

Network Attack Detection and Defense

Manifesto of the Dagstuhl Perspective Workshop
March 2nd – 6th, 2008

Editors:

Georg Carle

Falko Dressler

Richard A. Kemmerer

Hartmut Koenig

Christoph Kruegel

Pavel Laskov

University of Tuebingen, Germany

University Erlangen-Nuremberg, Germany

University of California, Santa Barbara, USA

Brandenburg University of Technology Cottbus, Germany

Technical University of Vienna, Austria

Fraunhofer Institute Berlin, Germany

Dagstuhl Seminar Proceedings 08102

Perspectives Workshop: Network Attack Detection and Defense

<http://drops.dagstuhl.de/opus/volltexte/2008/1491>

Table of contents

Executive Summary.....	2
I. Rationale.....	3
II. Objectives.....	3
III. Deliverables.....	3
IV. Scoping.....	4
V. Research Agenda and Topics.....	5
VI. Conclusions.....	12

Executive Summary

This manifesto is the result of the Perspective Workshop *Network Attack Detection and Defense* held in Schloss Dagstuhl (Germany) from March 2nd – 6th, 2008. The participants of the workshop represent researchers from Austria, France, Norway, the Switzerland, the United States, and Germany who work actively in the field of intrusion detection and network monitoring. The workshop attendee's opinion was that intrusion detection and flow analysis, which have been developed as complementary approaches for the detection of network attacks, should more strongly combine event detection and correlation techniques to better meet future challenges in future reactive security.

The workshop participants considered various perspectives to envision future network attack detection and defense. The following topics are seen as important in the future: the development of early warning systems, the introduction of situation awareness, the improvement of measurement technology, taxonomy of attacks, the application of intrusion and fraud detection for web services, and anomaly detection.

In order to realize those visions the state of the art, the challenges, and research priorities were identified for each topic by working groups. The outcome of the discussion is summarized in working group papers which are published in the workshop proceedings. The papers were compiled by the editors to this manifesto.

I. Rationale

The increasing dependence of human society on information technology (IT) systems requires appropriate measures to cope with their misuse. The growing potential of threats, which make these systems more and more vulnerable, is caused by the complexity of the technologies themselves and by the growing number of individuals that are able to abuse the systems. Subversive insiders, hackers, and terrorists get better and better opportunities for attacks. In industrial countries this concerns both numerous companies and the critical infrastructures, e.g. the health care system, the traffic system, power supply, trade (in particular e-commerce), or the military protection.

Reactive measures comprise beside the classical virus scanner intrusion detection and flow analysis. The development of intrusion detection systems began already in the eighties. Intrusion detection systems possess a prime importance as reactive measures. A wide range of commercial intrusion detection products has been offered meanwhile; especially for misuse detection. The deployment of intrusion detection technology still evokes a lot of unsolved problems. These concern among others the still high false positive rate in practical use, the scalability of the supervised domains, and explanatory power of anomaly-based intrusion indications.

In recent years network monitoring and flow analysis has been developed as a complementary approach for the detection of network attacks. Flow analysis aims at the detection of network anomalies based on traffic measurements. Their importance arose with the increasing appearance of denial of service attacks and worm evasions, which are less efficient to detect with intrusion detection systems. The flow analysis community developed two approaches for high speed data collection: flow monitoring and packet sampling. Flow monitoring aims to collect statistical information about specific portions of the overall network traffic, e.g. information about end-to-end transport layer connections. On the other hand, packet sampling reduces the traffic using explicit filters or statistical sampling algorithms.

II. Objectives

The objective of the Perspective Workshop *Network Attack Detection and Defense* was to discuss future challenges in reactive security, in particular in intrusion detection and flow analysis. New challenges arise as the functionality of network monitoring, attack detection and mitigation must be suitable for a large variety of attacks, and has to be scalable for high data rates and number of flows. Event correlation techniques can be used to combine results from both worlds. The workshop was the first one devoted to this topic in Dagstuhl. A particular objective of this workshop was to bring together both the intrusion detection and network monitoring communities, which still do their research relatively separated and are organized in different communities (e.g. WGs SIDAR and KUVS in the German Society of Informatics (GI) for reactive security and communication systems, respectively). The seminar was supposed to foster the coordination of the research activities in both communities.

III. Deliverables

The outcome of the workshop is a written manifesto, detailing the open issues and possible research perspectives for the coming 5 -10 years according to the objectives given above. The manifesto was compiled by the editors listed at the front page based on the working group papers. Pavel Laskov kindly added a section on anomaly detection. The seminar participants and the composition of the working groups are listed in the appendix.

