THREAT MODEL FOR USER SECURITY IN E-LEARNING SYSTEMS

Maria Nickolova, Eugene Nickolov

Abstract: This paper suggest a generic threat model for the processes in e-learning systems. After a breaf discussion of the particularities of e-learning systems from the security point of view, the role of threat models in system security is explained. Then a characterization of the attackers is made, together with the assets and entry points identification and threats identification and classification. The most probable attacks for e-learning system are described regardless of the specific implementation, to help the e-learning system developers to eliminate or mitigate these attacks if possible in the design stage not waiting for actual attacks to occur.

Keywords: Security, threat modeling, e-learning.

ACM Classification Keywords: D.4.6 Security and Protection

Introduction

During the past three decades we are witnesses to a radical and transformative technological revolution that has resulted in fundamentally new ways of seeking information, communicating and learning. E-learning has arised as a new way of distributing knowledge without geographic or temporal restrictions and has passed through a several distinct phases – from Computer Based Training through to Learning Management Systems and Courseware Management Systems. There are a number of factors for this evolution, including the following:

- Ongoing development of dedicated e-learning software applications;
- Wide adoption of portals in e-learning;
- Offerings from publishers of value added services to the e-learning market;
- The advent of mobile learning which became a significant area of research and development (e.g., through the European MOBILearn project).

However, this evolution of e-Learning also brings with it a new set of threats for user and application security that impact the design of e-learning content and tools.

Particularities of E-learning Systems from the Security Point of View

E-Learning systems allow multiple users or applications to download, upload and exchange distributed information. Communication issues between end-users' computers and e-learning site (portal) in these systems are very important, as the systems are defined by widely dispersed elements in terms of network topology and physical geography. Additionally, the systems often allow many-to-many communication which provides powerful capabilities and allows many system nodes to have the same communication at any given time.

As noted in [1], a system can be attacked only through its "entry points". Designers of many computing systems can then limit security risks to a large degree by reducing the number of entry points for example by using network communication as little as possible. E-learning system designers obviously could not use this method. Since by definition e-learning communication system is an entity which exists in many physical/logical locations simultaneously, entry points are so prevalent that the system itself can in some ways be defined by them.

Additional unique challenges arise from the dynamic nature of these systems. To be effective in their intended application, it is desirable to allow the most flexibility by enabling dynamic sessions (e.g. on-line tests) in which any process may join or leave the group session at any time. To contrast, in a static session (e.g. read-only processes) all end-users only receive information during the entire session.

Another concern is that we never know in advance which processes will exist on a system at any given time. So not only are there a large volume of potential entry points to a system, but it is impossible to know exactly how many or even where (in terms of physical/logical location) those entry points will be. Another concern is that a malicious process could always attempt to join a session (by a single user's computer or through the Internet) and conduct attacks from within the system.

Further, the security of the overall system depends not only on the security of the network and sub-networks within the system, but also on the security of each particular member process, the machine(s) on which the process resides, and the communication protocols used between participating processes. Because the identity of an e-learning system user is not necessarily tied to specific machines, it may be appropriate to implement strict authentication criteria and "credentials" verification to better control the intrusion of malicious processes and the possible loss of privacy and confidentiality. However, this may impose some performance decrease and result in a percentage of legitimate processes/users being denied participation, thereby reducing the system effectiveness.

Numerous entry points essentially force e-learning systems to rely heavily on encryption schemes for security, but if not implemented carefully they can lead to additional vulnerabilities. The difficulty in controlling access to entry points, or even identifying entry points, suggests the conclusion that it will be better to make security efforts in restricting an attacker's ability to do anything useful at each entry point. Two principles are used in cryptographic key distribution schemes for confidentiality in e-learning systems, refering to the notions that only current group of learners may possess key information for current traffic - backward secrecy and forward secrecy [2]. Backward secrecy ensures that new learners cannot decipher old traffic. Forward secrecy ensures that ex-learners cannot decipher new traffic. Adherence to these principles essentially requires new keys for the entire group of learners either at frequent intervals, or every time any learner leaves or joins the group.

The Role of the Threat Models in System Security

To design a secure e-learning system it is not sufficient to choose strong authentication and encription system and to implement new security solutions only in response to actual attacks that have recently occurred. The non-systematic nature of this approach could leave unprotected several points of the attack space and allow to break many e-learning systems with little effort. Therefore, a secure e-learning technology must address the issues of security as part of an organized process in the design phase. But a system can't be made "secure" in general, it may be secure only against a specific attack. Therefore, when designing a secure e-learning system, designers must have a clear idea of threats they have to prevent and of technical capabilities the attackers have, i.e. the necessary preliminary step in the secure design of e-learning system is to answer the question "secure against what ?" The answer is the threat model - a set of hypotheses about who (what) and how could attack the system.

