
- Storage scanner: The storage scanner may be hosted on the DLP manager platform or on networked devices that report back to the manager. Scanner agents scan data repositories, including databases, document repositories, networked file shares and cloud storage.
- Network monitor: These are network devices that passively monitor at strategic points on the enterprise network.
- Network gateway/prevent: These are active network devices that intercept traffic to determine if security/privacy policies are satisfied. Common applications are outbound e-mail traffic and web traffic.
- Endpoint agent: Endpoint agents are software modules on endpoint systems, such as personal computers (PCs) and workstations. These agents can perform both discover and protect functions.
- Mobile DLP: A typical DLP solution for mobile devices is to require that they connect to the corporate network through a virtual private network (VPN) to send corporate messages or access the corporate network. The device connects to a VPN

server, which can perform monitor and prevent functions.
- Cloud DLP component: There are two common elements of cloud DLP. One is the integration of DLP with cloud-based e-mail service. The other element is software that sits between cloud service consumers in the enterprise and cloud service providers to perform discover, monitor and prevent functions.

ANALYSIS OF DATA

A detailed discussion of the algorithms and approaches used to analyse data is beyond the scope of this paper. This section provides a brief overview. Data analysis techniques for DLP fall into two broad categories: context analysis and content analysis. Context analysis and content analysis work together to maintain security and privacy policies. Content analysis identifies sensitive data in real time, including data at rest, data in motion and data in use. Context analysis monitors sensitive data access and usage in applications, data repositories and network devices. DLP dynamically uses content and context analysis to automatically, and in real time, detect high-risk data access, transmission and usage. High-risk activity can trigger notification, remediation and quarantining types of responses.

Content analysis

All sensitive data and PII within an enterprise needs to be protected at all times and in all places. As a first step, the enterprise needs to determine what types of data are sensitive data and, if necessary, establish different levels of sensitive data. Then, there is a need to recognise sensitive data wherever it is encountered in the enterprise. Content analysis refers to methods of recognising sensitive data in real time. Key considerations in configuring a set of techniques are the processing power and time required for each technique, and the rate of false positives and false negatives. False positives (ie falsely identifying a data object as sensitive) create high costs in time and resources that are required to investigate and resolve apparent incidents that are not actual incidents. False negatives (ie failing to identify a data object as sensitive) obscure gaps in security by allowing data loss, the potential for financial losses, legal exposure and damage to the reputation of an organisation. False negatives are especially dangerous because the organisation does not know it has leaked sensitive data.

The following are common approaches to the recognition task[3,4]:

- Rule based: Regular expressions, keywords and other basic pattern-matching techniques are best suited for basic structured data, such as credit card numbers and social security numbers. This technique efficiently identifies data blocks, files, database records and so on that contain easily recognised sensitive data. This is the most common analysis technique available in DLP products and is best considered as a first-pass filter. It cannot be relied upon alone because of its high false positive rates.
- Database fingerprinting: This technique searches for exact matches to data loaded from a database, which can include multiple-field combinations, such as name, credit card number, and card verification value number. For example, a search could look only for credit card numbers in the customer base, thus ignoring employees buying online. This is a time-consuming technique but has a very low false positive rate.
- Exact file matching: This technique involves computing the hash value of a file and monitoring for any files that match that exact fingerprint. This is easy to implement and can check if a file has been accidentally stored or transmitted in an unauthorised manner. Unless, however, a more time-consuming cryptographic

hash function is used, this is trivial for an attacker to evade.

- Partial document matching: Looks for a partial match on a protected document. The method involves the use of multiple hashes on portions of the document, such that if a portion of the document is extracted and filed elsewhere or pasted into an e-mail, it can be detected. It is useful for protecting sensitive documents.
- Statistical analysis: Involves the use of a variety of statistical techniques and a collection of files and other data. The purpose is to find content that is similar to protected content. If the degree of similarity is high enough, then protection is extended to the new content. It is useful for protecting data in cases where it is difficult or infeasible to be specific and exact in defining what is to be recognised. These techniques are prone to both false positives and false negatives.
- Form recognition: Detects images of forms containing sensitive information, such as tax forms, medical forms and insurance forms. This capability should be able to operate on a variety of formats, including word documents, PDF (portable document format) documents, JPEG (joint photographic experts group) files, and PNG (portable network graphics) images.
- OCR-sensitive image recognition: Extracts text from images (scanned documents, screen shots, pictures etc) and from PDFs, enabling the use of new or pre-existing text-based detection rules on this content.
- File type detection: Detects file types according to their binary signature. The file signature is typically a short sequence of bytes at the beginning of the file used to identify the file type. An example of the use of this capability is to enforce that a certain type of document should never leave the organisation (such as an AutoCAD file).

These techniques are common in most commercial DLP products. Which ones to use and how to combine them depend on the nature and range of data to be protected and the magnitude of the estimated risk.

