TOPHAT: Topology-Based Host-Level Attribution for Multi-stage Attacks in Enterprise Systems Using Software Defined Networks

Subramaniyam Kannan[1(✉)], Paul Wood[1], Larry Deatrick[2], Patricia Beane[2], Somali Chaterji[1], and Saurabh Bagchi[1]

[1] Purdue University, West Lafayette, USA
{kannan5,pwood,schaterji,sbagchi}@purdue.edu
[2] Northrop Grumman, Falls Church, USA
{larry.deatrick,Patricia.Beane}@ngc.com

Abstract. Multi-layer distributed systems, such as those found in corporate systems, are often the target of multi-stage attacks. Such attacks utilize multiple victim machines, in a series, to compromise a target asset deep inside the corporate network. Under such attacks, it is difficult to identify the upstream attacker's identity from a downstream victim machine because of the mixing of multiple network flows. This is known as the attribution problem in security domains. We present TopHat, a system that solves such attribution problems for multi-stage attacks. It does this by using moving target defense, *i.e.*, shuffling the assignment of clients to server replicas, which is achieved through software defined networking. As alerts are generated, TopHat maintains state about the level of risk for each network flow and progressively isolates the malicious flows. Using a simulation, we show that TopHat can identify single and multiple attackers in a variety of systems with different numbers of servers, layers, and clients.

Keywords: Multi-stage attacks · Attack attribution
Software defined network · Moving target defense

1 Introduction

Multi-stage attacks (MSA) have plagued distributed system administrators for decades. In these attacks, multiple computers are used simultaneously to breach a particular target, and attackers often rely on a series of privilege escalation attacks to circumvent access controls protecting assets. One of the most challenging aspects of MSA comes as an attribution, mixing, or traceability problem [4]. Defenders wish to know what particular network traffic resulted in a privilege escalation, to prevent it in the future, but from a network perspective, the traffic output at each stage is not associated with any particular input. Consequently, defenders cannot distinguish legitimate from malicious network traffic,

and identifying or disrupting vulnerabilities remains a daunting task. In this paper we present TOPHAT (TOPology-based Host-level ATtribution), a technique for identifying malicious users and their network traffic.

Multi-stage attacks operate on top of distributed systems where each distributed layer has different access privileges to sensitive business assets. An attacker must penetrate multiple layers to access some protected information, a *crown jewel*. As the attacker progresses, she generates some intrusion alerts due to some traffic with a malicious signature passing through intrusion detection systems (IDS). These alerts, while useful for finding single stage attacks, are less useful in the MSA because the {source, destination} pairs are both machines inside of the distributed system, instead of an external attributable source (as would be the case for an Internet-facing web server, for example). Consequently, there is no obvious relationship between alerts deep in the distributed system and the outsider, and this problem is referred to as the attribution, traceback or un-mixing problem [3,20]. In this context, an attributable alert is one which identifies an external source directly, and an unattributable alert is one which identifies no source or identifies an internal or intermediate source, which cannot actually be the attacker.

Existing solutions [1,2,5,14–16,19] to the attribution problem have a few common shortfalls that TOPHAT addresses. First, solutions such as [1,14,16] rely on attack graphs to perform alert inferencing, where existing relationships between alerts are known via expert system knowledge. For example, an expert would claim that a port scanning alert deep in the distributed system follows from a wrong password alert in the Internet-facing layers. In practice, such relationships are complex, numerous, and difficult to derive. Furthermore, it is challenging to keep such information updated because systems are dynamic with new vulnerabilities being discovered, new digital assets being brought online, and new users being added. TOPHAT solves this issue without relying on attack graphs, thus providing a more general, robust, and adaptive solution to solving the attribution problem. Second, solutions such as [2,5,15,19] rely on causal links between stages or layers of the MSA. For example, inside the system, it is known that input I_1 causes output O_1, and these relationships are logged and analyzed so that network traffic can be effectively tagged and tracked in the system. This approach relies on application support, however, to provide the causal links. TOPHAT does not rely on such information from the underlying application and can identify attackers without this causality link.

