Humans as the Strong Link in Securing the Total Learning Architecture

Fernando Maymí[1], Angela Woods
and Jeremiah Folsom-Kovarik

[1] Soar Technology, Ann Arbor, Michigan, USA
{fernando.maymi, angela.woods, jeremiah}@soartech.com

Abstract. This paper describes a proposed approach, centered on human factors, for securing the Total Learning Architecture (TLA). The TLA, which is being developed for the United States Department of Defense, will rely on large stores of personal data that could be targeted by sophisticated adversaries. We describe the TLA and its envisioned users at a fairly high level before describing expected classes of attacks against it. We then examine existing and proposed controls that, if properly managed, should allow users and service providers to significantly reduce the risks to the system.

Keywords: Total Learning Architecture · Cybersecurity · Threat Modeling · Human-systems Integration

1 Introduction

The last twenty years have brought significant advances in educational technology, which is now a core element of life-long learning for many children and adults in the western hemisphere. Though learning management systems (LMS) and education management information systems (EMIS) can be credited for many breakthroughs, they are by no means the only sources of learning activities, particularly those related to career advancement. Increasingly, people are turning to a variety of online resources in order to learn new skills, improve or maintain existing ones, or otherwise further their education.

This growing demand for educational technology solutions is driving many organizations to supply a variety of learning activities, most of which exist independently or within proprietary environments. Learners in this environment are forced to maintain multiple online personas and track their progress on site-by-site basis. As the number of learning activity providers that longitudinally track learner skills continues to grow, there is an opportunity to enhance competency mastery by aggregating these information sources and providing tailored recommendations to individual learners. This is the promise of the Total Learning Architecture (TLA).

2 Architecture Overview

The TLA is not a software system; rather, it is a set of Application Programming Interface (API) specifications that create a learning framework wherein learning activity providers and others can responsibly share learning data. The learners, providers and other relevant organizations create an environment within which learners are able to avail themselves of new learning opportunities while leveraging all their historical data. For example, if a learner is subscribed to activities at sites X and Y and then chooses to also participate in site Z, the TLA would allow competencies from all three sites to be considered when making recommendations for new activities in all sites, including the newly added site Z.

To illustrate the use of these TLA APIs, we have developed reference implementations of certain components (e.g. the Learning Record Store or LRS). However, the reference implementations themselves are not part of the formal definition of the TLA. Any developer is free to create their own replacement implementation of any component for which a reference implementation is provided, all that is necessary is components that conform to the API specifications that govern that area of the TLA.

2.1 Interfaces

The interfaces are what define the TLA; without them, the environment could not exist. These APIs perform two key functions: they define *what* is shared, and they specify *how* it is shared. The first of these functions is accomplished by enforcing consistent data structures. This creates a shared language that allows information to be unambiguously interpreted by different entities in the TLA. The second function, which deals with how this information is shared, is made possible by standardizing the transfer methods used when exchanging information about learning experiences between entities that comply with the architecture.

The TLA comprises multiple optional APIs that regulate everything from learning activities to the use of different assessment frameworks to learner profiles, just to name a few. Among these, the most developed is the Experience API (xAPI), which is the principal means by which any component can understand exactly what the learner has done in the past or is currently doing. The key object within xAPI is the Statement, which describes fine-grained communication about learner experiences from minute to minute in order to help interpret learner performance in context.

2.2 Data

The learner record store (LRS) is where all the learner experiences are stored. This, together with the activity providers, are the most important components of the TLA; without them, the architecture would not accomplish much. An LRS can be centralized, or it can be distributed. In fact, an activity provider can provide its own LRS, which could then interface with other stores to provide holistic tracking of competencies as well as a richer set of recommended activities for a given learner. The exchange of this data can be regulated by the user or sponsoring organization.

The data in a single xAPI Statement object, which captures a learning experience, can be thought of as a sentence with a subject, verb, object and, optionally, other components such as the outcome. For example, if an activity provider would like to store an experience in the LRS, the statement would add to a historical stream describing what an individual learner has encountered, accomplished, and done in context across all activity providers.