Context analysis

Context analysis, as part of DLP, was developed to address the need for specific security requirements and solutions in a complex environment. In many organisations, data moves from location to location and is stored on a variety of devices and in cloud storage applications. It is accessed by employees, partners and customers via connections and devices from inside and outside the enterprise network. Thus, an equally dynamic security solution is required for monitoring and responding to these events, which is the role of context analysis.

Context information typically includes location and time of data use, transmission and storage; identity of user or device, such as user ID and internet protocol (IP) address; and type of application. DLP script authoring tools can be applied to define acceptable and unacceptable situations. For example, if a payroll employee is observed viewing someone else's remuneration package, this event is a normal behaviour and can be ignored. If, however, this event were to occur from another department, the DLP should raise a flag that will trigger an action.

Context analysis and content analysis work together to maintain security and privacy policies. Content analysis identifies sensitive data in real time, including data at rest, data in motion and data in use. Context analysis monitors sensitive data access and usage in applications, data repositories and network devices. DLP dynamically uses content and context analysis to automatically, and in real time, detect high-risk data access, transmission and usage. High-risk activity can trigger notification,

server, which can perform monitor and prevent functions.
- Cloud DLP component: There are two common elements of cloud DLP. One is the integration of DLP with cloud-based e-mail service. The other element is software that sits between cloud service consumers in the enterprise and cloud service providers to perform discover, monitor and prevent functions.

ANALYSIS OF DATA

A detailed discussion of the algorithms and approaches used to analyse data is beyond the scope of this paper. This section provides a brief overview. Data analysis techniques for DLP fall into two broad categories: context analysis and content analysis. Context analysis and content analysis work together to maintain security and privacy policies. Content analysis identifies sensitive data in real time, including data at rest, data in motion and data in use. Context analysis monitors sensitive data access and usage in applications, data repositories and network devices. DLP dynamically uses content and context analysis to automatically, and in real time, detect high-risk data access, transmission and usage. High-risk activity can trigger notification, remediation and quarantining types of responses.

Content analysis

All sensitive data and PII within an enterprise needs to be protected at all times and in all places. As a first step, the enterprise needs to determine what types of data are sensitive data and, if necessary, establish different levels of sensitive data. Then, there is a need to recognise sensitive data wherever it is encountered in the enterprise. Content analysis refers to methods of recognising sensitive data in real time. Key considerations in configuring a set of techniques are the processing power and time required for each technique, and the rate of false positives and false negatives. False positives (ie falsely identifying a data object as sensitive) create high costs in time and resources that are required to investigate and resolve apparent incidents that are not actual incidents. False negatives (ie failing to identify a data object as sensitive) obscure gaps in security by allowing data loss, the potential for financial losses, legal exposure and damage to the reputation of an organisation. False negatives are especially dangerous because the organisation does not know it has leaked sensitive data.

The following are common approaches to the recognition task[3,4]:

- Rule based: Regular expressions, keywords and other basic pattern-matching techniques are best suited for basic structured data, such as credit card numbers and social security numbers. This technique efficiently identifies data blocks, files, database records and so on that contain easily recognised sensitive data. This is the most common analysis technique available in DLP products and is best considered as a first-pass filter. It cannot be relied upon alone because of its high false positive rates.
- Database fingerprinting: This technique searches for exact matches to data loaded from a database, which can include multiple-field combinations, such as name, credit card number, and card verification value number. For example, a search could look only for credit card numbers in the customer base, thus ignoring employees buying online. This is a time-consuming technique but has a very low false positive rate.
- Exact file matching: This technique involves computing the hash value of a file and monitoring for any files that match that exact fingerprint. This is easy to implement and can check if a file has been accidentally stored or transmitted in an unauthorised manner. Unless, however, a more time-consuming cryptographic

Table 2: Portions of PCI-DSS requirements suitable for DLP solution

Requirements	DLP feature
3.2 Do not store sensitive authentication data after authorization (even if encrypted). If sensitive authentication data is received, render all data unrecoverable upon completion of the authorization process	The DLP discover function, including network, storage, and endpoint discovery, can scan and quarantine or delete PCI authentication data.
3.2.1 Do not store the full contents of any track (from the magnetic stripe located on the back of a card, equivalent data contained on a chip, or elsewhere) after authorization. Note: In the normal course of business, the following data elements from the magnetic stripe may need to be retained: • The cardholder's name • Primary account number (PAN) • Expiration date • Service code To minimize risk, store only these data elements as needed for business	There are pre-existent templates in most DLP tools to detect PCI data captured using a magnetic-stripe. Any such data that is discovered other than those that need to be retained can be eliminated by the DLP function.
3.2.2 Do not store the card verification code or value (three-digit or four-digit number printed on the front or back of a payment card used to verify card-not-present transactions) after authorization	The DLP discover function, including network, storage, and endpoint discovery, can scan and quarantine or delete PCI authentication data.
3.2.3 Do not store the personal identification number (PIN) or the encrypted PIN block after authorization	The DLP discover function, including network, storage, and endpoint discovery, can scan and quarantine or delete PCI authentication data.
4.1 Use strong cryptography and security protocols to safeguard sensitive cardholder data during transmission over open, public networks, including the following: • Only trusted keys and certificates are accepted. • The protocol in use only supports secure versions or configurations. • The encryption strength is appropriate for the encryption methodology in use	Set all traffic to Block Mode except secure protocols such as TLS and IPsec when PCI data is identified using PCI data identifiers
4.2 Never send unprotected (primary account numbers) PANs by end-user messaging technologies (for example, e-mail, instant messaging, SMS, chat, etc.).	Using a regular expression function, create a definition for PANs and block using DLP

