The problem of NIDS evasion in mobile networks

Michele Colajanni, Luca Dal Zotto, Mirco Marchetti, Michele Messori
Department of Information Engineering
University of Modena and Reggio Emilia
Modena, Italy
{michele.colajanni, luca.dalzotto, mirco.marchetti, michele.messori}@unimore.it

***Abstract*—This paper presents a novel NIDS evasion strategy that allows attackers to exploit network mobility to perform attacks undetectable by modern NIDSs. Mobility-based NIDS evasion works by combining traditional evasion techniques and node mobility. It represents a generally applicable evasion strategy that works on several protocols for node mobility, and it is effective against state-of-the-art and well configured signature-based NIDSs. We describe three evasion scenarios based on node mobility, and demonstrate the practical applicability of the proposed evasion strategy through a proof of concept attack in a realistic network environment. We conclude the paper by presenting some ideas addressing mobility-based NIDS evasion.**

***Index Terms*—network intrusion detection; mobility-based NIDS evasion;**

I. Introduction

We are witnessing a steady increase in both the number and the computational power of Internet-enabled mobile devices, such as smartphones, PDAs, Internet tablets and netbooks. While traditional IPv4 networks do not provide native support for transparent node mobility, several technologies and protocols exist that allow a mobile node to roam among different networks (or network segments) without interrupting established connections. Relevant examples are Mobile IPv4 [1]–[3] and the mobility extensions to IPv6 [4] at the network layer, as well as protocols for layer-2 handover across wireless access points [5], [6]. Transparent node mobility is required in order to allow nodes to roam seamlessly, and their use will likely grow in the near future. However, transparent node mobility can reduce the effectiveness of widely deployed network security technologies that have not been designed to cope with mobile nodes.

In this paper we focus on the detrimental effects that node mobility has on *Network Intrusion Detection Systems* (NIDSs). NIDSs represent a valuable and widely deployed technology whose effectiveness can be reduced by the diffusion of mobile protocols and devices. We present a new attack strategy, called *mobility-based NIDS evasion*, that allows a malicious user to exploit node mobility to perform stealth network attacks, undetectable by the state-of-the-art NIDSs having stateful capabilities.

We remark that mobility-based NIDS evasion does not represent an exploit against a bug of a specific NIDS implementation, and is not related to a design flaw of a protocol. The proposed evasion strategy can be easily adapted to leverage several network protocols and technologies for node mobility. Moreover, since it prevents the analysis of part of the attack payload, mobility-based NIDS evasion is effective against all stateful signature-based NIDSs. To the best of our knowledge, this is the first paper that describes how node mobility can be exploited by an attacker to perform stealth network attacks undetectable by modern NIDSs.

The applicability and effectiveness of the proposed evasion technique is demonstrated through a proof of concept stealth attack, carried out against a real network environment.

Section II summarizes the background by describing the stateful approaches to network intrusion detection, as well as the evasion techniques based on packet fragmentation from which mobility-based evasion evolves. Section III describes how an attacker can leverage node mobility to perform stealth attacks. A proof of concept of the new NIDS evasion technique proposed in this paper is described in Section IV. Section V outlines a viable solution to this problem based on cooperation among distributed NIDSs. Section VI highlights the main contribution of this paper with respect to previous works in the fields of NIDS evasion and distributed architectures for intrusion detection. Section VII draws the conclusions and outlines future works.

II. Background

Without loss of generality, it is possible to model any NIDS as a black box whose input is represented by an ordered and finite sequence of network packets $P = \{p_1, p_2, ..., p_n\}$, where each p_i represents the i_{th} network packet, and n is the total number of packets analyzed by the NIDS. All modern NIDSs are *stateful*, meaning that they are able to build and maintain *state* information extracted from the analyzed network traffic.