For concreteness, an example xAPI statement might add to Angela's record in the LRS to reflect that she scored a 97% on a graded task to configure a pfSense firewall. The activity sending the statement is recorded along with various metadata. Similarly, there can be a variety of assessments that could be mapped to the result. So, while the format of the Statement is fixed, it can be arbitrarily enriched by activity providers and others for the purpose of interpreting and understanding the learner's experience.

3 Use Cases

In order to better understand the uses and potential vulnerabilities of the TLA, it is helpful to describe some relevant use cases. The sections below provide a limited, but representative sample of cases in which this architecture would be used by learners both in their personal or work spaces.

3.1 Support a Job Task

Since learning On-The-Job (OTJ) is an important aspect of much adult training, the TLA could be used to support Just-In-Time (JIT) delivery of task support when a person needs help with a job task. The delivery device could be an embedded system or a personal mobile device, and the trigger for the TLA to offer help might be performance monitoring of job tasks through the same channels TLA uses to monitor instructional performance. Job task support is valuable for a range of populations including apprentices and novices, people responding to an infrequently occurring emergency, or people who need cognitive support for specific deficits.

3.2 Learn

We expect users of the TLA to spend most of their time learning in a structured, guided environment. The interaction of learners with these environments, which are developed and delivered by a variety of Learning Activity Providers, represents the most likely and frequent use case for the TLA. While these are perhaps best visualized by evolutions of the computer based training with which many of us are familiar, they will also involve novel modalities such as presenting flashcards on a smart watch right before the learner delivers a presentation.

3.3 Monitor Human Performance

A growing number of individuals are interested in monitoring and measuring information about their own daily lives for reasons of health, self-improvement, or simply personal interest [1]. In parallel, Defense researchers are carefully studying human performance and precursors or mediators that contribute to performance [2]. The TLA could help individuals learn about themselves by facilitating the empirical measurement and manipulation of individual experience. For example, a person who wants to compare their personal caffeine intake against their sleep patterns can record both in their own TLA Learner Record Store (LRS). By using the TLA, the person can gradually expand the data they collect as needed to learn more about themselves.

3.4 Integrate with Personal Assistant

At the time of publication, commercial assistants are available in most of the consumer computer and mobile platforms, including Siri, Google Assistant, and Cortana. Information in the TLA about learners could help tailor each of these assistants to individual needs. For example, when the learner engages in informal learning by searching through a commercial assistant, the TLA can help identify the appropriate reading level and the background knowledge the learner has. Of course, the TLA would also benefit from any information the assistants share about the learner's current context and life experience.

3.5 Track Progress

Some learning activities will evaluate the progress of learners automatically, while other activities will rely to some extent on inputs from other people. Instructors, supervisors, and perhaps others could interact with TLA components to directly assess or provide input into the assessment of learners. An instructor could manually grade exercises, whereas a supervisor could validate that a learner was (or was not) able to show evidence of a proficiency. In these cases, select individuals will have access to relevant components in order to track the progress of learners in their purview.

4 Threat Model

A threat model identifies threat sources and methods that can undermine the functionality of a system and result in losses to an organization. Our approach is to identify the classes of threat actors that would have the intent and capability to attack the TLA, and then infer the means by which they would accomplish their goals. What follows is the first threat model developed for the TLA.

4.1 Threat Actors

We focus our discussion of threat sources on four classes of actors: terrorists, nation states, insiders and criminals. These classes emerged from the misuse case analysis that is described in the following section, but we describe them here in order to facilitate our later description of their desired actions. We note that there are numerous other potential classes of threat actors who could attempt to compromise a TLA system; we simply focus on the ones that appear likeliest to threaten the target systems.

Terrorist. Terrorist threat actors could attempt to compromise a TLA system if they think that doing so would allow them to cause death or destruction. A potentially exploitable area are industrial processes that are increasingly automated and digitally connected, and present opportunities to remotely cause physical effects. An example would be a food processing plant in which a computer controls the amount of iron that is added to a popular breakfast cereal. If threat actors were to target industrial systems operators responsible for regulating the iron levels, they could cause iron poisoning on a national or perhaps international scale. The concomitant loss of public trust in the food supply would further magnify the effects of such an attack.

