DYNAMIC DATA PROTECTION SERVICES FOR NETWORK TRANSFERS: CONCEPTS AND TAXONOMY

Stefan Lindskog[1, 2], Anna Brunstrom[1], and Erland Jonsson[2]

[1] Department of Computer Science, Karlstad University, Sweden

[2] Department of Computer Engineering, Chalmers University of Technology, Sweden

Stefan.Lindskog@kau.se, +46 54 700 1152, SE-651 88 Karlstad, Sweden

Anna.Brunstrom@kau.se, +46 54 700 1795, SE-651 88 Karlstad, Sweden

Erland.Jonsson@ce.chalmers.se, +46 31 772 1698, SE-412 96 Göteborg, Sweden

ABSTRACT

Security should be thought of as a tuneable system attribute that allows users to request a specific protection level as a service from the system. This approach will be suitable in future networking environments with heterogeneous devices that have varying computing resources. The approach is also appropriate for multimedia applications that require tuning the protection level to maintain performance at levels that are acceptable to users. In this paper, we investigate data protection services for network transfers that are designed to offer variable protection levels and propose a taxonomy for such services. The taxonomy provides a unified terminology for dynamic data protection services and a framework in which they can systematically be inspected, evaluated, and compared. The taxonomy is also intended to provide a basis for development and identification of current and future user and/or application needs. It comprises four dimensions: protection service, protection level, protection level specification, and adaptiveness. On the basis of our taxonomy, we made a survey and categorization of existing dynamic data protection services for network transfers.

KEY WORDS

Network security, data protection, dynamic services, taxonomy.

DYNAMIC DATA PROTECTION SERVICES FOR NETWORK TRANSFERS: CONCEPTS AND TAXONOMY

1 INTRODUCTION

The overall goal of computer and network security is to protect data and/or resources from being unduly tampered with. Various protection as well as detection and response schemes have been proposed to enforce security. User identification and authentication, access control and auditing services are now integrated into modern operating systems. Security extensions [6] for detecting malicious programs and intrusions, network and personal firewalls to block unwanted network traffic and, various types of cryptographic systems to protect data that are either stored or in transit are widely available today.

The demand for security will vary heavily in future networking environments. Security should therefore be considered a tuneable system attribute that allows users to request a specific protection level as a service provided by the system. This approach will be suitable when heterogeneous devices that have varying computing resources are used. The approach is also appropriate for multimedia applications that require tuning the protection level in order to maintain performance at levels that are acceptable to users.

The lack of mechanisms by which system owners and users can request a specific level of protection as a service in the system makes it impossible to offer protection based on need. Instead all users are offered similar services regardless of whether it is the desired level of protection and all users are forced to bear the costs of either too much or too little protection. Furthermore, unnecessarily high levels of data protection can make systems more difficult to control, e.g., network management becomes harder, processor load on servers and clients increases, smaller hand-held devices are not able to encrypt and/or decrypt data in real-time etc. This results in applications with inadequate data protection and unnecessary costs to users.

In contrast, companies that can offer products where users know what to expect from the system will have a great competitive advantage and will be able to offer customers the most cost-effective, user-friendly and resource-efficient solutions. The field of applications that would benefit from such solutions would span all areas where computers are used–industrial automation, control systems, public service, transportation and traditional networking applications such as e-business, banking, government services and legacy systems.

With dynamic data protection services, security could eventually be introduced as a Quality of Service (QoS) parameter in current and future communication networks. However, dynamic services today are not sufficiently well understood in that context. Thus, a thorough analysis of existing services and their dynamic features is a first step towards integrating security as a QoS dimension.

In this paper we propose a taxonomy for dynamic data protection services in order to introduce a unified terminology and provide a framework in which they can be systematically inspected, evaluated, and compared. The objective of our taxonomy is twofold: first to provide a framework for classifying existing dynamic data protection services and second to provide a basis for development and identification of current and future user and/or application needs. The taxonomy comprises four dimensions: (1) protection service, (2) protection level, (3) protection level specification, and (4) adaptiveness.

Concepts and terminology are introduced in section 2. Section 3 discusses five different dynamic data protection services: IP security, transport layer security, Authenticast, a scalable encryption service, and a dynamic encryption service. Section 4 gives the proposed taxonomy. In

section 5, the five dynamic data protection services covered in section 3 are classified according to the taxonomy. Concluding remarks are given in section 6.