IV. Scoping

Intrusion Detection

The security function intrusion detection deals with the monitoring of IT systems to detect security violations. The decision which activities have to be considered as security violations in a given context is defined by the applied security policy. Two main complementary approaches are applied: anomaly and misuse detection. *Anomaly detection* aims at the exposure of abnormal system and/or network behavior. It requires a comprehensive set of data describing the normal system and network behavior. Although much research has been done in this area, it is difficult to achieve so that anomaly detection has currently still a limited practical importance. *Misuse detection* focuses on the (automated) detection of known attacks described by patterns, called *signatures*. These patterns are used to identify an attack in an audit data stream. This approach is applied by the majority of the systems used in practice. Their effectiveness, however, is also still limited. Intrusion detection systems are further classified in network- and host-based systems. Network intrusion detection systems analyze the network traffic to find suspicious attack patterns. They have proven to be robust and are preferably applied in today's commercial products. The development of field proven host-based systems seems to be more difficult. Today's solutions are mostly only able to capture simple attacks, especially by matching single step signatures in audit data streams, which have to be generated by special audit functions after a security relevant event took place.

The successful deployment of intrusion detection systems in practice still has to cope with a number of challenges. One problem is the accuracy of the detection models (such as signatures or specifications). When detection models are overly restrictive, false negatives are possible. This is particularly problematic for misuse detection systems that specify the properties of a particular attack. Here, care must be taken that the properties are not too specific and only valid for a very narrow set of instances of the complete class of attacks. When attack models are overly permissive, on the other hand, they will also match benign traffic. This is often the case with anomaly-based systems. A result of matching benign traffic is a large number of false positives. False positives undermine the trust in the intrusion detection system as they often cause lengthy investigations of valid network traffic. A second problem faced by intrusion detection systems is the large number of alerts that they produce. Network packets are at a very low level, and a single attack scenario run by an adversary (which includes scans, brute-force attacks against multiple services, etc) can quickly generate hundreds or even thousands of individual packets that match an attack specification. The result is hundreds or thousands of very similar alerts that actually refer to a single root cause. Alert correlation was proposed to infer high-level attack scenarios from a stream of low-level alerts. Unfortunately, the different alert formats and the difficulty of inferring strategies from low-level events make this problem challenging.

The main merit of anomaly-based intrusion detection techniques is their ability to detect previously unknown attacks. One might think that the collective expertise amassed in the computer security community and a sophisticated infrastructure for dissemination of security-related advice (e.g. vulnerability tracking systems and signature database) rule out major outbreaks of "genuinely novel" exploits. Unfortunately, signs are appearing that a wide-scale deployment of efficient tools for obfuscation, mutation, and simple encryption of attacks generate a huge variability of, strictly speaking, only "marginally novel" but nevertheless undetectable with modern signature-based tools attacks. This reality brings anomaly detection back into the limelight of scientific interest.

Network Monitoring

Network monitoring has become a major building block for various applications in the networking community. Examples range from accounting and charging to attack detection scenarios. The main challenges for network monitoring are the significant bandwidth growth compared to the processing speed of the monitoring probes. Thus, several solutions have been developed that allow reducing the processing requirements for network monitoring and analysis. The primary idea behind all these concepts is to split the monitoring and the subsequent analysis into two independent tasks. Hence, monitoring probes gather and export only the necessary information, keeping the amount of transferred monitoring data at an acceptably low level.

The first concept is the concept of *flow monitoring*. The key idea is to store information about packet flows and the corresponding statistics instead of individual packet information. In this context, a flow is defined as a unidirectional stream of IP packets identified by a common parameters such as the IP-five-tuple (protocol type, source IP address, destination IP address, source port, destination port). Doing this, a single measurement record can contain information of up to several thousand packets. For transmission of monitoring data to a remote analyzer, standard protocols such as Netflow.v9 or IPFIX (IP flow information export) have been developed.

In some kinds of analyses such as intrusion detection, flow statistics do not provide sufficient information about the monitored traffic. Although flow statistics may be used for anomaly detection, signature-based detection schemes usually work on the packet payload which is not contained in the flow records. Therefore, the applicability of flow accounting for intrusion detection is limited. To support the selection of single but complete packets and transporting them to an analyzer, *packet sampling* techniques were developed. They allow the selection of individual packets on the basis of filters and samplers. While filters are used for deterministic packet selection based on matching fields in the IP header, samplers select packets using a given sampling algorithm. The selected packets can be exported to the analyzer using the PSAMP protocol.