Generally, threat models have three main purposes [1]:

- To improve a design's security by anticipating specific attacks and implementing countermeasures in advance.
- To anticipate the varying outcomes of "successful" attacks (for example cracks) and their possible impact.
- To enable the creation of advance response plans to deal with significant attacks as and when they occur.

The intent is that through the threat model as many vulnerabilities as possible could be identified by the developers, rather than to be left for later identification by attackers.

Generally the threat model includes three high-level steps: characterizing the attacker, identifying assets and entry points, and identifying and classifying the threats [1]. Characterizing the attacker consists of identifying goals, motivation, and capabilities. Identification must be made of the system's assets and entry points through which an attacker will seek access to these assets. The threat profile of an e-learning system must describe potential threats arising from the overall threat model and this information must be submitted to risk assessment where a decision-maker will decide to eliminate, mitigate, or accept the risk associated with each threat.

Threat Model for the Processes in e-Learning Systems

Let's now based on the general threat models theory try to create a specific model describing the threats, the motives and the means that represent a risk for the security of e-learning systems. This will be a generic model that will have to be adapted to the real systems.

The Attacker. In reality, it is difficult to answer who the potential attacker to an e-learning system is and what his capabilities may be. Attackers can take many forms. Some of them may perform their actions deliberately. Others may simply be incompetent, legitimate users of the system. Attackers that act deliberately can be divided into many categories: proof-of-concept hackers, crackers, users wanting to get illegal free access to protected resources, disgruntled employees or even bored but technically savvy teenagers. We could try to enumerate common attackers for a specific system but it is impossible to overlook all potential attackers. Their capabilities embrace unbounded time for attack launch (teenagers), extensive knowledge, privileged access (insiders).

The Assets. An asset is any element of an e-learning system which provides critical functionality. Any threat could be defined in part by the asset which the attacker wants to get access to. Of course the goal of any protection solution is not to eliminate the assets but to protect them. However, to protect assets we must first identify them. For any generic e-learning system, the following assets could be targeted by an attacker:

- E-learning content;
- Cryptographic key content;
- User personal data;
- Messages between users;
- Different group membership data;
- Network bandwidth;
- Message integrity;
- Message availability.

Of course, threat sub-models should be created for each of these assets based on the particularities of the specific real system. These sub-models should be designed to answer the following questions:

- What is the damage that can be done to the particular asset?
- What are the vectors of attack that can be used to get access to this asset?
- What conditions would have to be in place for an attack to be successful?

The Entry Points. Entry points to an e-learning system could be defined as points an attacker must use to acquire access to assets. Here is a non-exhaustive list of potential entry points to a generic e-learning system:

- Used network protocols;
- Used communication channels;
- Computers of past, current, or future e-learners;
- Physical network infrastructure;
- Logs gathering data relevant to e-learning sessions.

Now, based on the described above e-learning system’s potential attackers, assets, and entry points we could build a threat model of an e-learning system. However, all systems cannot be threat modeled in the same way. Several different approaches exist [3] depending on the particular type of system and the intended purposes of the modeling. A threat model for an e-business system can be based upon the Data Flow Diagram (DFD) of the system. Other systems (e-commerce) may be better modeled using Layered Network Model Approach. We choose the well-known security aspects of Availability, Integrity, Confidentiality and Authentication (AICA) as the basis for our threat modeling of e-learning systems because these characteristics are the most important for the end-user security.

AICA Threat Modeling Approach for E-learning Systems

The four major security aspects in any computing system: Availability, Integrity, Confidentiality and Authentication (AICA) are well known and widely accepted. In this article we provide a description of the applicability of each element of AICA to an e-learning system and classify the different types of attacks in accordance to the related element. Because every attack may affect at least one aspect of the AICA model (and sometimes more), all threats can be covered by this approach (though our list is not exhaustive). Table 1 displays a summary of AICA aspects and attacks independent of the specific e-learning system implementation.

Availability	Integrity	Confidentiality	Authentication
Denial-of-Service	Malicious code attacks	Group session eavesdropping	Brute force attacks
Node attacks	Message injection	Group session traffic analysis	Dictionary attacks
Link attacks	Traffic modification	Group identity disclosure	Login spoofing attacks
Network infrastructure attacks	Traffic deletion		Key management attacks
	Traffic rerouting		Replay attacks
	Traffic misdelivery-rerouting		Man-in-the-middle attacks
	Forgery attacks		Session hijacking attacks
	Stack overflow attacks		Non-repudiation attacks

Let's now start the analysis of the specific attacks in the AICA threat model.