2 CONCEPTS AND TERMINOLOGY

It is well known that security is composed of a number of aspects: confidentiality, integrity, and availability. Confidentiality is the ability to prevent unauthorized disclosure of data, while integrity is the ability to prevent unauthorized modification of data and/or resources. Finally, availability is the ability to prevent unauthorized withholding of data and/or resources. These aspects describe different, and in some cases contradictory, requirements on the underlying systems and communication channels. For example, a system may be configured to offer high confidentiality and/or high integrity at the cost of reduced availability. Similarly, if availability is the key issue, the level of confidentiality and/or integrity must be reduced.

Security is typically implemented through one or more security services. A combination of protective and detective services is commonly used today. For example, firewalls are often used to block suspicious network traffic to and from internal networks, and Intrusion Detection Systems (IDSs) are used as a complement to firewalls to detect insider and outsider intrusion attempts as well as successful intrusions.

Neither firewalls nor IDSs are suitable security tools for protecting data that are transferred over an insecure network, such as the Internet. Instead some form of data protection service is needed. The type of service necessary depends on what is to be protected. There are protection services to achieve data confidentiality, data integrity and data authenticity [20][1]. A data confidentiality service ensures that transmitted data are accessible only for reading by authorized parties, while a data integrity service ensures that only authorized parties are allowed to modify transmitted data. Finally, a data authenticity service ensures that the origin and/or the source of data are correctly identified. Data protection services are typically based on various cryptographic algorithms.

In this paper, we consider services that aim to protect data (or message) transfers in a networking environment. Such services should be designed and implemented according to the **principle of adequate security**. Pfleeger and Pfleeger [14] define this principle as follows:

> "*Computer items must be protected only until they lose their value. They must be protected to a degree consistent with their value.*"

This principle specifies that data with a short lifetime can be protected by a protection service that is effective only for that short period. Our focus in this paper is on dynamic data protection services. We define such services as follows:

Definition 1. A dynamic data protection service is a service that has been explicitly designed to offer various protection levels that can be selected at run-time.

Support for dynamic data protection has also been integrated directly into several multimedia applications, e.g., Nautilus [12] and Speak Freely [19]. Such application level protection is not considered in this paper, however. Instead, we focus on generic data protection services that can be used by networked applications to achieve dynamic protection.

---

[1] Note that Stallings [20] defines also non-repudiation as a kind of data protection service. A non-repudiation service ensures that neither the sender nor the receiver of a message can deny the transmission. From this definition it is clear that a non-repudiation service protects the transmission but not the particular content. We have for that reason decided not to count such services as data protection.

3 EXISTING DYNAMIC DATA PROTECTION SERVICES

Five examples of dynamic data protection services are described below, starting with IP security and followed by transport layer security. A dynamic authentication service, called Authenticast, is then presented. A simple scalable encryption service and a dynamic encryption service are finally described.

3.1 IP Security

IP Security (IPSec) is a protocol that implements security at the IP level. It is described in [8]. IPSec is provided as an integrated part of IP version 6 (IPv6) [3] and can be added as an extension to IP version 4 (IPv4) [15]. Security services in IPSec are implemented through extension headers that follow the main IP header attached to each packet. Two types of such headers exist: Encapsulating Security Payload (ESP) headers and Authentication Headers (AHs). An ESP header is attached when data confidentiality is requested, while an AH is attached when data authenticity and/or data integrity is requested. When an ESP header is used, an optional AH can be attached as well.

A key issue in IPSec is the concept of Security Associations (SAs), which are stored in a database called a security policy database (SPD). An SA is a one-way relation between a sender and a receiver. An SA specifies what kind of security service is to be applied in a particular communication session and is normally defined by the following parameters: sequence number counter, sequence counter overflow, anti-replay window, AH information, ESP information, and lifetime of the SA. The AH information contains information about authentication algorithm, key, key lifetime, and other parameters related to AH. The ESP information contains information about encryption and authentication algorithms, keys, key lifetimes, and related parameters being used with ESP. The lifetime of the SA specifies a time interval or a byte count after which an SA must be replaced with a new SA or terminated.