TOPHAT is a network-based solution to the attribution problem. We represent incident flows from external clients to alert sources in a directed acyclic graph, where each node in the graph models the mixing property of intermediate servers and softwares. Some of these flows are malicious, and they generate one or more alerts at various nodes, and at various depths, on its path. For each alert, we generate and track partial attribution for all clients that can reach the alerted node as a stateful metric called *risk factor*, or equivalently, *risk value*. TOPHAT, taking into consideration current risk for each network flow, adjusts the servers that the flow will pass through, using a process called *shuffling* [10]. Through

the shuffling process, TOPHAT isolates the suspect flows and keeps adjusting the risk factor. With a sufficient number of shuffles, the risk factor of the malicious flows exceeds a user-set threshold, *i.e.*, the cumulative partial attributions for an attacker reaches a level of complete attribution, and the attacker is identified.[1]

In TOPHAT, we utilize detection techniques that resemble moving target defenses (MTD) [8], through our shuffling algorithms. Using software defined networks (SDN) [12], TOPHAT is able to manipulate or re-route the network flows to desired nodes that in turn helps in identifying the attacker in the distributed system. Using SDN-based load balancers [18], entering flows from external clients are mapped to any replica of an entry-level server in the distributed system. Then, whenever an alert is generated, by an IDS placed at a replica of any server in the system, some risk is attributed to all flows that are passing through that server replica. Using two different approaches corresponding to two different variants of TOPHAT, it tracks this risk and assigns clients so that the malicious flows have progressively increasing risk factor. Finally, those with risk values above a user-settable threshold can be isolated, blocked, or studied in a honey-pot.

Using this approach, TOPHAT is able to identify a single attacker in a system of 1000 clients and 3 servers at the entry layer in 6 shuffles, requiring 1000 seconds whenever the attacker repeats the attack for approximately every 150 seconds. In the same system with 4 attackers, all of the attackers are identified in 27 shuffles. We also show that the same system with 10 attackers, the shuffling mechanism requires the attacker to repeat their exploits over 1000 times before gaining access to the crown jewel, thus significantly increasing the attacker's efforts under TOPHAT. Finally, we demonstrate how TOPHAT impacts the legitimate clients, showing that after 3–4 shuffles a majority of clients can retain continuous connectivity while the attacker is still identified.

The main contributions that we present in this paper are:

1. TOPHAT can attribute multi-stage attacks on a distributed system to a single external source, without relying on attack graphs or modifying the server softwares.
2. The MTD-style defense significantly increases attacker's effort, and can support identification of multiple simultaneous attackers.
3. TOPHAT can support high availability for legitimate clients while still identifying attackers in the system.

2   Background and Assumptions

2.1   System Model

TOPHAT is designed to protect a distributed system where servers exist at multiple layers in a distributed system as shown in Fig. 1. Each layer has multiple

---

[1] *Terminology clarification*: In this paper, we will use the term "attacker" synonymously with "attacking flow" or "malicious flow".

instances of one specific kind of server for load balancing and is connected to the next layer by an open flow switch and managed by a SDN Controller [6]. As a running example, we consider a web-based e-Commerce system operated by a publicly traded company. External clients access a web front end instance (layer 1) that connects to a database back end to store orders, interact with inventory, and otherwise manage transactions. In layer 3, a corporate reporting server analyzes the database to create sales reports, track hot products, and manage inventory at a macro level. It interfaces with the database layer and stores reports on layer 4, the corporate file servers. Inside of the corporate file server is an upcoming earnings statement for the next quarter (the *crown jewel*), and its early release would allow for insider trading since the company's performance, relative to projections, can have a significant impact on stock prices. The attacker(s) wish to ex-filtrate the earnings report.
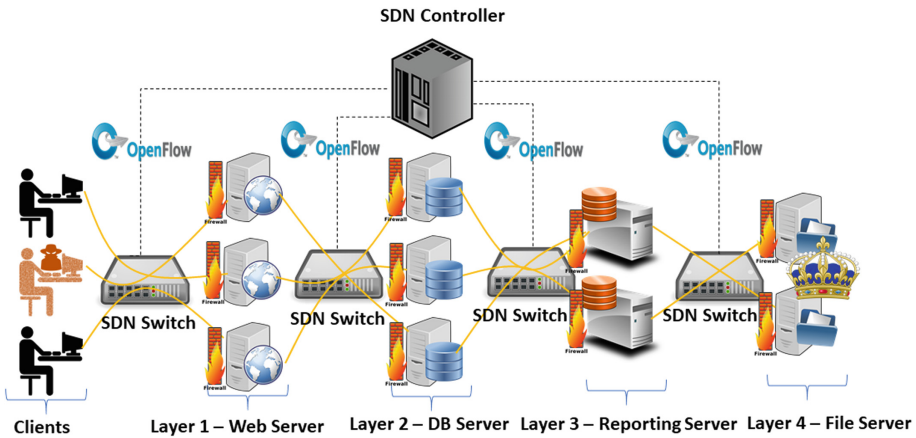

Fig. 1. A sample distributed system that can be protected by TopHat.

2.2 Network Structure

At its core, TopHat relies on intrusion detection systems to provide the alerts that drive its identification techniques. In this paper we make a simplifying assumption that each server instance has an IDS to remove resource management constraints.

Legitimate Client Model: We define a legitimate client as a system user that has no malicious intent and is using the target application for its designed purpose. The client connects to the application by sending a request to the outward facing service IP address. Further details on this process are described in Sect. 4.4.