State Actor. We assess that state actors represent the greatest threat to TLA systems. They could be interested in using TLA systems to cause physical destruction (with or without loss of life). We have seen at least one example of this in the 2013 breach of computer systems at the Bowman Avenue Dam in New York. The U.S. government indicted seven individuals who allegedly targeted the dam on behalf of the Iranian government [6]. While they were not able to cause damage, the event is an indicator of increasing proficiency and desire to damage cyber physical systems (CPS).

State actors could also want to alter the data within the TLA in support of information operations (IO), which involve deliberate attempts to influence what a population believes on specific issues. The 2016 compromise of George Soros' organizations [7], attributed to Russia, is a good example of a state actor altering stolen information in support of IO. In that operation, the actors modified some of the files to give the impression that his foundation was funding Russian dissidents [8]. As part of IO, one could imagine state actors implanting false information to influence how users perceive an issue of interest to the actors or for other purposes.

Criminal. Unlike the state actor, criminals are motivated by financial profits. Typically, monetization is accomplished by stealing large volumes of personal data and then selling them on online markets [9]. The most valuable targets for these actors are repositories of personal [10] or financial information [11], credentials (e.g., passwords) [12] and valid email addresses [13]. Depending on the specific system involved, TLA components will almost certainly contain at least one and perhaps all these types of valuable information. It would be reasonable to expect that these systems would almost certainly be targeted by criminal threat actors.

Insider. Insider actors also want to access data, albeit for a different purpose than the other threat actors. The insiders would most likely be interested in reading other users learning records, either out of misguided curiosity [14] or as a form of cyber stalking [15]. In fact, the news media has reported on many cases of employees with access to federated information systems similar to the TLA who have been disciplined or fired for improperly accessing the records of others.

A less likely but more damaging goal for an insider actor would be the unauthorized modification of learning records. There have been cases in which public officials have been accused or convicted of falsifying such information. The alleged motives range from financial gain [16] to avoiding public relations disasters [17]. As TLA systems become increasingly common, they could present opportunities for insider actors to modify the information contained in them.

4.2 Misuse Cases

A use case is a short story that describes the interaction of one user with the system in order to accomplish a specific task. Collectively, the collection of all use cases describes the entire functionality of the new system. Security professionals have adopted this modeling technique to include not only what authorized system users will do, but what threat actors will want to do also. We employ a common approach to threat modeling called misuse cases [18]. Figure 1 illustrates a partial use case model for the TLA that has been augmented to also show the misuse cases. By convention, the authorized actors and use cases are depicted in white, while the threat actors and misuse cases are shaded.

In the diagram we see four threat actors we discussed in the previous section. Their misuse cases encompass their main goals. In the following subsections, we look at how, specifically, these actors could leverage the TLA to accomplish their objectives.

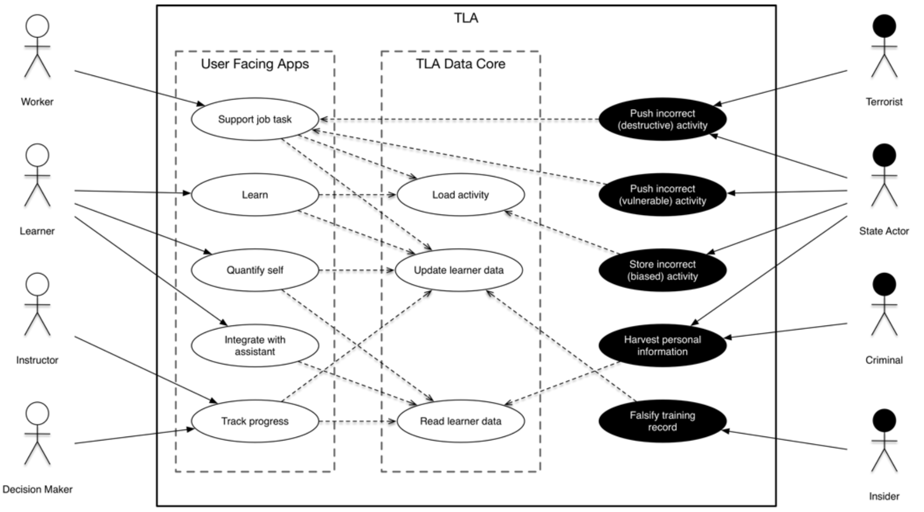

Fig. 1. TLA misuse case diagram showing the benign uses and users of the system filled in white (on the left) and the malicious uses and users of the system shaded black (on the right).