Instead of analyzing each network packet on its own (as early *stateless* NIDS did), the information contained within a network packet and relevant to the intrusion detection algorithm is used to create and update the NIDS state. As an example, in Snort [7] (a well known and widely deployed stateful NIDS) the state includes information about all the active transport level connections. For each connection the *Stream5* snort pre-processor maintains several metadata and two ordered lists of payloads (one for each direction of the communication) exchanged by the connection endpoints. The detection algorithm is then executed on the current state information, rather than on the information that can be extracted from a single packet. Hence, even if none of the individual p_i contains enough information to detect an intrusion attempt, it is still possible for a stateful NIDS to detect an intrusion

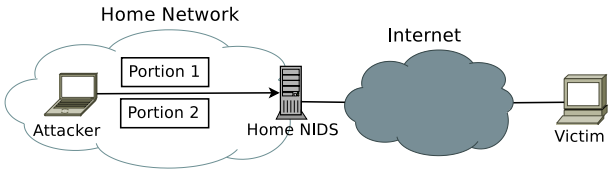

Fig. 1. NIDS evasion through attack fragmentation

by correlating information extracted from several different packets.

Several techniques [8]–[14] for evading NIDSs exist (see Section VI for a comparison between mobility-based evasion and traditional evasion techniques). Of particular relevance to this paper is NIDS evasion carried out through payload fragmentation, a well known technique that is effective against older stateless NIDSs but not against the state-of-the-art NIDSs with stateful signature-based capabilities.

We describe an example of fragmentation-based NIDS evasion through the scenario shown in Figure 1. The node labeled as *Attacker* is connected to its Home Network and is trying to exploit a remote vulnerability of the node *Victim*. We assume that all the network traffic that flows to and from the Home Network is analyzed by a NIDS (the Home NIDS in Figure 1). We also assume that the Attacker is trying to evade detection by fragmenting the attack payload.

Signature-based NIDSs work by comparing the analyzed network traffic against a known set of attack signatures. Each signature represents the fingerprint of an intrusion attempt, and is usually modeled as a regular expression. Any attacker with an approximate knowledge of the signature applied by a NIDS to model an intrusion attempt can identify the portion of the attack payload that matches the corresponding signature. This payload portion is then split into two or more portions (*Portion 1* and *Portion 2* in Figure 1) conveyed through separate network packets. This result can be achieved through several techniques, such as IP packet fragmentation or fragmentation of the payload into two different not-fragmented TCP packets.

Each of these network packets contains just a portion of the attack payload that does not match with the regular expression applied by the Home NIDS. Let $\{p_j\}, 1 \leq j < m$, denote m different packets, each crafted by an attacker to contain only a portion of the malicious payload. After having updated the state with the information extracted from all these m packets, the complete attack payload is contained in the state of the NIDS. Since the detection algorithm is executed on the state, rather than on the features extracted from a single packet, a stateful Home NIDS is able to detect the attack. We can conclude that all modern NIDSs are not vulnerable to evasion through payload fragmentation. In the following section we demonstrate that this positive conclusion does not hold anymore if node mobility is supported.

III. Mobility-based NIDS evasion

In this section we describe a novel evasion strategy, defined as *mobility-based NIDS evasion*, that an attacker can exploit

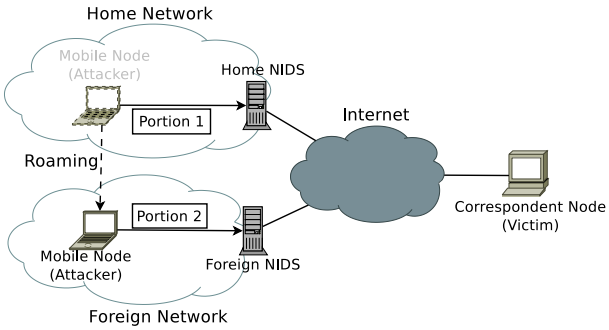

Fig. 2. First mobile-evasion scenario: mobile attacker

to perform "stealth" attacks, undetectable by state-of-the-art stateful NIDSs. The main idea behind mobility-based evasion is to coordinate node mobility and traditional NIDS evasion techniques [8] based on fragmentation of the attack payload. In particular, we describe how it is possible to exploit mobility-based evasion in three different attack scenarios, presented in sections III-A, III-B and III-C, respectively.