Availability Attacks. Availability attacks attempt to make e-learning services and data (or metadata) unavailable to legitimate users for a period of time. Here, we briefly discuss the two basic types of availability attacks: blocking attacks and flooding attacks. A blocking attack stops authorized users from accessing a resource by physically or electronically destroying the route to the resource. The non-malicious blocking attack (physical denial-of-service) embrace communications lines failure, crash of software/hardware, or inadvertent reconfiguration of systems in a way that prevents access. Malicious blocking attacks may occur when an attacker destroys or delete critical files, invalidates accounts or changes access control protocols.

A flooding attack overloads the e-learning system with a large number of requests to stop authorized users from accessing its resources. A flooding availability attack typically exploits a flaw in the design/implementation of a network protocol, operating system or commonly used application. Common examples of flooding attacks are:

Denial-of-Service. An e-learning system should have reasonable capacity (in terms of bandwidth and connectivity) to meet the peak demands, however, this capacity is finite and can be exhausted. DoS attacks could be very dangerous for e-learning systems because a single message/packet may be replicated to many receivers over many links. These attacks may be malicious, but they can also be caused unintentionally or carelessly. Some are extremely complex, while others rely upon very simple methodologies. These attacks can be conducted in several ways on an e-learning system:

Masquerading Sender DoS Attacks. The attacker may gain access to the e-learning system by joining it through an authentication attack. Once authentificated, a misbehaving source can flood traffic to all other users to disrupt current and future sessions.

Masquerading Receiver DoS Attacks. Once authentificated through an authentication attack, the attacker may join the traffic by creating many end-user processes, thus greatly increasing the overhead of the system but not the traffic into sessions. The size of the session must therefore expand to handle the increased traffic which consumes bandwidth and processing resources.

Insider DoS Attacks. A legitimate end-user becomes a traitor by flooding traffic to all learners or subverting the e-learning system by signaling or creating many receiver processes.

Transit DoS Attacks. Without being authenticated, the attacker may inject unauthorized transit traffic on the same network(s) in order to disrupt communications. The widespread use of UDP as a transport protocol for real-time communications is a vulnerability since currently there is no UDP mechanism to prevent traffic congestion [4].

Availability to Present. A specific learner is denied from participiging to a session through a targeted DoS attack [5].

Availability to Receive. A specific learner is denied from receiving information in a session through a targeted DoS attack [5].

SYN flood attacks. They consist in bombarding the e-learning site with SYN packets used to open a connection and could either overload the e-learning system servers or cause them to crash.

S*murf attacks*. They use a PING packet sent by an end-user computer with the forged return address of the intended victim. This PING is broadcast to a large number of other users (intermediaries) and if many machines respond by sending PING packets to the intended victim, the result can be severe network congestion or outages, that could potentially make the network unusable. The intermediary can be victimized in his turn.

D*istributed denial-of-service attacks* multiply the effect of an ordinary denial-of-service attack by launching it simultaneously from multiple addresses. The attacker gains access to a number of other end-user computers in advance and puts code in each to launch the attack at a preset time or on a signal. Using such remote computers makes the attacker invisible, that's why distributed denial-of-service attacks are difficult to trace and prevent.

Node Attacks. By definition, in e-learning communications there are more nodes (senders/receivers) involved in a session communications than point-to-point communications and thus more potential exposure to node-targeted attacks. Each legitimate user node may be attacked through a node-specific vulnerability and be degraded, subverted or otherwise non-functioning member in the e-learning process. This decreases the availability of any information or contribution to joint projects and activities from attacked end-user nodes.

Link Attacks. Communication traffic in an e-learning system uses many more links than point-to-point communications so it is exposed to link attacks [4]. As some legitimate nodes are connected to the network by a single or few links, if these links are successfully degraded by an attack, the affected nodes may no longer be available to the system and the system is no longer available to them. If link attacks are strategically placed, network partitioning and even general system degradation may occur.

Network Infrastructure Attacks. If any part of the e-learning network infrastructure which directly or indirectly supports the communication sessions in progress, is physically/logically attacked or otherwise damaged (power, routers, switches, hubs, servers for DNS, etc.), the system will degrade and may not remain functional.