Push Incorrect (Destructive) Activity. One of the features supported by the TLA is providing just-in-time (JIT) training to learners. This is useful when users need to perform a task at which they are not proficient. TLA activity providers would then provide tutorials, step-by-step guides or similar activities to guide the user in performing the task. If the activity had been altered with malicious content, it could be used to direct the user to perform actions that would result in physical destruction or loss of life. An example of this would be an operator of a dam performing an infrequent remote test of the floodgate actuators. If a threat actor modifies the procedure in the JIT activity so that instead of testing it actually causes the floodgates to open, the users actions could result in flooding. Since the learners are unskilled at the task, they would be particularly vulnerable.

In order to modify the activities to contain destructive instructions, the terrorist or state actor would have to gain access to the repository of activities and modify the data without being detected. Stealing credentials would be a feasible way to accomplish this, provided the stolen identity has authority to edit the materials. This would not be so much a TLA attack as one against a specific TLA compliant system. It would also be a broad attack, since everyone who accesses that activity (including knowledgeable authors, administrators or auditors) could notice the modification.

Push Incorrect (Vulnerable) Activity. The illicit modification of JIT activities described above is not limited to destructive purposes. A nation state actor could leverage this feature to induce learners to misconfigure their information systems in order to make it easier for the threat actor to compromise them. An example of this would be a system administrator with limited proficiency in configuring rules for an intrusion detection system (IDS). That administrator could turn to the TLA for JIT training when updating malware signatures. Since an IDS can have cryptic rules, it would be unlikely that the administrator would notice that the rule being typed, instead of generating alerts, would cause the IDS to suppress them. This very simple modification of probably a line or two would make it inordinately easier for the threat actor to successfully attack the target system with very little risk of detection.

Push Incorrect (Biased) Activity. Both of the previous misuses of the TLA provide the learner with intentionally incorrect information. For reasons discussed above, it is preferable for the threat actor to not store the incorrect activity in a legitimate Activity Provider's data stores. There is, however, at least one scenario in which it would make sense for a nation-state actor to store incorrect information so that large numbers of learners have access to it. This scenario involves information operations, which are deliberate activities carried out in order to influence the thinking of target group. A benign example of this is modern marketing practices designed to persuade you to purchase a particular good or service. Less benign examples are misinformation campaigns carried out by oppressive regimes in order to pacify their citizens. Increasingly, however, we are seeing information operations carried out in large scale by nation state actors against citizens of other nations.

The effectiveness of incorrect and biased information is proportional to its reach and volume. This is unlike the previous examples of destructive and vulnerable activities. For this reason, nation state actors would want to store, as opposed to surgically push, this information with multiple Activity Providers and, specifically, for popular activities. These actors can accomplish this objective through a variety of means including stealing credentials for content editors, recruiting legitimate content editors to alter the information, and contaminating sources used by content editors so they contain the incorrect biased information. Note that not all these means are technical, but they all exploit vulnerabilities in the human component of the TLA.

Harvest Personal Information. Exploiting people is often facilitated by gaining access to volumes of personal information. Whether the goal is to recruit foreign agents [19] or to sell personal information online [9], nation state and criminal actors can put a lot of effort into harvesting as much information as possible about their human targets. By providing a means to aggregate a very large set of data on learners (many of them associated with the government) the TLA could provide a lucrative target to these actors.

The likely target within the TLA in this misuse case is the Learner Profile. Whether this data set is contained in one database (as in the prototype implementation) or in a federation of data stores, it is the heart of the TLA and, as such, must be accessible to most other components. This degree of connectivity could present a significant vulnerability if left unattended, so we will provide some recommended controls later in this paper. For now, it is important to keep in mind that the Learner Profile will require a more comprehensive set of controls than other entry vectors discussed in this section.