We remark that the following scenarios are independent of the technology that allows mobile nodes to roam across different networks. The only requirement is for the roaming process to be transparent, meaning that active transport level connections are not interrupted by the handover process. This assumption is satisfied by several technologies and network protocols, such as IPv6 [4], Mobile IPv4 [1]–[3] and protocols for layer-2 handover across wireless access points [5], [6].

A. Mobility-based NIDS evasion: mobile attacker

Let us consider a simple network scenario, represented in Figure 2, in which two nodes are communicating through the Internet. We consider a Mobile Node, installed in a network that allows node mobility, communicating with a Correspondent Node. No assumption is made on the nature of the Correspondent Node. It can be a fixed Internet node, as well as another mobile node. We assume that both the Home Network and the Foreign Network include a stateful NIDS (Home NIDS and Foreign NIDS, in Figure 2) that monitors all the Internet traffic coming to and from the monitored network, including the traffic generated and received by the Mobile Node. In this first example, the Mobile Node (Attacker) aims to exploit a remote vulnerability of the Correspondent Node (Victim) by sending network packets containing a malicious payload. Moreover, the Mobile Node is trying to evade detection by separating the attack in two different portions (*Portion 1* and *Portion 2*, in Figure 2), each transmitted in a separate network packet. Since the fragmentation of IP packets is discouraged in the IPv6 protocol [15] and easily detected by modern NIDSs as an anomalous network activity, the attacker sends the two attack portions inside two not-fragmented TCP packets having consecutive sequence numbers.

The sequence of activities performed by the attacker is as follows:

1) the Mobile Node sends the first attack portion;

2) the Mobile Node roams to the Foreign Network;
3) the Mobile Node sends the second (last) attack portion.

Similarly to the case presented in Figure 1, in which mobility was not involved, the first portion of the attack is sent through the Home Network, and is intercepted and analyzed by the Home NIDS. Being only a portion of the attack, the Home NIDS updates its state information, but it does not have enough information to detect an intrusion attempt.

The second portion of the attack is sent from the Foreign Network to the Correspondent Node. It is intercepted by the Foreign NIDS that has not received the previous portion and does not have the state information necessary to reconstruct and recognize the whole attack. This state information is possessed by the Home NIDS, but it is useless, since the Home NIDS does not receive the second attack portion.

As a result, none of the two NIDSs deployed in the Home and Foreign Networks can detect the attack. This inability is not the result of a bug in a NIDS implementation. Since none of the two NIDSs receive the whole malicious payload, intrusion detection by either of the two isolated NIDSs is impossible. Hence node mobility allows the malicious Mobile Node to perform a stealth attack, that would have been easily detected if mobility had not been supported.

Depending on how node mobility is implemented, only a stateful NIDS installed in the Correspondent Node's network may be able to detect the intrusion attempt. In any case, both the Home and the Foreign Network infrastructures can be exploited by a mobile attacker to damage third parties, without the network administrators being able to prevent, stop, or even detect the attack.

B. Mobility-based NIDS evasion: mobile victim

With respect to the case study presented in Section III-A, in this scenario the roles are reversed: the Correspondent Node is the Attacker that leverages mobility to evade detection, while the Mobile Node is the Victim. We use Figures 3 and 4 for describing this new type of evasion.

We assume that the Foreign Network and the Home Network are monitored by two stateful NIDSs. We also assume that the Correspondent Node (Attacker) knows when the Mobile Node (Victim) roams across different networks. Depending on the topology of the involved networks, and on how mobility is implemented, the attacker can use several strategies to gain this knowledge. For example, if Mobile IPv4 is used the round trip time between the Correspondent and the Mobile Node is bound to increase after the Mobile Node roams from the Home Network to a Foreign Network. The additional delay is a direct consequence of how triangular routing works. It cannot be eliminated, and it is easily measurable by the attacker. In the case of IPv6, the Attacker can detect any mobility event of a victim to which it is already connected by analyzing the Mobile IPv6 control messages exchanged among the Mobile Node, the Correspondent Node and the Home Agent [4].