V. Research Agenda and Topics

This section presents the visions and research requirements, as seen by the workshop participants, to improve the quality and effectiveness of network attack detection and defense in the next years. They are the result of multiple discussions during the workshop as well as the outcome of separate working groups formed during the workshop. The following topics were discussed in the working groups:

- Early Warning Systems
- Situation Awareness
- Attack Taxonomy
- Measurement Requirements
- Requirements for Network Monitoring from an IDS Perspective
- Intrusion and Fraud Detection for Web Services

Moreover, a topic *Anomaly Detection* was added by the editors. This part was written by Pavel Laskov.

Early Warning Systems

Recently the interest in *Early Warning Systems* arose that inform about suspicious activities in the monitored system or network. Early warning systems aim at detecting unclassified but potentially harmful system behavior based on preliminary indications. They are considered complementary to intrusion detection systems. Both kinds of systems try to detect, identify and react before possible damage occurs and contribute to an integrated and aggregated situation report. A particular emphasis of early warning systems is to establish hypotheses and predictions as well as to generate advises in still not completely understood situations. Thus the term *early* has two meanings: a) to start early in time aiming to minimize damage, and b) to process uncertain and incomplete information.

We see an early warning system consisting of the following process chain:

1. observation of system behavior,
2. pre-classification in order to concentrate on relevant observations,
3. learning a suitable classification framework,
4. applying the learned classification framework on actual observations and evaluation of current system behavior, and
5. triggering appropriate countermeasures.

This process chain is meant as a continuously working pipeline with feedback to adjust preceding steps.

State-of-the-Art

For each step of the process, first proposals are known to the community, but **the** process chain has not been studied yet and needs future research.

Challenges

In order to implement the early warning process a number of challenges have to be addressed. An early warning system has to deal with still unclassified not well-understood but potentially harmful system behavior. Further, a set of countermeasures must be defined which have to be appropriately customized and initiated in order to detect and prevent a threat. Since effective early warning can only be realized by a cooperative approach, the different interests of all involved parties have to be considered. Moreover, non-technical challenges include the establishment of trust relations among parties participating in the early warning system as well as compliance with legislation. How to motivate parties to contribute to an early warning system and what are suitable business models to operate and maintain such a system are additional open questions.

The challenges of early warning systems require the cooperation of different communities especially that of network measurement, machine learning, intrusion detection, and information engineering.

Situation Awareness

The impact of security events on a mission can be understood only if cyber situation awareness is achieved. That is, the infrastructure security officer has to be aware of what is happening in the protected environment, and he/she needs to understand how information, events, and possible actions can impact the overall objectives of the mission. Cyber situation awareness allows the security officer to make critical decisions, such as how an attack is affecting the enclave and/or impacting the mission. Given the answers to these questions the security officer can determine the appropriate course of action (COA) to take.

In terms of computer network security and defense, *situation awareness* (SA) refers to the operational picture that consolidates all available information that is actually needed for identifying attacks and for selecting and applying appropriate countermeasures. In addition to making human beings “situation-aware,” systems and their implementations must also be made aware.

To get the “big picture,” the consolidation of information needs to be performed in an appropriate model that is comprised of resources (or assets), actors, and their inter-dependencies. This model needs to be aligned to the actual mission or high-level objectives of running the computer network and to support different layers of granularity (e.g. from the business process layer of a network operations center’s point-of-view down to low-layer configuration information).

State-of-the-Art

From the IT security community’s point-of-view, there have been many different substantial contributions to obtaining common operational pictures of computer networks. Examples include vulnerability management, intrusion detection, security information & event management, and intrusion response modelling and selection metrics. There has also been related work identified in other areas of interest and other research communities: system & network management, IT service management, multi-sensor data fusion, trust modelling, belief modelling and propagation, modelling of military strategies and business processes, and game theory. The current state-of-the art in these areas needs to be further evaluated in order to estimate synergy benefits that can be expected.

Challenges

One of the biggest challenges in research and deployment is the development of an information consolidating model. There are unsolved questions on how to model assets, actors, and their dependencies, honouring the network’s objectives on different abstraction levels.

Even more challenges arise as soon as the necessary information for creating and updating the situation model is *not available* (e.g., due to missing ability or willingness to share) or *incorrect* (e.g., due to malfunction, misconfiguration, or forgery). Non-available or incorrect information may not only lead to an incomplete operational picture, but also to contradictions that need to be resolved.

Another challenge is how to formulate hypotheses or predictions about detected phenomena and their update (support, rejection) in an iterative fashion. Including potential adversaries introduces multiple magnitudes of complexity to the situation awareness model.