Falsify Training Record. We conclude our discussion of misuse cases with what is perhaps the least damaging of all: the falsification of learning records. Apart from causing perception and trust challenges, this case is fairly contained both in terms of actions and effects. The likeliest actors to engage in falsification are insiders seeking to modify their own or someone else's records to show proficiencies that are not real. The self-serving version of this act is easier to mitigate by controlling the ability of learners to modify their own records. The technical controls to accomplish this are already part of the TLA.

The challenge is in detecting when an insider inappropriately modifies someone else's records. At issue is the ability of supervisors, trainers and others to certify proficiencies of those under their watch. Technical controls alone are unlikely to prevent this type of misuse because it would require the TLA to differentiate between a legitimate and inappropriate certification by an otherwise authorized user. We will focus on this challenge in our later discussion of procedural controls.

5 Technical Controls

In this section, we address technical controls that can mitigate the risk of compromise in the five misuse cases we have developed. We stress, however, that each technical control's effectiveness can be either undermined or enhanced by appropriate user behaviors. In the discussion that follows, we consider protections to data while it is at rest in some component of the TLA as well as when it is in transit between components.

5.1 Data in Transit

In the first two misuse cases we presented (pushing destructive and vulnerable activities), terrorists or nation states replace or modify a legitimate stream of activity data intended for a specific user. Their goal is to destroy or otherwise exploit a system that the learner is attempting to configure using the JIT training functionality of the TLA. This attack is unlikely to be attempted by modifying the activity data in the providers' data stores (i.e., data at rest). Instead, the attacker would either intercept the learner request and directly provide the malicious activity, or selectively modify portions of the activity as they flow between a legitimate activity provider and the learner.

In the first case the threat actor prevents the learner from connecting to the activity provider and impersonates the latter. Once the learner is connected to the impostor provider, the threat actor is provides tailored content for that learner that results in the destruction of assets or in rendering them vulnerable to follow-on attacks. This case requires a significant amount of preparation since the threat actor must recreate the entire learning environment that the learner expects.

An alternative approach, which is illustrated in Figure 2, is one in which the threat actors simply inject themselves between the learner and the activity provider. From this position, they can selectively intercept, edit and then forward any part of the activity. The advantage is that the threat actors would not need to replicate the learning environment, but simply ensure all traffic flows through them. When the right content goes across the connection, say the instruction to "turn the knob slightly to the left," the actors could replace it with "turn the knob fully to the right. Ignore any alarms."

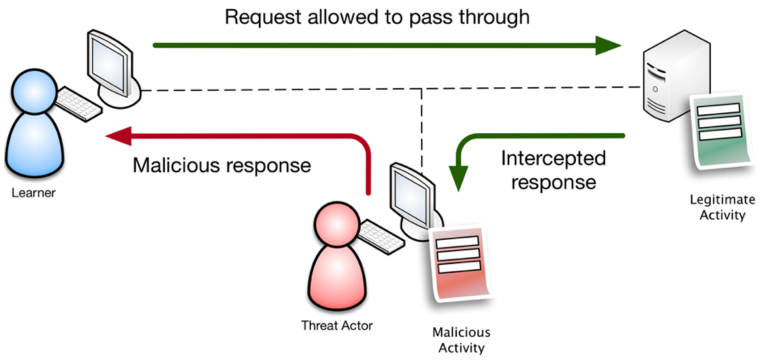

Fig. 2. Threat actor allowing a request from the user to reach the server, but intercepting and modifying the server's response intended for the user.

This type of attack is commonly called a Man-in-the-Middle (MitM), and requires a fair amount of sophistication to succeed. The key to mitigating a MitM attack is to ensure there is a secure link between the learner and provider. An example of this is the establishment of a secure hypertext transfer protocol (HTTPS) session in a client's web browser connecting to a financial institution's server. In order to set it up, the server must present evidence of its identity to the client. This evidence is almost always in the form of a public key infrastructure (PKI) certificate, which is tied to a specific internet domain (e.g., soartech.com). PKI certificates are signed or validated by a trusted third-party certificate authority (CA). The process by which this is done ensures that it is difficult for a threat actor to impersonate a legitimate site over HTTPS. Still, there are ways in which a threat actor might counter this technical control, which makes the user our next line of defense.