Moreover, it may be possible for the Attacker to exploit some previous knowledge about the victim's behavior. If the attacker knows that the Mobile Node roams frequently among different networks, it is sufficient to wait for a given amount of time before sending the last attack portion.

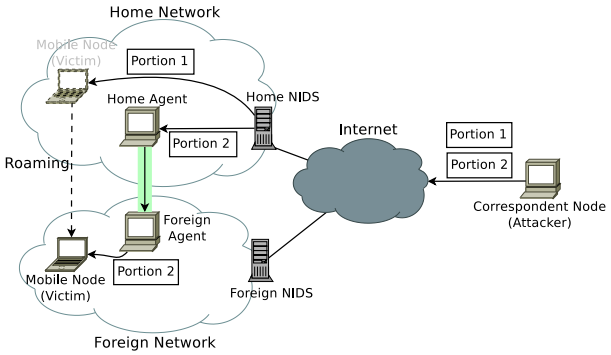

Fig. 3. Second mobile-evasion scenario: mobile victim - triangular routing

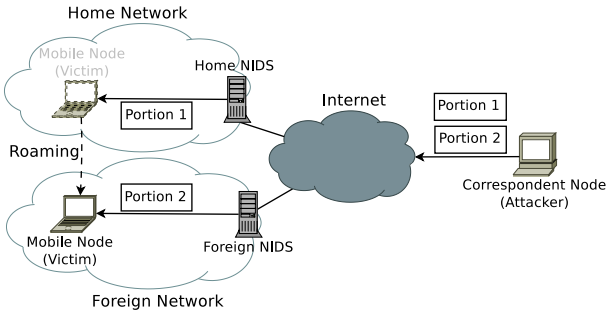

Fig. 4. Second mobile-evasion scenario: mobile victim - route optimization

The attack proceeds as follows:

1) the Correspondent Node sends the first attack portion;
2) the Correspondent Node waits for the Mobile Node to roam to the Foreign Network;
3) the Correspondent Node sends the second attack portion.

Depending on how mobility is implemented, the second attack fragment may or may not be routed to the Mobile Node through the home network. We distinguish between two cases: triangular routing and route optimization.

1) Triangular routing: If Mobile IPv4 is used with the default triangular routing scheme, the two attack portions follow the paths shown in Figure 3. The first attack portion (*Portion 1*, in Figure 3) is sent directly to the Mobile Node, hence it is received and analyzed by the Home NIDS. The second attack portion (*Portion 2*, in Figure 3) is sent by the Correspondent Node after the Mobile Node roamed to the Foreign Network, and is routed to the Mobile Node through the Home Agent. In this setup, the Home NIDS is able to inspect both the attack portions, and to detect the attack. On the other hand, the Foreign NIDS can only analyze the second attack portion, thus being unable to detect the intrusion attempt. As a result, the Mobile Node has been compromised while connected to the Foreign Network, without any chance for the network administrator to detect the attack.

2) Route optimization: The situation is worse if route optimization (the suggested scheme for IPv6 mobility) is

implemented. In this case, the second attack portion flows directly from the Correspondent Node (Attacker) to the Mobile Node (Victim), as shown in Figure 4. Hence neither the Home NIDS nor the Foreign NIDS are able to receive enough information to detect the intrusion attempt.

C. Mobility-based NIDS evasion: mobile target and attacker

The last mobility-based evasion scenario is a combination of the two scenarios described in Sections III-A and III-B. In this case, both the Attacker and the Victim are mobile nodes. We assume that all the four networks involved in this scenario (Victim's Home Network, Victim's Foreign Network, Attacker's Home Network and Attacker's Foreign Network) are monitored by stateful NIDSs, and that the Attacker knows when the Victim roams to the Foreign Network.

The attack proceeds as follows:

1) the Attacker sends the first attack portion;
2) the Attacker waits for the Victim to roam to the Foreign Network;
3) the Attacker roams to a (different) Foreign Network;
4) the Attacker sends the second attack portion.

Note that the order of steps 2 and 3 can be inverted without changing the attack outcome.