Integrity Attacks. Integrity attacks attempt to actively modify or destroy information in the e-learning site without proper authorization. Modification may include creating, changing, appending, and deleting both data and metadata. With an integrity attack, authorized users can gain access, but what they find when they get there is not what is supposed to be there. We shall briefly discuss some integrity attacks:

Malicious Code Attacks. Integrity attacks can be non-malicious in origin, as many unintentional causes may corrupt or modify data. However, malicious integrity attacks are becoming both more common and destructive. Malicious code comes in a variety of forms – virus, Trojan horse, worm etc. E-learning system administrators as well as end-users should protect and regularly chech their systems to assure they are malware-free.

Message Injection Attacks. An attacker who joins an e-learning communication, for example via an authentication attack, may then freely inject messages into the system which will be viewed as legitimate system traffic by the other end-users.

Traffic Modification Attacks. An attacker may intercept packet data, rearrange or delete specific bits, and/or forward data as if no changes occurred. This attack does not require knowledge of key data or any particular understanding of the e-learning data itself.

Traffic Deletion Attacks. Similarly to the previous attack an attacker may simply delete data on the communication channels. Again, this requires no particular understanding of the e-learning data itself.

Traffic Mirror-Rerouting Attacks. An attacker may mirror the traffic by rerouting it to unauthorized end-users. This is analogous to inserting a mirror into the stream of observable traffic and re-directing it elsewhere, without affecting the original stream's destination – thus making this attack difficult to detect.

Traffic Misdelivery-Rerouting Attacks. If the attacker reroutes traffic to unauthorized receivers without mirroring, some messages are lost and not received by some or all end-users. This attack is more detectable than traffic mirror-rerouting attacks since end-users and messaging protocols can detect lost messages.

Forgery (counterfeit) Attacks. They make a false representation of data that has come from another address. An attacker can also hijack a session by intercepting a communication session and continuing it in the name of (and with the privilege of) the victimized end-user.

Stack Overflow Attacks. The attacker intentionally supplies a very large amount of input data (for example, 3,000 characters in a limited length field), in order to exceed the space allocated and spill over into adjacent data or code areas, either corrupting other values or inserting new commands to be executed.

Confidentiality Attacks. Confidentiality attacks are passive attacks that expose confidential data to the view of unauthorized readers [4].

Unlike integrity attacks, confidentiality attacks don't alter the e-learning content, but only affect the security level and dissemination of this content, as well as learners' personal data. However, confidentiality attacks may be used as a first step in availability or integrity attacks, as where the attacker obtains confidential passwords by defeating encryption or simply by password guessing. Confidentiality attacks to e-learning systems may take the form of commercial appropriation of confidential information. Most involve intrusions and disclosure of private facts that are not commercially motivated, but are undertaken solely for the amusement of the attacker.

In many e-learning systems, the unauthorised access to confidential data may require the relatively difficult task of gaining access to the system on the administrator level. However, an attacker may also read information without illegitimate privilege escalation (e.g. insider attackers). Storage and timing leaks via covert channels also fall into this class of attack. In an e-learning system exist more eavesdropping opportunities in comparison to most networked systems, because of the multiple channels used through the multicast communication protocols [4]. Further, attackers don't need eavesdrop if they can instead simply become learners – at which point the system views them as legitimate and purposely sends them all relevant traffic. Common protection techniques for this type of attack require extensive use of encryption with cryptographic keys to ensure privacy. The cryptography used may or may not be the same as that used for authentication of the learners. Let's now discuss in short some confidentiality attacks:

Session Eavesdropping Attacks. Traffic between particular learner's processes in an e-learning system is virtually impossible to hide. Because of the dispersal of the processes throughout a typically large network, as well as because the volume of communication channels used, it does not seem feasible to prevent attackers from observing traffic. This is why encryption methods are critical for communications in an e-learning system. The goal is not to prevent enterely the attacker from observing traffic, but rather to ensure that the observable content is unusable or meaningless to him.

Group Session Traffic Analysis. Even if appropriate encryption measures are taken, the attacker may still observe the traffic flow and make corresponding deductions based on when and where messages are sent, message type and volume.

Identity Disclosure Attacks. The identity of e-learners is attacked for unauthorized disclosure [6].

Authentication Attacks. Authentication attacks occur when an attacker masquerades as a legitimate end-user (using a stealed password, key, or credential) or an attack device masquerades as a legitimate device participating in e-learning scheme both usually aiming to get free access to paid e-learning networks and services. But a masquerader can also launch insider attacks to access data/metadata (confidentiality attack), modify data/metadata (integrity attack), and/or deny others data/metadata (availability attack) based on the legitimate user identity authorization capabilities they have taken over. Let's discuss briefly the following attacks:

Brute force attacks. These attacks are unsophisticated and can be effective only if power computer is used. The principle is to attempt every possible combination of characters to satisfy password (key) authentication.

