OpenFlow Random Host Mutation: Transparent Moving Target Defense using Software Defined Networking

Jafar Haadi Jafarian, Ehab Al-Shaer, Qi Duan
Department of Software and Information Systems
University of North Carolina at Charlotte
Charlotte, NC, USA
{jjafaria, ealshaer, qduan}@uncc.edu

Static configurations serve great advantage for adversaries in discovering network targets and launching attacks. Identifying active IP addresses in a target domain is a precursory step for many attacks. Frequently changing hosts' IP addresses is a novel proactive moving target defense (MTD) that hides network assets from external/internal scanners. In this paper, we use OpenFlow to develop a MTD archi- tecture that transparently mutates IP addresses with high unpredictability and rate, while maintaining configuration integrity and minimizing operation overhead. The presented technique is called OpenFlow Random Host Mutation (OF-RHM) in which the OpenFlow controller frequently assigns each host a random *virtual* IP that is translated to/from the *real* IP of the host. The real IP remains untouched, so IP mutation is completely transparent to end-hosts. Named hosts are reachable via the virtual IP addresses acquired via DNS, but real IP addresses can be only reached by authorized entities. Our implementation and evaluation show that OF-RHM can effectively defend against stealthy scanning, worm propagation, and other scanning-based attack.

Categories and Subject Descriptors

C.2.3 [**Computer-Communication Networks**]: Network Operations

IP mutation, software defined networking (SDN), Moving target defense (MTD), Security

1. INTRODUCTION

Static assignment of IP addresses gives adversaries significant advantage to remotely scan networks and identify their targets accurately and quickly. Scanning tools and worms usually send probes to random IP addresses in the network in order to discover their targets. When a target responds, it can then be identified and attacked. Otherwise the probed addresses will be considered unused. Despite firewall deployment, most enterprise networks have many public and private hosts accessible from outside. The IP address assignment scheme can become more dynamic by using approaches based on DHCP or NAT, but they are insufficient to pro- vide proactive countermeasure because the IP mutation is infrequent and traceable.

In this paper, we introduce a moving target technique called *OpenFlow Random Host Mutation* (OF-RHM) which mutates IP addresses of end-hosts randomly and frequently so that the attackers' premises about the static IP assignment of network fails. OF-RHM has two main objectives. Firstly, the IP mutation must be transparent to the end- host. To provide transparency, OF-RHM keeps the actual or real IP addresses of hosts (called *rIP*) unchanged, but associates each host with random short-lived virtual IP ad- dresses (called *vIP*) at regular intervals which are translated to rIPs right before the host. Secondly, the IP mutation must be performed with high unpredictability and speed to maximize the distortion of attackers' knowledge about the network and increase deterrence of attack planning. To optimize IP mutation with respect to unpredictability and speed, the mutant vIPs are selected randomly from the entire unused address space in the network. The unused address ranges must be assigned to hosts such that it satisfies several constraints including mutation unpredictability and minimum required mutation rate of all hosts. We formulate this problem as a constraint satisfaction problem and solve it using Satisfiability Modulo Theories [1] (SMT) solvers.

Implementation of this technique requires two major components: (1) subnet gateways to perform rIP-vIP translation, and (2) a central management authority which coordinates mutation across network. In a traditional network these components must be incorporated in the network architecture. This incorporation could be disruptive and costly. Furthermore, it poses serious network management challenges such as real-time global reconfiguration, and synchronization of several network devices in a decentralized environment.

Software-defined networking (SDN) provides flexible infrastructure for developing and managing random host mutation efficiently and with minimal operational overhead. In SDN, the network controller (*e.g.*, NOX [2]) monitors and controls the entire network from a *central* vantage point via an interface, such as OpenFlow [3] and defines the forwarding and address translation behavior of switches distributed in the network accurately and synchronously.

In OF-RHM, the controller preforms the following tasks:

HotSDN 12, August 13, 2012, Helsinki, Finland.

OpenFlow Random Host Mutation: Transparent Moving Target Defense using Software Defined Networking

Jafar Haadi Jafarian, Ehab Al-Shaer, Qi Duan
Department of Software and Information Systems
University of North Carolina at Charlotte
Charlotte, NC, USA
{jjafaria, ealshaer, qduan}@uncc.edu

Static configurations serve great advantage for adversaries in discovering network targets and launching attacks. Identifying active IP addresses in a target domain is a precursory step for many attacks. Frequently changing hosts' IP addresses is a novel proactive moving target defense (MTD) that hides network assets from external/internal scanners. In this paper, we use OpenFlow to develop a MTD architecture that transparently mutates IP addresses with high unpredictability and rate, while maintaining configuration integrity and minimizing operation overhead. The presented technique is called OpenFlow Random Host Mutation (OF-RHM) in which the OpenFlow controller frequently assigns each host a random *virtual* IP that is translated to/from the *real* IP of the host. The real IP remains untouched, so IP mutation is completely transparent to end-hosts. Named hosts are reachable via the virtual IP addresses acquired via DNS, but real IP addresses can be only reached by authorized entities. Our implementation and evaluation show that OF-RHM can effectively defend against stealthy scanning, worm propagation, and other scanning-based attack.