Since the Attacker roams to a Foreign Network before sending the second attack fragment, none of the two NIDSs deployed in the Attacker's Home and Foreign Networks receive both the attack portions, as shown in Figure 2. Hence, they are not able to detect the attack, as described in Section III-A.

Moreover, since the Victim is also roaming, the detection ability of the stateful NIDSs that monitor the Victim's Home and Foreign Networks is reduced as described in Section III-B. Again, we distinguish between node mobility implemented through triangular routing and route optimization.

1) Triangular routing: If triangular routing is implemented, the second attack fragment is forwarded to the Victim through the Home Agent, as shown in Figure 3. In this case the Home NIDS of the Victim receives both the attack portions, and it is able to detect the attack. However, it is worth to note that only one out of four stateful NIDSs is able to detect an attack that would have been easily detected without mobility.

2) Route optimization: If route optimization is in place, the two attack portions are routed directly to the Victim, as shown in Figure 4. In this situation neither the Victim's Home NIDS nor the Victim's Foreign NIDS receive both the attack portions, thus being unable to detect the attack. Hence, none of the four stateful NIDSs can detect the attack.

IV. Proof of concept

We demonstrate the applicability of mobility-based NIDS evasion through a proof of concept attack implementing the scenario described in Section III-B on an IPv6 network.

Our experimental testbed comprises two IPv6 networks, the Home Network and the Foreign Network. Both contain a GNU/Linux host, connected to a wireless access point. Each host (Home Agent and Foreign Agent for the Home and Foreign Network, respectively) act as mobility agent and as NIDS for its network. Both NIDSs are instances of Snort 2.8.6.1 (the latest stable version at the moment of writing). Support for IPv6 and node mobility is provided by the Linux kernel (version 2.6.34 on both the Home and Foreign Agent), and the *radvd* [16] and *mip6d* software daemons. *Radvd* is used to advertise network prefixes, while *mip6d* handles Mobile IPv6 control messages and enables route optimization, thus providing higher performance to mobile nodes. Home Agent and Foreign Agent are connected to a network switch, that provides connectivity to a third machine used as Correspondent Node. The Correspondent Node runs the Linux kernel 2.6.34 and the *mip6d* daemon. The wireless access points are two Cisco Aironet 1100 and allow the Mobile Node to roam between the Home and the Foreign Networks by associating to one of the two access points. The Mobile Node is a laptop provided with a wireless network interface, it runs the 2.6.35 Linux kernel and the *mip6d* daemon.

We test the network environment by verifying that the Mobile Node can roam between Home and Foreign Networks without interrupting open TCP connections. We also test the ability of the Home and Foreign NIDSs to detect fragmented attack payloads. For the sake of simplicity, we chose an easily readable Snort signature that is used to detect NOOP sled commonly used as part of shellcodes (signature id: 1394). This signature matches any string composed of more than 31 consecutive "A". We use *netcat6* [17] to open a TCP connection to the Mobile Node and send the malicious payload, composed of 40 "A". We try to evade detection by splitting this payload into two portions, each containing 20 "A". These two portions are sent as the payload of two consecutive TCP packets, without using IP packet fragmentation. We execute this experiment two times. In the first run the Mobile Node is permanently connected to the Home Network. Since Snort is stateful, the Home NIDS is able to reassemble the malicious payload and to detect the attack. In the second run, the Mobile Node is permanently connected to the Foreign Network, and the attack is successfully detected by the Foreign NIDS. As expected, without mobility it is not possible to evade a modern NIDS through fragmentation of the malicious payload.

To demonstrate the effectiveness of mobility-based NIDS evasion, we test the scenario described in Section III-B. At the beginning of the experiment, the Mobile Node is associated to the Home Wireless Access Point. Its IPv6 address (2001:db8::beef) belongs to the Home Network's address space (2001:db8::/64).

The Correspondent Node starts the attack by using *netcat6* to open a TCP connection to the Mobile Node, and to send the first half of the attack, consisting of 20 "A". However, since this payload does not match any signature, no alert is raised. The Correspondent Node then starts to ping the Mobile Node and waits for the control messages used by the Mobile IPv6 protocol to initiate route optimization.