Notes: DLP, Data loss prevention; PCI-DSS, Payment Card Industry Data Security Standard; SMS, short message service; TLS, Transport layer security.

becomes sensitive PII. The filtering hierarchy to be searched for in order is:

(1) Social security number
(2) Passport number
(3) Driver's license/state identification number
(4) Bank account/credit card number
(5) Medical/Health Insurance Portability and Accountability Act (HIPAA) information
(6) Date of birth
(7) Mother's maiden name

For example, the passport number filter scans for the word 'passport' in English or Spanish followed by a string of digits.

All outgoing messages from a network are subject to filtering for sensitive PII. The DLP system quarantines any e-mail with a match, and sends the e-mail sender an auto-generated e-mail message describing the possible violation, the quarantine of the e-mail message and the steps necessary to release the e-mail message to the intended recipients. The options are encrypt the e-mail, remove the sensitive data or contact the privacy office. If the employee suspects the DLP quarantine is in error (false positive) and contacts the privacy office, another e-mail message is sent stating the results of the review.

If the sender does not act upon the quarantine notice within a predetermined period of time, the DLP system alerts the user that the e-mail message and any attachments have been deleted and not sent.

The DLP system also scans incoming e-mail. The DLP system blocks any incoming e-mail message containing PII from entering the DOC network. The DLP system sends an automated notification to the sender describing the policy prohibition, with instructions for using DOC-approved encryption software. In addition, the DLP system electronically notifies the intended recipient of the block message.

CONCLUSION

DLP is a mature technology deployed by many enterprises to support information security requirements. DLP is also a powerful tool for providing information privacy protection that will satisfy recent regulations and standards, such as the GDPR[8] and ISO 27701.[9]

Key features of DLP that respond to information privacy requirements are the following:

- The ability to automatically analyse and classify the content of transmitted, used and stored data of many formats and types
- The ability to enforce real-time preventive security controls in a wide range of data-leakage channels and scenarios
- The ability to control data operations based on their context
- Central management by IT security/privacy administrators rather than be local systems administrators

References

1. Stallings, W. (2020) 'Information Privacy Engineering and Privacy by Design', Pearson, Upper Saddle River, NJ.
2. Liu, S. and Kuhn, R. (2020) 'Data loss prevention', *IT Pro*, March/April, 2010, available at: https://www.computer.org/csdl/magazine/it/2010/02 (accessed 17th June, 2020).
3. Mogull, R. (2007) 'Understanding and selecting a data loss prevention solution', *SANS Institute White Paper*, 3rd December, available at: https://securosis.com/assets/library/publications/DLP-Whitepaper.pdf (accessed 30th April, 2020).
4. Symantec (2019) 'Symantec data loss prevention administration guide', 19th August, available at: https://support.symantec.com/us/en/article.doc9261.html (accessed 30th April, 2020).
5. Payment Card Industry (2018) 'Data security standard: Requirements and security assessment procedures version 3.2.1', May, available at: https://www.pcisecuritystandards.org/document_library (accessed 17th June, 2020).
6. Leadvue (2016) 'How a Data Loss Prevention (DLP) solution can help achieving PCI 3.0 compliance', Symantec Connect, 29th February, available at: https://www.symantec.com/connect/articles/how-a-data-loss-prevention-dlp-solution-can-help-achieving-pci-30-compliance?page=1 (accessed 30th April, 2020).
7. U.S. Department of Commerce (2014) 'Privacy data loss prevention working group recommendations', 17th December, available at: http://osec.doc.gov/opog/privacy/Memorandums/DLP-Memo_04152016.pdf (accessed 30th April, 2020).
8. European Parliament. 'Regulation (EU) 2016/679 OF THE European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)', available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679 (accessed 17th June, 2020).
9. International Organization for Standardization (2019) 'Security techniques — Extension to ISO/IEC 27001 and ISO/IEC 27002 for privacy information management — Requirements and guidelines', ISO/IEC 27701, August, available at: https://www.iso.org/standard/71670.html (17th June, 2020).