Using a Fake PKI Certificate. The PKI certificate is tied to a specific domain, but an actor could use a mismatched certificate (i.e., one whose domain doesn't match the uniform resource locator or URL). The learner would request a connection with the activity providerto which the threat actor would respond with its own certificate. This would cause the learner's browser to show an alert. Unfortunately, many naïve users will simply click "OK" on the warning and proceed with the connection to the bogus site. This risk is best mitigated through security awareness training in which the users are taught to recognize these certificate warnings as a serious threat. Furthermore, we should provide a simple technical means of notifying the appropriate security personnel if any links exhibit this behavior. Lastly, users who accurately report insecure conditions like this one should be recognized or otherwise rewarded for doing so.

Stripping the Secure Connection. The learner may request to connect to an HTTPS server, but it is possible for a threat actor to respond to that request in such a way that the browser accepts an unsecure connection instead. This is known as stripping the connection and is remarkably easy to do. An alert user would notice that the connection is not secure, since the browser would not display the secure icon on or near the address bar. Again, many users would not notice one way or another, but this can be remedied with security awareness, appropriate notification mechanisms, and a system of incentives for reporting anomalies such as this one.

5.2 Data at Rest

The data stored within the TLA could be a target to nation states that exploit information operations (IO). The main purpose of IO is to influence, disrupt, corrupt or usurp adversarial human decision making [20]. The means of carrying out an information operations attack would differ from the preceding discussion on pushing destructive or vulnerable activities. In the prior case, the attack was targeted, while in the case of IO the desire would be to maximize the number of affected individuals. The misinformation would have to be in the activity providers' stores.

A way to accomplish this would be to modify the information in the activity data stores to suit the threat actors' needs. While keeping sophisticated nation state actors from gaining access to a computer network is beyond the means of most organizations, detecting them or their actions is a much more reasonable expectation. Altering large amounts of information would doubtless require a prolonged interactive operation. Implementing best practices for data protection, including extensive logs of information access and modification, would significantly reduce the risk of these activities remaining undetected by content authors and system administrators. In this case, the users who would serve as the strong link would be authors and administrators.

Another way of inserting misinformation would be to target the authors directly or indirectly. A direct means would be to have persons friendly to the threat actor secure employment as content developers. They would then insert the desired content in a way that would be almost impossible to detect through technical means. Alternatively, the threat actor could persuade or coerce legitimate content developers. This would likely not be detected using technical controls either, but alert colleagues could provide early warning. Many government organizations have counter-intelligence programs that aim to identify insider threats. The adoption of the TLA would reinforce the need for these programs in both government and private sector organizations.

Finally, as our misuse cases show, threat actors would be interested in reading TLA information about learners and activities. Whether the actor is a nation state trying to surveil an individual or organization or a criminal trying to sell user data, the personal information within the component systems of the TLA represents a lucrative target to multiple threat actors. This is one class of threats in which we cannot rely on users or content authors for enhanced protection. Already the TLA community is rallying around robust protocols for protecting the confidentiality of learner information within its systems, which will certainly help, but we will rely almost exclusively on systems administrators and security personnel to protect this data at rest.

6 Other Considerations

Apart from the technical and procedural controls we have described for protecting the TLA, there are other considerations that can help protect this environment if properly addressed. Reinforcing the right behaviors among both users and providers can further enable a safe and secure environment that is critical to realizing the promise of the TLA. As the amount of personal data that is stored in networked nodes increases, so must the awareness of the individuals described by that data. Security awareness training programs are intended to help users be aware of threats and how to mitigate them. The TLA has the potential to store very intimate data about its users, which furthers the case for effective security awareness for everyone.