Dictionary attacks. They could be considered as a subset of brute force attacks. Instead of trying all password combinations, a dictionary attack attempts to satisfy a password prompt by trying commonly used passwords from a list or dictionary.

Login spoofing attacks. An attacker can use a fake login program that prompts the legitime user for an ID and password. Instead of logging the user into the requested e-learning system, the bogus program stores or forwards the stolen credentials and returns a notice that the login has failed.

Key Management Attacks. In a system in which the key management scheme does not adequately protect both forward and backward secrecy, an attacker may use information obtained from a former legitime user to gain access to the e-learning system. The attacker may in fact be a former legitime learner aiming to find enough information to break the current encryption scheme and to access the content for free. A basic example may be procedural mishandling of keys exposing them to disclosure. A centralized server that generates and disseminates keys for legitime users represents a single entry point and an attractive attack target.

Replay Attacks. In certain encryption schemes, it may be possible for an attacker to eavesdrop on a communication channel and to record an encrypted message. This recorded message may later be played back in such a way as to allow the attacker to masquerade as the original sender. The attacker may appear legitimate to the recipient and authorized to access data to which he hasn't a right to.

Man-in-the-Middle Attacks (MITM). In MITM an attacker actively intercepts messages between end-users and servers of the e-learning site and is able to read, modify and retransmit these messages without legitime participators in the communication either knowing that their session has been compromised. The MITM attack may include: eavesdropping (including traffic analysis and known plaintext attacks), chosen ciphertext attacks, substitution attacks, replay attacks, and DoS attacks.

Session Hijacking Attacks. In this kind of attack an attacker intercepts and continues a session started by a legitime user of the e-learning site. In contrast to MITM session hijacking occurs after a session is established and the legitime user's session is terminated.

Non-Repudiation Attacks. Non-repudiation requires the ability to prove to a third party that a particular message is sent or received preventing the source from later denying transmission of the message. As a consequence of existing vulnerabilities in access control, e-learning user may deny (legitimately or illegitimately) having participated in some or all communication sessions and having sent/received specific messages.

Source Non-Repudiation Attacks. An attack targeting an end-user in order to deny sending communications.

Receiver Non-Repudiation Attacks. An attack targeting an end-user in order to deny receiving communications.

Conclusions and Future Work

E-learning systems present a unique challenge to security engineers because of their nature and inherently complex architecture. That's why, it is especially important to use a systematic formal methodology of identifying threats during the development phase. We propose the AICA threat modeling approach for focusing e-learning systems' design efforts on the elimination or mitigation of the threat associated risk, regardless of the specific implementation. The goal of this work is not to provide an exhaustive list of all possible attacks for e-learning systems; it rather tries to provide a conceptual framework by which developers of e-learning systems can decrease the number of overlooked security vulnerabilities at the design stage. The presented AICA threat model needs to be tested in practical e-learning implementations so as to really ascertain its performance.

References

1. Swiderski ,F. and W. Snyder, Threat Modeling (Microsoft Professional). Microsoft Press, 2004.
2. Judge, P. and M. Ammar, "Security Issues and Solutions in Multicast Content Distribution: A Survey," IEEE Network, January/February 2003.
3. Myagmar, S., A. J. Lee, W. Yurcik, "Threat Modeling as a Basis for Security Requirements," In Symposium on Requirements Engineering for Information Security (SREIS), 2005.
4. Ballardie, T. and J. Crowcroft, "Multicast-Specific Security Threats and Counter-Measures," Symposium on Network and Distributed System Security, February 1995.
5. Tannenbaum A., M. van Steen, Distributed Systems: Principles and Paradigms, Prentice-Hall, Inc., 2002.
6. Nickolova, M. and E. Nickolov, User Privacy Problems in DRM protected E-learning Objects – fron Standards to Practice, proceedings or the Second International Workshop on Compiter Science and Education in Computer Science, June, 5-7 2006, Borovetz – Sofia, Bulgaria.

Authors' Information

Maria Nickolova – National Laboratory of Computer Virology; BAS, Acad.G.Bonthev St., bl.8, Sofia 1113, Bulgaria; e-mail: maria@nlcv.bas.bg

Eugene Nickolov – National Laboratory of Computer Virology; BAS, Acad.G.Bonthev St., bl.8, Sofia 1113, Bulgaria; e-mail: eugene@nlcv.bas.bg