With key lifetimes and the lifetime of the SA, both keys and algorithms might change at run-time. A protocol referred to as the Internet Security Associations and Key Management Protocol (ISAKMP) [11] is used to establish, modify and delete SAs. ISAKMP defines a standardized packet format used to agree on an SA. The Oakley [5] key exchange protocol, which is derived from the Diffie-Hellman key exchange protocol [18], is typically used to exchange keys in IPSec.

3.2 Transport Layer Security

Transport Layer Security (TLS) [1, 4] is an Internet protocol that implements security features above the Transmission Control Protocol (TCP) [16] or the Stream Control Transmission Protocol (SCTP) [21], which implies that a reliable end-to-end security service is provided. The first version of TLS was derived from the Secure Socket Layer (SSL) protocol, introduced originally by the Netscape Corporation. Version 3.1 of SSL and TLS 1.0 are essentially the same protocol. Two types of services are offered by TLS: confidentiality and message integrity. Confidentiality is achieved through the use of conventional encryption algorithms such as DES, 3DES, IDEA, etc., and message integrity is achieved through the use of a hash algorithm (SHA-1 or MD5). In addition, a compression algorithm can be used to compress data before adding the message authentication code (MAC) produced by the hash algorithm, which in turn is added before encryption of data is performed. In TLS, both the application data and the MAC are encrypted.

A key component in TLS is the session concept. A session in TLS is an association between two communicating parties, i.e., between a client and a server. Sessions are created by the handshake protocol, which is an essential component of the TLS specification. The handshake protocol is executed before any application data are transmitted. On successful completion of the handshake protocol, encryption algorithm, hash algorithm, and compression method, together with related parameters, such as keys, initialization vectors (IVs), hash size, etc., have been exchanged and agreed upon.

3.3 Authenticast

Authenticast [17] is a dynamic authentication protocol designed and implemented by Schneck and Schwan. The main idea in this protocol is to address security and performance trade-offs in client-server environments.

Authenticast is a user-level communication protocol that provides variable levels of security that can be changed at run-time. Each connection has an associated security level, where the security level is defined as the percentage of data packets that are authenticated (i.e., verified) before they are processed by the client. Two different authentication algorithms are supported: RSA and DSA. A user selects one of these algorithms together with related parameters such as keys and key lengths, but may later change to the other algorithm. With a component called "security thermostat", users are able to specify a desired security level and a security level range.

The specification of a security level range is used by the protocol to adaptively change the security level at run-time. The range specifies the minimal and maximal percentage of packets that must be authenticated. It is used to decrease the security level, within user specification, during times of increased load, in favor of other, more important computations.

Note, however, that Authenticast does not support variations of the security level at the server side. This implies that every packet sent by the server is signed.

3.4 Scalable Encryption

A content-independent scalable encryption (SE) service is proposed by Lindskog et al. in [9]. The basic idea in this service is to apply strong encryption to only parts of the content. The remaining parts are either encrypted with a weak and fast algorithm or left unencrypted.

The SE service utilizes existing symmetric encryption algorithms and is block-oriented. An "m-out-of-n" selection mechanism is proposed, where m and n are two integer variables that must be agreed upon in advance. Variable m specifies the number of blocks to be encrypted with the strong algorithm out of n blocks. Hence, if m = 3 and n = 4, three out of four blocks, i.e., 75 % of the blocks, are encrypted with the strong encryption algorithm.

The main advantage of the SE service is that it allows users to make trade-offs between security and performance. However, security level adaptation is not part of the model. This implies that the security level is specified during session initiation and thereafter remains fixed during the whole lifetime of the session.

3.5 Dynamic Encryption

Magnusson and Nilsson [10] recently proposed a dynamic encryption (DE) service that is implemented on top of SCTP.

Each connection (or association) in SCTP consists of one or more logical data streams. Streams are unidirectional and independent of each other. When a DE connection is established, three streams are created, see Figure 1. Stream 0 is used for data transfers and in-band signalling, while streams 1 and 2 are only used for out-of-band signalling. The specification of the current encryption level is signalled in-band. Out-of-band signalling can for instance be used to negotiate on the initial encryption level.