Finally, if it becomes feasible to obtain a comprehensive operational picture and reasonably verified hypotheses about adversaries and their intentions, several new challenges are expected to arise concerning how to actually implement SA in tools and systems (e.g. situation-aware IDS).

Attack Taxonomy

Attack taxonomy attempts to characterize the classes of attacks that can be detected (well) by inspecting network traffic.

State-of-the-Art

Detecting scan activity in network traffic has attracted a lot of interest in the research community over the last few years. As a result, there are several systems and algorithms that can be used to detect either port scans or to identify worm propagation. Moreover, current traffic monitoring techniques are useful to detect effects of attacks in order to identify hosts that have been compromised. For example, a bot that has been installed by a user because of a social engineering attack or with the help of a successful exploit can be detected by monitoring the network for suspicious behavior: such hosts commonly generate either lots of scan traffic, are suspicious due to a large amount of mails sent via these hosts, or generate lots of DNS queries. All such effects can be easily detected with different traffic monitoring strategies. Finally, current monitoring techniques allow us to detect artifacts of attacks. For example, we can detect common attack tools, ready-made exploits, or worms based on specific signatures in the network traffic.

Challenges

We identified four major problems that network-based intrusion detection systems are facing: encrypted network traffic, application-level attacks, performance, and evasion attacks.

1. An obvious problem in this area is payload inspection of encrypted traffic. Since the network-based intrusion detection system (NIDS) commonly has no access to the encryption keys, it cannot decrypt the captured data and, therefore, no analysis is possible.
2. Traditional attack venues, such as buffer overruns or exploits of input validation errors, have been known for a long time and are widely understood. As a result, a large number of defense mechanisms have been devised. For client-side attacks, however, only a few viable defense solutions have emerged so far. A distinctive feature of these client-side attacks is that security problems often cannot be traced to a particular vulnerability that can be easily fixed. That is, the client's security policy is not obviously and immediately violated. Furthermore, sending of code from server to client becomes more and more common (e.g., AJAX sends JavaScript over the network) and this new interaction model poses further challenges, since a NIDS would need to inspect and verify the code. By monitoring network traffic, such attacks are not easy to identify as they occur at the application-level: the NIDS would need to understand the context of requests and also keep track of the application state. That is, one needs to understand the application logic and try to detect attacks, which is hard even given the current network speed.
3. Given the fact that networks are getting faster at a higher pace than processing power is increasing, this is also clearly a problem.
4. Finally, traffic blending attacks and similar evasion attacks pose several challenges for NIDS.

To improve the detection capabilities from a network point of view and to cope with future challenges in this area, we developed some recommendations for future work in this area. An application should support a NIDS such that it becomes easier to check for ongoing attacks. This could, for example, be achieved by developing protocols in such a way that the NIDS can verify – without too much overhead – whether or not a given packet is legitimate.

We require a deeper analysis and metrics to define the complexity of an attack. Here, we understand the complexity of an attack as the difficulty to see this attack at the network level. In particular, attacks that target application-level flaws could be perfectly legitimate from a network perspective, but might arrive at an unexpected point in time or in an unintended order. These facets should be captured by the proposed complexity metric.

An additional area of future work is behavior based detection of attacks: if we can understand the current, normal configuration of a system, then we can detect deviations from this profile as an attack. To achieve this goal, we need to develop algorithms to understand what operations are normal based on the current configuration (e.g. information about network configuration, running services, clients that make use of these services).

Measurement Requirements

In recent years network monitoring and flow analysis has been developed as complementary approach for the detection of network attacks. Flow analysis aims at the detection of network anomalies based on traffic measurements. Their importance arose with the increasing appearance of denial of service attacks and worm evasions which are less efficient to detect with intrusion detection systems.

State-of-the-Art

There are a wide variety of tools available for packet and flow measurement. The MOME database (www.ist-mome.de) lists more than 400 different tools. Configurable measurements are possible to a certain extent. Standardization efforts are in progress (IETF IPPM, IPFIX, PSAMP, etc.).

It is difficult to provide labeled data, i.e. data that is precisely pre-classified, because it is hard to classify data from existing networks in attack/non-attack traffic, i.e., it cannot be decided exactly, whether the traffic includes new attacks or not. On the other hand, it is extremely difficult to generate artificial traces due to the complexity of nowadays networks. Data from honeypots or honeynets may help to collect useful input for training intrusion detection systems.