Categories and Subject Descriptors

C.2.3 [**Computer-Communication Networks**]: Network Operations

IP mutation, software defined networking (SDN), Moving target defense (MTD), Security

1. INTRODUCTION

Static assignment of IP addresses gives adversaries significant advantage to remotely scan networks and identify their targets accurately and quickly. Scanning tools and worms usually send probes to random IP addresses in the network in order to discover their targets. When a target responds, it can then be identified and attacked. Otherwise the probed addresses will be considered unused. Despite firewall deployment, most enterprise networks have many public and private hosts accessible from outside. The IP address assignment scheme can become more dynamic by using approaches based on DHCP or NAT, but they are insufficient to provide proactive countermeasure because the IP mutation is infrequent and traceable.

In this paper, we introduce a moving target technique called *OpenFlow Random Host Mutation* (OF-RHM) which mutates IP addresses of end-hosts randomly and frequently so that the attackers' premises about the static IP assignment of network fails. OF-RHM has two main objectives. Firstly, the IP mutation must be transparent to the end-host. To provide transparency, OF-RHM keeps the actual or real IP addresses of hosts (called *rIP*) unchanged, but associates each host with random short-lived virtual IP addresses (called *vIP*) at regular intervals which are translated to rIPs right before the host. Secondly, the IP mutation must be performed with high unpredictability and speed to maximize the distortion of attackers' knowledge about the network and increase deterrence of attack planning. To optimize IP mutation with respect to unpredictability and speed, the mutant vIPs are selected randomly from the entire unused address space in the network. The unused address ranges must be assigned to hosts such that it satisfies several constraints including mutation unpredictability and minimum required mutation rate of all hosts. We formulate this problem as a constraint satisfaction problem and solve it using Satisfiability Modulo Theories [1] (SMT) solvers.

Implementation of this technique requires two major components: (1) subnet gateways to perform rIP-vIP translation, and (2) a central management authority which coordinates mutation across network. In a traditional network these components must be incorporated in the network architecture. This incorporation could be disruptive and costly. Furthermore, it poses serious network management challenges such as real-time global reconfiguration, and synchronization of several network devices in a decentralized environment.

Software-defined networking (SDN) provides flexible infrastructure for developing and managing random host mutation efficiently and with minimal operational overhead. In SDN, the network controller (*e.g.*, NOX [2]) monitors and controls the entire network from a *central* vantage point via an interface, such as OpenFlow [3] and defines the forwarding and address translation behavior of switches distributed in the network accurately and synchronously.

In OF-RHM, the controller preforms the following tasks:

HotSDN 12, August 13, 2012, Helsinki, Finland.

(1) coordinates mutation across OpenFlow switches based on host mutation requirements and available unused address space, (2) determines optimal set of new vIPs for hosts using SMT [1], (3) manages active connections by installing flows in OF-switches along with required address translations actions, and (4) handles DNS updates. Each OpenFlow switch (OF-switch) performs the vIP-rIP translations as specified by the controller.

We implemented the OF-RHM approach on OpenFlow managed by a NOX controller. To facilitate the development and analysis of our approach, we used Mininet [4] to generate fairly large networks of OpenFlow switches and hosts. Our theoretical analysis and implementation results show that OF-RHM can reduce the accuracy of information gathering via scanning up to 99%. Moreover, up to 90% of the network hosts are saved from vicious scanning worms. One limitation of OF-RHM is that a named host can still be reached via DNS. However, most existing scanners use IP address to collect information in order to avoid too many queries to DNS and thus detection.

The rest of the paper is organized as follows. Section 2 discusses related works. Section 3 defines and formulates the IP mutation problem. In Section 4 we describe the architecture and protocol details of OF-RHM. Section 5 describes the implementation of OF-RHM using OpenFlow and its evaluation against several attack models.