Attacker Model: The attacker begins as a normal client and starts exploring for vulnerabilities in the outward facing layer 1. Once an exploit is found in layer 1, the attacker stages an attack on layer 2 by leveraging elevated access

privileges that she has gained at the outer layer 1. If an ongoing attack is flagged by an IDS at any layer then TOPHAT is activated. If the attack is undetected by any of the IDS, then it may proceed to the next layer, until reaching the *crown jewel*. TOPHAT works by making a few assumptions about the nature of the multi-stage attackers:

– **Persistent Attacks (PA)** if a server is reset, or the attacker connects to a new server, then the attack must be repeated.
– **Strong Alerts (SA)** the attacker will generate at least one strong alert during a MSA for which TOPHAT responds. The strong alert is known to be part of an attack with high certainty (*e.g.*, brute force attacks, known exploit signatures, or other high priority[2] alerts).

3   Solution Overview

TOPHAT utilizes software defined networks (SDN) and intrusion detection systems (IDS) to monitor and attribute alerts to specific attackers. At its core, TOPHAT sits along side a SDN controller such as an OpenDaylight [13] where it can observe the network flows and make decisions about changes to the network. The algorithm chooses which clients will be connected to which outward-facing servers, and which downstream servers are connected to which upstream servers in the distributed application. TOPHAT's algorithm operates by maintaining a risk factor for each connected client and then modifying that risk factor whenever alerts are generated. As more alerts are generated, the attacker's stateful *risk factor* is increased until she can be discriminated from the other connected clients. Whenever an alert is generated, the risk is increased for all the flows that are passing through the alerting service. The clients are then shuffled based on their risk so that over time, the attacker ends up with the maximum risk. The risk factor is initialized to zero for all clients and this monotonically increases with alerts in the system, till the attacker is identified and isolated. Then the risk factors of all the clients that are found to be legitimate in retrospect are reduced (Risk Rebalancing as explained in Sect. 4.3). We classify our protocol as an instantiation of Moving Target Defense (MTD), though it is somewhat different from the traditional notion of MTD. Here we are moving the clients and the assignment of flows to servers, while in traditional MTD, the protected system is "moved", *i.e.* reconfigured [7].

3.1   TopHat's Intuition

Several challenges exist in protecting a distributed system that has the structure shown in Fig. 1. First, alerts generated at any layer $(i + 1)$ look as if they are coming from layer i, not from an external attacker. This argues against the

---

[2] http://manual-snort-org.s3-website-us-east-1.amazonaws.com/node31.html#Snort_Default_Classifications in Snort, rules are tagged with priority where "high" priority correlates with strong in our solution.

simple solution of blocking flows from a particular source because that would create a service interruption—if a server in layer $(i+1)$ blocks a server in layer i, then the application stops working for all the clients connected to that particular server in layer i. TOPHAT overcomes this limitation by attributing an attack to all clients that are connected to the alerting server in layer i and then stopping the ongoing attack using the MTD approach. When an alert event happens, *all* of the clients are disconnected from the servers in layer 1 (for purposes of randomization), assigned to new servers, and the alerting server is refreshed to a clean state and restarted. TOPHAT then constantly tracks the attack history of each client with the help of the *risk factor* as we describe in detail in Sect. 4, so that the attacker is identified due to multiple alerts, which in turn is due to the persistence of the attack (as assumed in our attack model). The persistent attack property fundamentally allows TOPHAT to converge given a sufficient number of alerts.

3.2   Legitimate Client Impacts

The SDN-based shuffling in TOPHAT can have some negative impacts on legitimate client connections. First, whenever a shuffle involves a client, the client's connection is reset. This overhead cannot be avoided since the attackers and legitimate clients share the same network flow paths—a connection reset that disrupts an attacker's flow also disrupts the legitimate client's flow. Its impact can be mitigated, however, with state management approaches [17]. Second, when a server is being reset and restarted (to clear the infected status), the clients assigned to that server cannot function. This case can be minimized by using fast restart hardware or by keeping hot spares for the server instances.

4   Detailed Design of TopHat

TOPHAT probabilistically identifies attackers in the system by repeatedly attributing alerts to suspect sets until the likelihood of a client being the attacker is sufficiently large to certify identification. In the case of multiple attackers, this process is repeated so that multiple identification events occur until all of the attackers are exhausted.