An ideal measurement system would be the one that captures everything everywhere. With this, attack detection algorithms (e.g. machine learning) could extract the information needed to detect anomalies. Nevertheless, measuring everything everywhere would mean to fully capture each packet at each network node. This is not feasible, especially to not slow down the Internet or to not make its usage much more expensive. Furthermore, current IDS are not capable to handle the high amount of data that would result from full measurements at all nodes.

The main constraints for a measurement system nowadays are:

(1) *Resource Limitations*

Measurements are limited to resources that we can afford for network measurements. Limitations apply for processing power, memory and transmission capacity. The situation gets worse in environments with wireless transmission and small mobile devices.

(2) *Privacy*

Users have a high interest that their privacy is protected. Capturing flow data and especially packet contents (header, payload) clearly violates this. Anonymization techniques have the disadvantage of removing information which might contradict analysis rules. Providers have an interest that others do not gain knowledge about their network and users. As a consequence getting network traces and sharing of data is extremely difficult nowadays.

Challenges

In order to cope with resource constraints we should focus on the following topics:

- *Resource Management*

Methods have to be developed to reduce resource requirements, e.g. by providing improved algorithms for measurement tasks (e.g. for capturing, classification) or by applying smart aggregation and data selection techniques.

- *Data Sharing*
Reference traces should be provided for the research community, in ideal case labeled data. Network operators should be motivated by incentives to share data. Further standardized measurement methods and result representation should be applied to ensure the comparability of measurement results.
- *Privacy-preserving methods*
Methods have to be developed which provide a trade-off between anonymous information representations and applicable analysis demands.
- *Encrypted traffic*
More and more communication is encrypted. Methods have to be developed which meet measurement requirements also for this kind of traffic.

It is further important to gather cross-layer and meta information and associate it with pure measurement results. This includes information about specific events that may influence traffic, measurements from end system and multiple layers.

Moreover, measurement should be an integral part of future network design. A single measurement system for different task (intrusion detection, accounting, SLA validation) should be aspired, since none of the current approaches will meet the various demands sufficiently.

Requirements for Network Monitoring from an IDS Perspective

Detection of malicious traffic is based on its input data, the information that is coming from network-based monitoring systems. Best detection rates would only be possible by monitoring all data transferred over all network lines in a distributed network. Monitoring and reporting this amount of data is feasible in neither today's nor will be in future's systems. Later analysis like stateful inspection of the traffic imposes even more processing costs. But only at this level of monitoring and analysis there may be a chance to capture all attacks inside a system. So there needs to be a trade-off between detection success and the processing costs.

State-of-the-Art

Malicious traffic is mostly generated by compromised systems. Catching attackers during the process of taking over a vulnerable host is complicated, as such attacks only use very low traffic volumes. For higher monitoring data rates, the effects of attacks may be easier monitored like scans for vulnerabilities, large file transfers, or mass spam distribution. Most of these attacks are initiated by script kiddies using downloaded tools, maybe slightly modified. So far, these automated attacks focus on the mass market and do not implement any sophisticated anti-IDS techniques. This fact improves the chance of detecting these attempts dramatically.

Netflow monitoring has become widely accepted as standard to create statistics about network traffic transferred by routers. The IETF defined a protocol to carry flow information over the network called IP flow information export (IPFIX). It includes a standard which allows the transfer of per-packet information like payload data. Standard 5-tuple flow aggregation produces around 8000 flows/s for a 1 GBit/s link. Only few methods are available for mitigating DoS attacks on monitors based on a hash table. In order to enable the processing of high data rates, dynamic reconfiguration of monitors become an issue for attack detection. Currently, only few monitoring systems support seamless reconfiguration without packet losses. To solve the problem in environments with even higher speeds, several methods of packet sampling and filtering are employed.

Challenges

The location of monitoring systems also poses several unsolved challenges: monitoring and analysis can be run in a combined way on end systems. This way, the attack detection is heavily distributed and would enable full payload inspection, although correlation of the analysis results will be difficult in this solution. The more conventional way of placing monitoring systems on the network backbone implies that high data rates only allow a coarse analysis of the data. Detailed analysis would only be feasible for a portion of the traffic. Therefore, we have to work on methods that find the best balance between the massively distributed case and the more centralized high-speed monitoring approach. Such adaptive monitoring techniques would allow focusing attack detection algorithms on selected suspicious data for more detailed inspection.

For an attacker, it is always easier to avoid detection by coarse grained anomaly-based algorithms compared to detailed inspection. We believe that evasion will always be possible unless detection methods analyze full packet payload data and detect anomalies in the semantic content of the application layer. The tradeoff becomes visible in IP and TCP fragmentation issues: connection reassembly does not offer the speeds required for the data rates that occur in backbone networks.