2. RELATED WORK

A few research proposals on dynamically changing IP addresses for proactive cyber defense have been presented in the literature. The APOD (Applications That Participate in Their Own Defense) scheme [5] uses *hopping tunnels* based on address and port randomization to disguise the identity of end parties from sniffers. However, this approach is not transparent as it requires cooperation of both client and server hosts during the IP mutation process. DyNAT [6] provides a transparent approach for IP hopping by translating the IP addresses before packets enter the core or public network in order to hide the IP address from man-in-the-middle sniffing attacks. Although this technique will make network discovery infeasible for sniffers, it does not work for scanners who rely on probe responses for discovering the end-hosts. A network address space randomization scheme called NASR [7] was proposed to offer an IP hopping approach that can defend against hitlist worms. NASR is a LAN-level network address randomization scheme based on DHCP update. NASR is not transparent to the end- hosts because DHCP changes are applied to the end-host itself which results in disruption of active connections dur- ing address transition. Moreover, it requires changes to the end-host operating system which makes its deployment very costly. Also, NASR provides very limited unpredictability and mutation speed because its IP mutation is limited on the LAN address space and will require DHCP and host to be reconfigured for this purpose (the maximum IP mutation speed is once every 15 minutes).

In summary, none of the previous techniques provide a deployable transparent mechanism for IP mutation that can defend against external and internal scanning attacks without changing the configuration of the end-hosts. OF-RHM exploits the power for software-defined networking to implement an efficient IP mutation in term of unpredictability, mutation speed and configuration management. Unlike the previous techniques, OF-RHM uses the entire address space to increase unpredictability and updates configurations at real-time while preserving network operation integrity.

3. PROBLEM DEFINITION AND FORMULATION

In OF-RHM, each host is associated with an unused address range of the network based on its specific requirement. At each mutation, a vIP is chosen from this range and associated with the host. The vIP of each host is mutated after each *mutation interval*.

The main objective of OF-RHM is to maximize both mutation unpredictability and mutation rate. The proposed technique must address both IPv4 and IPv6 address schemes. Scarcity of IP addresses in IPv4 networks makes the unused address space small and highly fragmented. Therefore, major challenge of OF-RHM is to guarantee that, even with a limited and fragmented unused address space, each host would mutate with its required rate such that no IP ad- dress is reused (assigned to any host more than once) for a reasonably long time.

These objectives can be achieved by choosing each vIP from largest possible unused address space such that the same vIP is not assigned more than once to any host in many consecutive vIP mutations. This problem can be di- vided into two sub-problems: (1) allocating unused ranges to hosts, and (2) mutation within allocated ranges.

3.1 Range Allocation

Suppose we have a set of n hosts $\{h_1, \ldots, h_n\}$. Minimum required mutation rate (R_i) for each host h_i is provided as input. In general, sensitive hosts are supposed to have higher mutation rates. Each host belongs to a subnet in the set $\{s_1, \ldots, s_z\}$, where subnet is a group of hosts that are physically connected through an OF-switch.

For mutation, we need to determine the unused address ranges in the network address space. Given used address ranges $A_1, \ldots, A_u$, we determine the contiguous blocks of unused address ranges of the network by simply masking the full network address space A as follows using Boolean operations:

$$\{r_1, r_2, .., r_m\} \leftarrow A \wedge \neg(A_1 \vee \ldots \vee A_u) \quad (1)$$

If a range is larger than a maximum size, it is divided into smaller ranges.

Sharing ranges among hosts allows us to increase mutation unpredictability and rate, because the host can mutate in a larger range. However, routing limitation does not allow us to share all unused ranges between all hosts, because each range can only be routed to one subnet. Based on these considerations, the OF-RHM problem is:

Given unused ranges $r_1, \ldots, r_m$ and subnets $s_1, \ldots, s_z$, what is the appropriate assignment scheme such that the following objectives are achieved:

- *Objective I*: the ranges assigned to the subnet must include enough IP addresses to satisfy the minimum required mutation rate of all hosts in that subnet during an interval, T, such that no IP addresses is assigned twice in one interval.
- *Objective II*: unpredictability and mutation rates must be maximized by firstly allocating all unused address

(1) coordinates mutation across OpenFlow switches based on host mutation requirements and available unused address space, (2) determines optimal set of new vIPs for hosts using SMT [1], (3) manages active connections by installing flows in OF-switches along with required address translations actions, and (4) handles DNS updates. Each OpenFlow switch (OF-switch) performs the vIP-rIP translations as specified by the controller.

We implemented the OF-RHM approach on OpenFlow managed by a NOX controller. To facilitate the development and analysis of our approach, we used Mininet [4] to generate fairly large networks of OpenFlow switches and hosts. Our theoretical analysis and implementation results show that OF-RHM can reduce the accuracy of information gathering via scanning up to 99%. Moreover, up to 90% of the network hosts are saved from vicious scanning worms. One limitation of OF-RHM is that a named host can still be reached via DNS. However, most existing scanners use IP address to collect information in order to avoid too many queries to DNS and thus detection.