We then simulate a roaming event by using *iwconfig* to associate the wireless NIC of the Mobile Node to the Foreign Wireless Access Point. The Mobile Node receives network advertisements from the *radvd* daemon deployed in the For-

eign Agent, notifying that it is now connected to the Foreign Network (2001:db8:1::/64). In compliance with the IPv6 protocol specifications, the Mobile Node uses its MAC address to generate a *Care-of Address* (2001:db8:1:0:218:deff:fe25:599) that is unique and that belongs to the Foreign Network. It then issues *Binding Update* messages to the Home Agent and to the Correspondent Node to notify them of its new network address.

After receiving the *Binding Update* from the Mobile Node, the Correspondent Node sends the last portion of the attack over the same TCP connection used to send the first attack fragment. Since route optimization is implemented, the second attack portion is routed directly to the Mobile Node and is analyzed by the Foreign NIDS, whose state does not contain the first attack fragment. As a result, the Mobile Node receives the complete malicious payload, while the two stateful NIDSs that monitor the Home and Foreign Networks are unable to recognize the attack. During the experiment, all the network traffic as seen by the Home Agent, the Foreign Agent, the Mobile Node and the Correspondent Node have been recorded in four different network traffic traces, in `.pcap` format. These traces are available for download at http://cris.unimore.it/MobSec.

V. Solution through NIDS cooperation

The problem of mobility-based NIDS evasion is caused by the fragmentation of relevant state information among geographically distributed NIDSs, deployed within independent networks. This fragmentation prevents modern, stateful NIDSs to build a complete state, thus exposing them to the same evasion strategies that were effective only against obsolete stateless NIDSs. To address this issue we propose a cooperative solution, based on the exchange of state information among distributed NIDSs.

Our proposal is to extend the mobility protocols that allow a mobile node to roam by adding three main steps:

1) extraction and serialization of all the state information that are relevant to the Mobile Node from the NIDS that monitors the network that the Mobile Node left;
2) transmission of the serialized state to the NIDS that monitors the destination network;
3) deserialization and insertion of the transmitted state information within the state of the NIDS that monitors the destination network.

In this cooperation scheme, the state information that is related to a Mobile Node "follows" the Mobile Node in the new network, thus preventing an attacker to exploit mobility to evade detection.

Similar cooperation schemes have already been proposed in the context of parallel NIDS architectures [18]–[20] to allow the exchange of state information among NIDSs deployed within the same network, each analyzing a small fraction of traffic gathered from the same network link. However, state exchange among geographically distributed NIDSs poses several new challenges.

First of all, different networks may leverage heterogeneous NIDSs, whose internal state representations are not compatible. Hence, new protocols and standards need to be designed. Moreover, it is necessary to establish trust relationships among cooperative NIDSs, as well as mechanisms to provide confidentiality, authentication and non repudiation of exchanged state information. Finally, the delays related to state management operations need to be compatible with live analysis of network traffic, and the state migration process has to be robust with respect to network delays. The design of a distributed and cooperative NIDS that is able to address all these issues is still ongoing.

VI. Related work

NIDS evasion is a well known attack technique in the network security literature. The seminal work in this area is [8], and its results are extended in [9], [12]–[14] describing several evasion techniques.

The simplest and oldest class of evasion techniques is based on the fragmentation of the attack payload. For example, it can be implemented by splitting the malicious payload across several packets, or by dividing a single packet into several IP fragments. This attack is effective only against a stateless NIDS that does not reassemble network packets. This weakness is the main reason that motivated the design of more complex, stateful NIDS architectures.