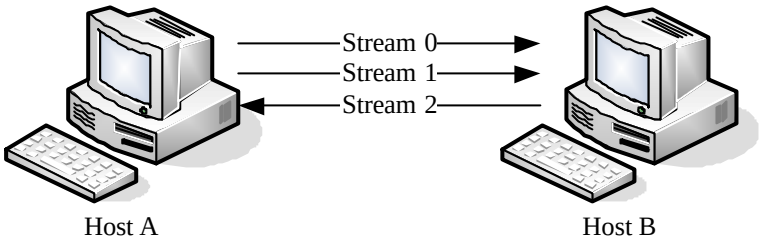

Figure 1. An established connection between hosts A and B in the DE service. Host A is the sender of data and host B the receiver.

The encryption level is specified through an encryption mask, which is transferred before the actual data are sent. The encryption mask consists of a length field and a bit vector. The length field indicates the length of the following bit vector, and the bit vector specifies which of the following data blocks are encrypted. For example, a mask with a length of 4 and the bit vector 1101 specifies that blocks 1, 2, and 4 following the mask are encrypted, while block 3 is not. The mask itself is always encrypted. Whenever the sender would like to either change the amount or the order of blocks being encrypted, a new mask is transferred to the receiver. This implies that the encryption level can easily be changed at run-time.

4 THE TAXONOMY

We propose here a taxonomy to introduce a unified terminology for dynamic data protection services and provide a framework in which such services can be systematically inspected, evaluated and compared. Our intention is also to provide a basis for development and identification of user and/or application needs. The taxonomy consists of four dimensions: protection service, protection level, protection level specification, and adaptiveness. An overview of the taxonomy is given in Table 1.

Table 1. Taxonomy of dynamic data protection services.

Protection service	Data confidentiality
	Data integrity
	Data sender authenticity
	Combined
Protection level	Algorithm selection
	Selective protection
	Combined
Protection level specification	User-defined
	System-defined
Adaptiveness	Per-session
	In-session

4.1 Protection Service

Protection of data that are transferred in a networking environment may be achieved in different ways depending on the security requirement(s) that must be met. We distinguish between the following three basic types of data protection services:

- Data confidentiality

– Data integrity

– Data sender authenticity

A data confidentiality service prevents unauthorized disclosure of information that is transferred between a sender and a receiver. A data integrity service, on the other hand, prevents unauthorized modification of information. Finally, a data sender authenticity service allows a receiver to verify who the sender actually is.

Common to all these services is that they are implemented using some kind of cryptographic system. For example, symmetric encryption systems, such as DES [18] and AES [2], are often used to achieve data confidentiality. Data integrity is often implemented using cryptographic hash functions, such as MD4, MD5, or SHA [18]. Data sender authenticity, on the other hand, is typically implemented using some form of asymmetric cryptographic system [17].

Note also that some data protection services, e.g., IPSec and TLS, implement two or more of the basic protection services mentioned above. Such services are referred to in our taxonomy as combined data protection services.

4.2 Protection Level

A protection level can be specified in different ways. In our taxonomy, we distinguish between two fundamentally different ways of specifying a particular protection level:

– Algorithm selection

– Selective protection

In some services, a protection level can simply be specified through the selection of a particular protection algorithm together with its related parameters. For example, in a confidentiality service, it might be possible to specify a particular algorithm (DES, 3DES, AES, etc.), mode (Electronic Code Book (EBC), Cipher Block Chaining (CBC), etc.), key length, block length and number of rounds. Ong et al. [13] suggested that quality of protection for an encryption service can be specified in the following form: `<content type, interval of security, encryption algorithm, encryption key length, encryption block length>`. The strength of the encryption, expressed through the encryption algorithm, key length, and block length, depends on the content type and on the time interval to keep the data secure. The last three parameters correspond to the category of algorithm selection in our taxonomy. However, we do not consider the parameters of content type and interval of security part of the protection level. Rather we see them as criteria on which to base the selection of a protection level. Irvine and Levin [7] similarly suggested the use of algorithm selection to achieve different levels of security.

The protection level could alternatively be specified on a selective basis. One way to express selective protection is through the usage of a percentage-based protection level. For example, a protection level of 0 % means that nothing of the content is protected, while 100 % means that the whole content is protected. Accordingly, a protection level of 50 % means that half of the content is protected.