Flow monitoring and analysis techniques, which are operating only on the header information, are often not sufficient for proper attack detection, especially for misuse detection techniques. A trade-off would be the use of flow data accompanied with payload information, like the first N byte of a stream. This solution would be challenging in terms of high-speed operation, but feasible. In addition, hashes and sampling techniques must be developed that avoid overloading monitors in distributed denial-of-service attack scenario.

So far we have only considered plain, unencrypted traffic whose payload may be analyzed directly. But current trends in networking show that the amount of encrypted traffic, tunnels and, in general, overlays is increasing. Monitoring this data introduces more problems: Content is obscured and only statistical features may be used for detection of malicious data. Usage of other networks like 3G networks for mobile devices also increases and those offer entire new types of attacks, e.g. power depletion. The structure of these networks is often fundamentally different from the Internet: the operator has complete control over the network. As only little information is available about those networks, we did not include them in our considerations.

In today's Internet protocol architecture the application layer protocol number (i.e. destination port) is meaningless. Furthermore, the specification of lower protocol layers allows "interesting" protocol use like using only one byte of payload per packet. These limitations can be countered in the future: The introduction of a separate control plane which allows application layer protocol checks by lower layer processing entities would lessen this problem. IANA could assign IDs for protocols that have a defined specification also involving lower layer packet structures. An example of this idea could be that a specific data request (e.g. a HTTP request) must lie inside the first one or two packets of the corresponding connection. Connection reassembly would not be needed any more. The challenge for this problem is it to be designed properly without limiting the flexibility of having stacked protocol layers. The first step of starting this development could be the implementation of application layer checks and lower level protocol restrictions in currently used protocols, so that both approaches remain compatible to each other for easy migration. The next step would be the definition of a new clean-slate approach for future protocols.

Intrusion and Fraud Detection for Web Services

Web services (WS) technology bears the promise to finally bring the power of SOA (Service-oriented Architecture) middleware to the road on a large scale and across organizational domains. Big players such as Google, Amazon, SAP, and IBM have already adopted the technology. European funding agencies are strongly believing and heavily investing into WS-related technological developments and application scenarios. A growing adoption and widespread use of Web services for different application areas can be expected, among them e.g. value added service composition, Web 2.0-enhanced communication systems (e.g. based on Ajax), and focused service offerings from specialized small or medium sized enterprises (SMEs).

State-of-the-Art

More and more technical aspects of WS technology are being standardized, among them standards for security and policy enforcement. Many research groups directly focus on insufficiencies of the existing standards. However, these standards do neither address the availability of Web services, nor intrusion detection, nor fraud detection. A visible trend in the amount of work invested is the application of formal methods for model driven policy generation and verification. Another cluster of interest revolves around the problem of securing service choreography and orchestration. Finally, a lot of work is seen in the area of authorization modeling and enforcement. As a main conclusion we observe that the overwhelming amount of work addresses threat prevention. Reactive aspects have to date been largely ignored, such as detecting and mitigating attacks and frauds at the service layer of the computing system.

Challenges

Clearly, there is an urgent need for methodologies and tools to counter intrusions and fraud geared towards Web services. We state that this need is urgent even though vulnerabilities of Web services are not yet exploited on a large scale. Analyses of potential attacks against Web services revealed that Web services are very vulnerable especially against DoS attacks [2]. On the other hand, Web services and technology not only open a new window for vulnerabilities; they also offer unprecedented opportunities for the detection of intrusions and fraud. The novel idea we propose is to leverage available formal descriptions of system behavior, such as provided formal interface and policy descriptions, to generate models of acceptable behavior and detect deviations thereof by dedicated security services. Concrete approaches might be:

- Extension of WS policy specifications beyond confidentiality and integrity to enable intrusion and fraud detection,
- Transformations of formal WS descriptions into systems for deviation detection to leverage existing specifications,
- Investigation of existing IDS approaches to re-use available technology for the WS environment,
- Design of methods for the integration of service policies across domains to monitor composite services,
- Development of new efficient XML processing methods and algorithms to counter the effects of resource exhaustion attacks.

Thus, detection is comprehended as a part of the life cycle of Web services, influencing the trust evaluation of services and thereby guiding service selection. As service providers realize that their services are used less frequently as a consequence of lax internal and external security, they may implement better safeguards in order to re-gain the trust of the user community.

Anomaly Detection

The main idea of modern anomaly detection methods is to build reliable models of normal behavior and detect deviations thereof. The models can be built at a network or host level, or, generally speaking, for any kind of events observed in a system.