With respect to a stateless NIDS, the main improvement of a stateful NIDS is the ability to reconstruct a complete, reordered information flow by reassembling several fragmented and possibly out-of-sequence network packets. In networks that do not provide support for node mobility, we can assume that a NIDS is able to analyze all the traffic that an attacker sends to its victim. Hence, a stateful NIDS receives all the information that is needed to detect the attack. Attackers can still evade detection by a stateful NIDS through several known strategies that exploit misconfigurations or bugs in the packet reassembly algorithms, such as partially overlapping and out-of-sequence network packets, or time-to-live manipulations. If the algorithms used by a stateful NIDS to reassemble network traffic were flawless, this NIDS would be immune to the evasion strategies known in literature.

On the other hand, the transparent node mobility opens new possibilities for an attacker. Since network nodes are free to roam without interrupting open connections, the assumption that a NIDS receives all the traffic that the attacker uses to compromise the victim does not hold anymore. An attacker can then leverage the evasion technique based on attack fragmentation and combine it with node mobility to prevent a NIDS from receiving all the fragments of the malicious payload. Even a stateful NIDS with an ideal packet reassembly algorithm is vulnerable to mobility-based evasion, hence mobility-based evasion clearly differentiates from all the NIDS evasion strategies previously described in literature. To the best of our knowledge, this is the first paper that describes how node mobility can be exploited for novel forms of NIDS evasion.

To defeat mobility-based NIDS evasion we advocate a cooperative approach, in which state information is exchanged by cooperative NIDS architectures whenever a mobile node roams to a different network. Several distributed NIDS systems that cooperate by exchanging alerts have been proposed in literature [21]–[25]. However, this form of cooperation cannot prevent mobility-based evasion because it does not include exchange of state information [19], [20]. This possibility is explored in the context of parallel NIDS architectures [18], [26], but not yet applied to distributed NIDS architectures. Hence, no distributed architecture for intrusion detection proposed in literature can solve the problem of mobility-based NIDS evasion highlighted in this paper.

VII. Conclusion

In this paper we describe a new attack strategy, called mobility-based NIDS evasion, that attackers can use to perform stealth network intrusions, undetectable by state-of-the-art NIDSs. The main idea behind mobility-based NIDS evasion is to combine the fragmentation of a malicious payload (a well known NIDS evasion strategy, ineffective against any modern NIDS) and node mobility. Hence, the relevance of mobility-based NIDS evasion is bound to grow as the number of Internet-enabled mobile devices increases. We remark that NIDS evasion is not the result of design or implementation flaws in a network protocol or in a NIDS implementation; it is an attack strategy that takes advantage of node mobility to prevent stateful NIDS from building the state information that is necessary to detect an attack. Hence, mobility-based NIDS evasion is applicable to several protocols for network mobility, and is effective against all stateful signature-based NIDSs. We described three different scenarios in which an attacker can leverage mobility-based NIDS evasion, and we demonstrated its applicability through a proof of concept attack carried out against an IPv6 network.

Future works will focus on the design of cooperative NIDS architectures that are able to cope with mobility-based NIDS evasion by exchanging state information. We will also explore the applicability of mobility-based NIDS evasion in new contexts, such as live migration of virtual machines.

Acknowledgments

This research has been funded by the IST-225407 EU FP7 project CoMiFin (Communication Middleware for Monitoring Financial Critical Infrastructures).

References

[1] C. E. Perkins, "Mobile networking through mobile ip," *IEEE Internet Computing*, vol. 2, no. 1, pp. 58–69, January 1998.

[2] ——, "Mobile ip," *IEEE Communications Magazine*, vol. 40, no. 5, pp. 66–82, May 2002.

[3] ——, "Mobility support for ipv4," *Request For Comments 3344, Internet Engineering Task Force*, August 2002.

[4] D. Johnson, C. E. Perkins, and J. Arkko, "Mobility support in ipv6," *RFC 3775, Internet Engineering Task Force*, June 2004.

[5] "Ieee standard for information technology-specific requirements part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications amendment 2: Fast basic service set (bss)," *IEEE Std 802.11r-2008*, jul. 2008.

[6] "Cisco Fast Secure Roaming Application Note. Available online at http://www.cisco.com/en/US/products/hw/wireless/ps4570/prod_technical_reference09186a00801c5223.html."

[7] "Snort home page, available online at http://www.snort.org."