A dynamic data protection service in which the protection level could be specified through both algorithm selection and selective protection is referred to as a combined service with respect the protection level.

4.3 Protection Level Specification

In some situations no data protection at all is needed. In other cases, one or more of the data protection services described in subsection 4.1 are needed to protect data in transit. Ideally, the protection services should be designed and configured in such a way that they can be invoked

selectively and with a desired protection level. We distinguish between two classes of specifications:

- User-defined
- System-defined

In a user-defined specification, the protection level is determined directly by the user. The protection level for a dynamic data confidentiality service could simply be a specification of the amount of blocks or packets to be encrypted or a specification of the encryption algorithm to be used with related parameters. The protection level for a dynamic data integrity service could be the amount of blocks or packets that are signed and/or verified.

In system-defined specifications, the protection level is either determined by the system to enforce a certain security policy or by an application to guarantee a pre-specified protection level. In a banking application, for example, it might be pre-specified that account numbers, passwords, etc. must be encrypted with a certain strength. Additionally, in an organizational security policy, there might be a rule that specifies that emails sent externally must be encrypted to a certain minimal degree.

4.4 Adaptiveness

The last dimension in our taxonomy is referred to as adaptiveness. An adaptive service is one that allows changes of the protection level at run-time. We distinguish between two main classes of adaptive services:

- Per-session
- In-session

In per-session adaptive protection services, the protection level is specified at the inception of a communication session, and, once specified, remains fixed during the lifetime of the session.The highest degree of adaptiveness is offered by in-session adaptive protection services. In an in-session adaptive service, the protection level can vary during the lifetime of a session.

5 A CLASSIFICATION OF EXISTING DYNAMIC DATA PROTECTION SERVICES

The five dynamic data protection services described in section 3 are classified in this subsection according to the taxonomy we presented in section 4. Table 2 shows how IPSec, TLS, Authenticast, the scalable encryption service, and the dynamic encryption service are classified.

Table 2. Classification of the five dynamic data protection services described in section 3.

	Dimension			
Protection service	**Protection service**	**Protection level**	**Protection level specification**	**Adaptiveness**
IPSec	Combined (data confidentiality, integrity, and authenticity)	Algorithm selection	User- or System-defined	Per-session
TLS	Combined (data confidentiality and data integrity)	Algorithm selection	User-defined	Per-session
Authenticast	Data authenticity	Combined (algorithm selection and selective protection)	User-defined	In-session

Scalable Encryption	Data confidentiality	Selective protection	User-defined	Per-session
Dynamic Encryption	Data confidentiality	Selective protection	User-defined	In-session

IPSec and TLS are classified as combined services with respect to the protection service, while Authenticast is a pure authenticity service, and the two encryption services are pure data confidentiality services.

The protection level in IPSec and TLS is algorithm selection. Authenticast is the only service investigated that offers both algorithm selection and selective protection and is thus classified as a combined service with respect to the protection level. Both the scalable encryption and dynamic encryption service offer only selective protection.

Protection level specifications are user-defined in all five services. However, IPSec can also be configured and used in such a way that the protection level specification is system-defined. IPSec is very flexible and can be configured to use a particular data protection service based on a number of parameters or parameter combinations. The configuration could for instance be based on destination address, source address, userID or transport layer protocol.

IPSec, TLS, and the scalable encryption service are all classified as per-session adaptive data protection services, which means that a protection level is essentially selected at the inception of a communication session. The possibility to exchange new security parameters during a session through a handshake protocol, as can be done for example in IPSec, is not considered sufficient for a classification of in-session adaptivness. In Authenticast and the dynamic encryption service, an initial protection level is specified at the inception of a session. The initial protection level can however be changed at any time within a particular session. For that reason these two services are classified as in-session adaptive services.

6 CONCLUDING REMARKS

This paper introduces a taxonomy for dynamic data protection services and defines and discusses the corresponding security concepts. The purpose of the taxonomy is to establish a unified terminology and a framework in which dynamic data protection services can be developed, evaluated, and compared in a quantitative and systematic way. The ultimate goal is that this should help us to introduce security as a QoS parameter. Five existing dynamic data protection services have been classified according to the taxonomy in order to illustrate some of its benefits. We are aware that more work remains to be done to arrive at complete and fully comprehensive coverage of the area. Still, we hope that our work is a step towards integrating security as a full-fledged QoS dimension.