Reducing the amount of personal user data stored in the system naturally creates a tension between functionality and privacy. At this stage in the development of the TLA, the set of requirements that would support a deliberate tradeoff analysis are not specific enough. If the community were to allow these functional requirements to emerge and morph in a naturalistic way, their impacts on privacy would be much more difficult to ascertain. Instead, we propose a deliberate dialog about the tradeoffs that should be considered. This conversation, which has already started informally within the TLA community, should continue as the architecture matures.

7 Conclusion

In this paper, we have presented a detailed threat model for the TLA, together with reasonable controls that could mitigate the risks posed by these threats. While technical controls are always needed in an information system, we have presented procedural counter-measures that can further improve the security of the environment. However, the final and critical layer of protection consists of engaged, aware, and alert users who understand their stake in the process and take appropriate steps to enhance the effectiveness of the security controls that have already been or soon will be built into the TLA.

Acknowledgement

This material is based upon work supported by the Advanced Distributed Learning (ADL) Initiative under Contract No. W911QY-16-C-0019. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Advanced Distributed Learning (ADL) Initiative.

References

1. Swan, Melanie: Emerging patient-driven health care models. International journal of environmental research and public health 6.2, 492--525 (2009)
2. Blackhurst, J. L., J. S. Gresham, and M. O. Stone: The quantified warrior: How DoD should lead human performance augmentation. In: Armed Forces Journal (2012)
3. Kenney, Michael: Cyber-terrorism in a post-stuxnet world. Orbis 59.1, 111--128 (2015)
4. U.S. House of Representatives. Committee on Homeland Security. Hearing on Emerging Cyber Threats to the United States. (2016)
5. Cyber terrorism seen as the BIGGEST single future threat. http://securesense.ca/cyber-terrorism-seen-biggest-single-future-threat
6. U.S. charges Iranians for cyberattacks on banks, dam. http://www.cnn.com/2016/03/23/politics/iran-hackers-cyber-new-york-dam
7. Soros hacked, thousands of Open Society Foundation files released online. https://www.rt.com/usa/355919-soros-hacked-files-released
8. Turns Out You Can't Trust Russian Hackers Anymore. https://foreignpolicy.com/2016/08/22/turns-out-you-cant-trust-russian-hackers-anymore
9. Holt, Thomas J., Olga Smirnova, and Yi-Ting Chua. Data Thieves in Action: Examining the International Market for Stolen Personal Information. Springer, New York (2016)
10. Michigan State University hacked, personal information stolen. http://nbc4i.com/2016/11/18/michigan-state-university-hacked-personal-information-stolen
11. How was your credit card stolen? https://krebsonsecurity.com/2015/01/how-was-your-credit-card-stolen
12. Hackers selling 117 million LinkedIn passwords. http://money.cnn.com/2016/05/19/technology/linkedin-hack
13. Hohlfeld, Oliver, Thomas Graf, and Florin Ciucu: Longtime behavior of harvesting spam bots. In: Proceedings of the 2012 ACM Conference on Internet Measurement (2012)
14. Celebrities' Medical Records Tempt Hospital Workers To Snoop. http://www.npr.org/sections/health-shots/2015/12/10/458939656/celebrities-medical-records-tempt-hospital-workers-to-snoop
15. NSA officers spy on love interests. http://blogs.wsj.com/washwire/2013/08/23/nsa-officers-sometimes-spy-on-love-interests/
16. Ky. Fire Commission investigating Southeast Bullitt Fire Dept. http://www.wdrb.com/story/24643774/ky-fire-commission-investigating-southeast-bullitt-fire-department
17. Report: Tulsa Sheriff's Office falsified training records for reserve deputy who fatally fired gun instead of Taser. https://www.washingtonpost.com/news/morning-mix/wp/2015/04/16/report-tulsa-sheriffs-office-falsified-training-records-for-reserve-deputy-who-fatally-fired-gun-instead-of-taser/?utm_term=.32c56ea0498e
18. Harris, Shon and Maymi, Fernando: CISSP all-in-one exam guide, 7th Edition. McGraw-Hill, Inc., San Francisco (2016)
19. Burkett, Randy. "An alternative framework for agent recruitment." (2013)
20. US DoD: Pub 3-13: Joint Doctrine for Information Operations 9, 1--9 (1998)