[8] T. H. Ptacek and T. N. Newsham, "Insertion, evasion, and denial of service: Eluding network intrusion detection," Secure Networks, Inc., Suite 330, 1201 5th Street S.W, Calgary, Alberta, Canada, T2R-0Y6, Tech. Rep., 1998.

[9] Y. Yoon, J. Y. Oh, and Y. M. Yoon, "Nids evasion method named SeolMa," *Phrack Magazine*, vol. 0x0b, no. 0x39, June 2001.

[10] U. Shankar and V. Paxson, "Active mapping: Resisting nids evasion without altering traffic," in *Proceedings of the 2003 IEEE Symposium on Security and Privacy (SP 03)*. Washington, DC, USA: IEEE Computer Society, 2003, p. 44.

[11] R. Smith, C. Estan, and S. Jha, "Backtracking algorithmic complexity attacks against a nids," Dec. 2006.

[12] S. Siddharth, "Evading nids, revisited. available online at http://www.symantec.com/connect/articles/evading-nids-revisited," 2005.

[13] M. Handley, V. Paxson, and C. Kreibich, "Network intrusion detection: Evasion, traffic normalization, and end-to-end protocol semantics," in *Proc. of the 10th conference on USENIX Security Symposium (SSYM'01)*, Nov. 2001.

[14] V. Paxson and M. Handley, "Defending against nids evasion using traffic normalizers," in *Proc. of the Second International Workshop on the Recent Advances in Intrusion Detection (RAID99)*, Nov. 1999.

[15] R. H. S. Deering, "Internet protocol, version 6 (ipv6) specifications," *Request For Comments 2460, Internet Engineering Task Force*, December 1998.

[16] P. Savola, "radvd homepage." [Online]. Available: http://www.litech.org/radvd/

[17] "Netcat6." [Online]. Available: http://www.deepspace6.net/projects/netcat6.html

[18] M. Colajanni and M. Marchetti, "A parallel architecture for stateful intrusion detection in high traffic networks," in *Proc. of the IEEE/IST Workshop on "Monitoring, attack detection and mitigation" (MonAM 2006)*, Tuebingen, Germany, September 2006.

[19] R. Sommer and V. Paxson, "Exploiting independent state for network intrusion detection," in *Proc. of the 21st Annual Computer Security Applications Conference*, Tucson, AZ, USA, December 2005.

[20] M. Colajanni, D. Gozzi, and M. Marchetti, "Enhancing interoperability and stateful analysis of cooperative network intrusion detection systems," in *Proc. of the ACM/IEEE Symposium on Architectures for Networking and Communication Systems (ACM/IEEE ANCS 2007)*, Orlando, FL, USA, Dec. 2007.

[21] "Prelude hybrid intrusion detection system, available online at http://www.prelude-ids.org/."

[22] P. A. Porras and P. G. Neumann, "Emerald: Event monitoring enabling responses to anomalous live disturbances," in *In Proceedings of the 20th National Information Systems Security Conference*, 1997, pp. 353–365.

[23] V. Yegneswaran, P. Barford, and S. Jha, "Global intrusion detection in the domino overlay system," in *Proc. of the ISOC Symposium on Network and Distributed Systems Security*, Feb. 2004.

[24] R. Janakiraman, M. Waldvogel, and Q. Zhang, "Indra: A peer-to-peer approach to network intrusion detection and prevention," in *Proc. of the 12th IEEE International Workshops on Enabling Technologies*, Linz, Austria, Jun. 2003.

[25] M. Marchetti, M. Messori, and M. Colajanni, "Peer-to-peer architecture for collaborative intrusion and malware detection at a large scale," in *Proc. of the 12th Information Security Conference (ISC 2009)*, Pisa, Italy, September 2009.

[26] M. Vallentin, R. Sommer, J. Lee, C. Leres, V. Paxson, and B. Tiereney, "The NIDS cluster: Scalable, Stateful Network Intrusion Detection on Commodity Hardware," in *Proc. of the International Symposium on Recent Advances in Intrusion Detection (RAID)*, September 2007.