State-of-the-Art

Two main challenges arising in anomaly detection systems are:

1. Dealing with contaminated “normal data”, and
2. Measuring similarity between monitored events.

Both of these challenges have been addressed, although in an abstract way, in the machine learning community. The impact of contamination on normality models can be diminished by regularization, a technique that penalizes overly complex models resulting from trying to accommodate contaminated data. Recently developed similarity-based anomaly detection algorithms offer a possibility to abstract mathematical characterizations of normality behind the notion of similarity. In other words, once a reasonable notion of similarity is defined for pairs of observed events, the algorithms provide a meaningful measure of abnormality.

Challenges

The application of anomaly detection methods in the realm of network intrusion detection still contains many unsolved challenges. The definition and especially the efficient computation of similarity measures for highly structured events arising from network monitoring are quite nontrivial, especially if a detailed protocol analysis beyond byte sequences is desired. Another serious challenge is dealing with “chaff”, unusual but benign events that account for the majority of false alarms. Here some kind of integration of additional knowledge – possibly in the form of small amounts of labeled data – may be necessary. Finally, in order to be recognized in practice anomaly detection has to provide explanations of its predictions to be digested in a wider context of situation awareness.

VI. Conclusions

The increasing dependence of human society on information technology (IT) systems requires appropriate measures to cope with their misuse. The growing potential of threats, which make these systems more and more vulnerable, is caused by the complexity of the technologies themselves and by the growing number of individuals that are able to abuse the systems. Subversive insiders, hackers, and terrorists get better and better opportunities for attacks. In industrial countries this concerns both numerous companies and the critical infrastructures. A major challenge of modern IT security technologies is to cope with an exploding variability of attacks which stems from a significant commercial motivation behind them.

Reactive measures will prove in future as the most efficient mean to ward off these threats. Intrusion detection and network monitoring as complementary approaches will remain the most important approaches in reactive security. They face new challenges caused by an increasing variety of attacks and rising data rates.

Where is the field going? We can identify the following trends:

- The potential of threats in networked systems will further grow as well as the number of individuals which are able to abuse these systems. Increased efforts in research and society are required to protect critical infrastructures such as the health care system, the traffic system, power supply, trade, military networks and others in industrial countries.
- Reactive measures will further turn out as the most efficient countermeasures to ward off these threats in future. Intrusion detection and network monitoring as complementary approaches will remain the most important approaches in this context..
- A careful analysis is necessary to critically review the attacks that can be captured by inspecting different input sources. For example, certain attacks are easier to identify when analyzing network traffic (e.g., scanning, denial of service). Others are more readily visible at the host (such as application-specific attacks). Thus, the community has to invest research effort in analyzing the available data sources and the attacks that can be detected based on these data sources. This is of increasing importance because the frequency of different types of threats is shifting towards client-side and application-level attacks.
- The success and the acceptance of intrusion detection systems essentially depend on the accurateness and the topicality of the applied signatures. Imprecise signatures heavily confine the detection capability of the systems and lead to false positives or negatives, respectively. In order to improve the accuracy of the analyses more attention should be paid to approaches for the systematic design and validation of signatures and attack models.
- An important factor for the success of reactive measurements is the effort and time that is required to create a signature that captures a novel threat. Clearly, it is necessary to quickly generate and deploy signatures because the increasing automation of attack tools significantly shortens the available response time.
- Careful analysis of evasion mechanisms, deployed to undermine detection, will be needed. Numerous practical tools are currently known for subverting anomaly detection, automatic signature generation, and other monitoring instruments. Development of counter-evasion techniques will require a fundamental re-thinking of many detection methods, for example the ones based on machine learning, in order to cope with a potential adversarial impact.
- Given the fact that many attacks affect a large number of distributed hosts, cooperation between sensors and coordination of response mechanisms is of paramount importance. Thus, we require mechanisms that facilitate the exchange of attack-related data in an efficient manner while protecting the privacy of the actors whose information is concerned. This is particularly relevant when data is exchanged between different organizations. Moreover, the analysis power of intrusion detection can be accelerated or improved by moving operations into the monitoring environment. Such cross-domain optimizations need to be adaptive in terms of network load and attack awareness

The workshop results have demonstrated a growing interest to the problems of reactive security in Germany. As the success of reactive measures strongly depends on a seamless interaction between data acquisition and analysis, there is an urgent need to coordinate the research activities of the network measurement and intrusion detection communities.