The rest of the paper is organized as follows. Section 2 discusses related works. Section 3 defines and formulates the IP mutation problem. In Section 4 we describe the architecture and protocol details of OF-RHM. Section 5 describes the implementation of OF-RHM using OpenFlow and its evaluation against several attack models.

2. RELATED WORK

A few research proposals on dynamically changing IP addresses for proactive cyber defense have been presented in the literature. The APOD (Applications That Participate in Their Own Defense) scheme [5] uses *hopping tunnels* based on address and port randomization to disguise the identity of end parties from sniffers. However, this approach is not transparent as it requires cooperation of both client and server hosts during the IP mutation process. DyNAT [6] provides a transparent approach for IP hopping by translating the IP addresses before packets enter the core or public network in order to hide the IP address from man-in-the-middle sniffing attacks. Although this technique will make network discovery infeasible for sniffers, it does not work for scanners who rely on probe responses for discovering the end-hosts. A network address space randomization scheme called NASR [7] was proposed to offer an IP hopping approach that can defend against hitlist worms. NASR is a LAN-level network address randomization scheme based on DHCP update. NASR is not transparent to the end-hosts because DHCP changes are applied to the end-host itself which results in disruption of active connections during address transition. Moreover, it requires changes to the end-host operating system which makes its deployment very costly. Also, NASR provides very limited unpredictability and mutation speed because its IP mutation is limited on the LAN address space and will require DHCP and host to be reconfigured for this purpose (the maximum IP mutation speed is once every 15 minutes).

In summary, none of the previous techniques provide a deployable transparent mechanism for IP mutation that can defend against external and internal scanning attacks without changing the configuration of the end-hosts. OF-RHM exploits the power for software-defined networking to implement an efficient IP mutation in term of unpredictability, mutation speed and configuration management. Unlike the previous techniques, OF-RHM uses the entire address space to increase unpredictability and updates configurations at real-time while preserving network operation integrity.

3. PROBLEM DEFINITION AND FORMULATION

In OF-RHM, each host is associated with an unused address range of the network based on its specific requirement. At each mutation, a vIP is chosen from this range and associated with the host. The vIP of each host is mutated after each *mutation interval*.

The main objective of OF-RHM is to maximize both mutation unpredictability and mutation rate. The proposed technique must address both IPv4 and IPv6 address schemes. Scarcity of IP addresses in IPv4 networks makes the unused address space small and highly fragmented. Therefore, major challenge of OF-RHM is to guarantee that, even with a limited and fragmented unused address space, each host would mutate with its required rate such that no IP address is reused (assigned to any host more than once) for a reasonably long time.

These objectives can be achieved by choosing each vIP from largest possible unused address space such that the same vIP is not assigned more than once to any host in many consecutive vIP mutations. This problem can be divided into two sub-problems: (1) allocating unused ranges to hosts, and (2) mutation within allocated ranges.

3.1 Range Allocation

Suppose we have a set of n hosts $\{h_1, \ldots, h_n\}$. Minimum required mutation rate (R_i) for each host h_i is provided as input. In general, sensitive hosts are supposed to have higher mutation rates. Each host belongs to a subnet in the set $\{s_1, ..., s_z\}$, where subnet is a group of hosts that are physically connected through an OF-switch.

For mutation, we need to determine the unused address ranges in the network address space. Given used address ranges $A_1, \ldots, A_u$, we determine the contiguous blocks of unused address ranges of the network by simply masking the full network address space A as follows using Boolean operations:

$$\{r_1, r_2, .., r_m\} \leftarrow A \wedge \neg(A_1 \vee \ldots \vee A_u) \quad (1)$$

If a range is larger than a maximum size, it is divided into smaller ranges.

Sharing ranges among hosts allows us to increase mutation unpredictability and rate, because the host can mutate in a larger range. However, routing limitation does not allow us to share all unused ranges between all hosts, because each range can only be routed to one subnet. Based on these considerations, the OF-RHM problem is:
Given unused ranges $r_1, ..., r_m$ and subnets $s_1, ..., s_z$, what is the appropriate assignment scheme such that the following objectives are achieved:

- *Objective I*: the ranges assigned to the subnet must include enough IP addresses to satisfy the minimum required mutation rate of all hosts in that subnet during an interval, T, such that no IP addresses is assigned twice in one interval.

- *Objective II*: unpredictability and mutation rates must be maximized by firstly allocating all unused address