7 ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their insightful comments and suggestions. This research is supported in part by grants from the Knowledge Foundation of Sweden.

8 REFERENCES

1. Simon Blake-Wilson, Magnus Nystrom, David Hopwood, Jan Mikkelsen, and TimWright. *RFC* 3546*: Transport Layer Security (TLS) Extensions.* June 2003, Status: Standard.

2. Joan Daemon and Vincent Rijmen. The Design of Rijndael: AES-The Advanced Encryption Standard. Springer-verlag, 2002.

3. Stephen E. Deering and Robert M. Hinden. *RFC 2460: Internet Protocol, Version 6 (IPv6) Specification*. December 1998, Status: Standard.

4. Tim Dierks and Christopher Allen. *RFC 2246: The TLS Protocol Version 1.0*. January 1999, Status: Standard.

5. Dan Harkins and Dave Carrel. *RFC* 2409*: The Internet Key Exchange (IKE).* November 1998, Status: Standard.

6. Hans Hedbom, Stefan Lindskog, and Erland Jonsson. Risks and dangers of security extensions. In *Proceedings of Security and Control of IT in Society-II (IFIP SCITS-II)*, pages 231-248, Bratislava, Slovakia, June 15-16, 2001.

7. Cynthia E. Irvine and Timothy E. Levin. Quality of security service. In Proceedings of the New Security Paradigms Workshop, Ballycotton, County Cork, Ireland, September 19-21, 2000.

8. Stephen Kent and Randall Atkinson. *RFC* 2401*: Security Architecture for the Internet Protocol*. November 1998, Status: Standard.

9. Stefan Lindskog, Johan Strandbergh, Mikael Hackman, and Erland Jonsson. A content-independent scalable encryption model. In *Proceedings of the 2004 International Conference on Computational Science and its Applications (ICCSA 2004)*, part I, pages 821-830, Assisi, Italy, May 14-17, 2004.

10. Robert Magnusson and Tomas Nilsson An in-session dynamic encryption service. Master's thesis, Karlstad University, Karlstad, Sweden, 2004. To appear.

11. Douglas Maughan, Mark Schertler, Mark Schneider, and Jeff Turner. *RFC 2408: Internet Security Association and Key Management Protocol (ISAKMP).* November 1998, Status: Standard.

12. Nautilus secure phone homepage. http://nautilus.berlios.de/. June 9, 2004.

13. Chui Sian Ong, Klara Nahrstedt, and Wanghong Yuan. Quality of protection for mobile applications. In *Proceedings of the 2003 IEEE International Conference on Multimedia & Expo (ICME 2003)*, Baltimore, Maryland, USA, July 6-9, 2003.

14. Charles P. Pfleeger and Shari Lawrence Pfleeger. Security in Computing, third edition. Prentice Hall Inc., Upper Saddle River, New Jersey, USA, 2003.

15. Jon Postel. *RFC 791: Internet Protocol*. September 1981, Status: Standard.

16. Jon Postel. *RFC 793: Transmission Control Protocol*. September 1981, Status: Standard.

17. Phyllis A. Schneck and Karsten Schwan. Dynamic authentication for high-performance network applications. In *Proceedings of the Sixth IEEE/IFIP International Workshop on Quality of Service (IWQoS'98)*, Napa, California, USA, May 18-20, 1998.

18. Bruce Schneier. Applied cryptography: Protocols, algorithms, and source code in C, second edition. John Wiley & Sons Inc., New York, USA, 1996.

19. Speak Freely homepage. http://www.speakfreely.org/. June 9, 2004.

20. William Stallings. Cryptography and Network Security: Principles and Practice, second edition. Prentice-Hall Inc., Upper Saddle River, New Jersey, USA, 1999.

21. Randall R. Stewart, Qiaobing Xie, Ken Morneault, Chip Sharp, Hanns Juergen Schwarzbauer, Tom Taylor, Ian Rytina, Malleswar Kalla, Lixia Zhang, and Vern Paxson. *RFC 2960: Stream Control Transmission Protocol*. October 2000, Status: Standard.