Appendix A: Workshop Participants List

Joachim Biskup, University of Dortmund	biskup@cs.uni-dortmund.de
Lothar Braun, University of Tuebingen	braunl@informatik.uni-tuebingen.de
Torsten Braun, University of Bern	braun@iam.unibe.ch
Roland Büschkes, RWE AG	roland.bueschkes@rwe.com
Georg Carle, University of Tuebingen	carle@uni-tuebingen.de
Marc Dacier, Institut Eurécom, Sophia-Antipolis Cedex	dacier@eurecom.fr
Hervé Debar, France Telecom R&D - Caen	herve.debar@orange-ftgroup.com
Falko Dressler, University of Nueremberg/Erlangen	dressler@informatik.uni-erlangen.de
Anja Feldmann, Technical University Berlin	anja@net.t-labs.tu-berlin.de
Ali Fessi, University of Tuebingen	fessi@informatik.uni-tuebingen.de
Ulrich Flegel, SAP - Karlsruhe	ulrich.flegel@sap.com
Dirk Haage, University of Tuebingen	haage@informatik.uni-tuebingen.de
Peter Herrmann, Univ. of Science & Technology Trondheim	herrmann@item.ntnu.no
Ralph Holz, University of Tuebingen	holz@informatik.uni-tuebingen.de
Thorsten Holz, University of Mannheim	thorsten.holz@gmail.com
Bernhard Hämmerli, Acris GmbH, Luzern	bmhaemmerli@acris.ch
Marko Jahnke, FGAN - Wachtberg	jahnke@fgan.de
Richard Kemmerer, Univ. California - Santa Barbara	kemm@cs.ucsb.edu
Engin Kirda, Technical University of Vienna	engin@seclab.tuwien.ac.at
Jan Kohlrausch, DFN-CERT Services GmbH	jan_kohlrausch@gmx.de
Christopher Kruegel, TU Wien	chris@cs.ucsb.edu
Hartmut König, BTU Cottbus	koenig@informatik.tu-cottbus.de
Pavel Laskov, Fraunhofer Institute Berlin	laskov@first.fhg.de
Tobias Limmer, University of Nueremberg/Erlangen	tobias.limmer@informatik.uni-erlangen.de
Norbert Luttenberger, University of Kiel	nl@informatik.uni-kiel.de
Michael Meier, University of Dortmund	meier@ls6.cs.uni-dortmund.de
Konrad Rieck, Fraunhofer Institute FIRST Berlin	konrad.rieck@first.fraunhofer.de

Sebastian Schmerl, Brandenburg University of Technology Cottbus
sbs@informatik.tu-cottbus.de

Radu State, INRIA – Nancy
Radu.State@loria.fr

James P. G. Sterbenz, The Univ. of Kansas - Lawrence
jpgs@eecs.ku.edu

Jens Tölle, FGAN/FKIE, Wachtberg – Bonn
toelle@fgan.de

Michael Vogel, Brandenburg University of Technology Cottbus
mvogel@informatik.tu-cottbus.de

Stephen Wolthusen, RHUL - London
Stephen.Wolthusen@rhul.ac.uk

Tanja Zseby, Fraunhofer Institute FOKUS Berlin
tanja.zseby@fokus.fraunhofer.de

Appendix B: Composition of Working Groups

WG Early Warning Systems

Michael Meier (*chair*), Joachim Biskup, Bernhard Hämmerli, Sebastian Schmerl, Jens Tölle, Michael Vogel (*note taker*)

WG Measurement Requirements

Tanja Zseby (*chair*), Lothar Braun, Thorsten Braun, Georg Carle, Falko Dressler, Anja Feldmann, Dirk Haage (*note taker*), Tobias Limmer

WG Situation Awareness

Richard Kemmerer (*chair*), Roland Bueschkes (*vice chair*), Ali Fessi (*note taker*), Hartmut Koenig, Peter Herrmann, Stephen Wolthusen, Marko Jahnke, Hervé Debar, Ralph Holz, Tanja Zseby, Dirk Haage

WG Attack Taxonomy

Christopher Kruegel (*chair*), Marc Dacier, Herve Debar, Thorsten Holz, Engin Kirda, Jan Kohlrausch, Konrad Rieck, James Sterbenz

WG Intrusion and Fraud Detection for Web Services

Ulrich Flegel (*chair*), Marc Daciér, Ralph Holz, Norbert Luttenberger

WG Requirements for Network Monitoring from an IDS Perspective

Falko Dressler (*chair*), Lothar Braun, Thorsten Holz, Engin Kirda, Jan Kohlrausch, Christopher Kruegel, Tobias Limmer, Konrad Rieck, James Sterbenz