Alert Group Attribution: In TOPHAT, there always exists a mapping between a server in any particular layer and the clients that, through any possible path, have access to that server. For example, if clients 1–5 are assigned to server S1 in layer L1, and S1/L1 is connected to S2/L2, then an alert sourced from S2/L2 will be attributed to *all* the clients 1–5. The relationship of how any given flow passes through the servers at the different layers is itself controlled by the SDN controller and thus this relationship is always known to our algorithm. Now we define a term *client group*. Consider that an alerting server has flows $F_1, F_2, ..., F_{N_G}$ going through it. By tracing each flow back to layer 1 servers, we can map each flow F_i to the client generating that flow C_i. The clients $C_1, C_2, ..., C_{N_G}$ form the client group here. Each such client has its stateful parameter, *risk factor*,

increased by $\frac{1}{N_G}$, where N_G is the number of clients in that particular group. In the earlier example, each client would have its risk increased by 1/5.

Likelihood of a Client Being the Attacker: We define the likelihood as follows:

$$P(C_i = A) = \frac{R(C_i)}{\sum R(C_j) \,\forall\, R(C_j) \geq R(C_i)} \tag{1}$$

where C_i is client i, $C_i = A$ is the indicator that C_i is an attacker, $R(C_i)$ is the risk factor of client i, and $\forall\, R(C_j) \geq R(C_i)$ implies that client j has a risk factor at least as large as client i, and $i = j$ is allowed. When a client has the highest risk factor of any client, then this probability value becomes 1. The control and convergence of TOPHAT is discussed in detail in Sect. 4.1 in [9].

4.1 Uniform Assignment Algorithm ("Uniform")

The uniform assignment algorithm is responsible for assigning arriving client flows to different servers at layer 1. There are N_S assignment pools available, where N_S is the number of servers in layer 1. For each client i, an assignment is made: $A : C_i \to [1, N_S]$ such that the imbalance in risk between any two servers is minimized. At the beginning of the operation of the system, each client will have the same risk factor and so this will be a uniform random assignment. However, in subsequent mappings (which happen after an alert arrives at TOPHAT) the risk factors will be different and the mapping A will be a weighted random assignment, using the risk factors as the weights. The goal is to balance the aggregate risk at any of the servers in level 1. The assignment process proceeds as follows:

1. A client seeks an assignment, either when it is connecting to the protected system for the first time, or in response to a disconnection forced by TOPHAT.
2. The client is given the assignment to $[1, N_S]$ according to the assignment function A.
3. When a new alert is received, the assignments between clients and servers are reset, and all clients return to step 1 and re-assigned to new servers.

This algorithm effectively assigns clients such that there is a uniform aggregate risk assigned to any particular server. Each attribution event reduces the set of ties ($N_T = |\forall\, R(C_j) \geq R(C_i)|$) to $\frac{N_T}{N_{AS}}$, where N_{AS} is the number of servers in the alert layer. For example, if there are 100 clients and 4 alert groups and every client has a risk of 1, then by Eq. 1, $P(C_i = A) = \frac{1}{100} \,\forall\, i$. After an attribution event, given uniform assignment (each server having balanced risk of 25, thus 25 clients per server), then the likelihood for those 25 becomes $\frac{1.04}{1.04 \times 25}$ because the 25 clients that were attributed with risk have an additional 0.04 added. The size of the set $|R(C_j) \geq R(C_i)|$ is now $\frac{100}{4} = 25$, following the reduction.

4.2 Low-Risk Isolation Algorithm ("LRA")

This variant of the algorithm shelters low risk clients into a *safe zone*, defined as a set of servers in layer 1 such that clients which are assigned to this set are not

shuffled around by TOPHAT. These clients do not suffer from any disconnections, and their risk factors do not change. Each alert/attribution event tells TOPHAT something about who may be the attacker, but it can also indicate who is *not* an attacker. In the uniform case, the legitimate clients are mixed in with the attackers, and this causes them to rise in risk, whenever they share a server with the malicious clients. It also dilutes the attribution power of a single attack since N_G remains near-constant. The LRA variant avoids this issue by placing some portion of the clients with the lowest risk into a safe zone:

1. Clients are assigned as in the uniform risk case, except for clients that exist in a safe set S_S, initialized as empty.
2. After an attribution event a portion of the clients, $I_R, I_R \in (0, 1)$, is moved from the active set S_A to the safe set S_S. The $|S_A| \cdot I_R$ clients with the lowest risk are moved to S_S.
3. The assignment of the safe clients S_S is fixed to a particular server, and then the clients in the active set S_A are redistributed among the remaining $N_S - N_{safe}$ servers using the uniform risk approach. N_{safe} is the number of servers used for the safe zone.
4. In the event an alert is generated from any of the S_S clients, then the entire set of clients is moved back to the S_A set.

Much less risk is assigned to the legitimate clients in the system using LRA with a single attacker. Additionally, uninterrupted connection paths to the protected application are provided, decreasing the probability of failed transactions. The convergence condition for both of the algorithms are explained in detail in Sects. 4.2 and 4.3 in [9].

4.3 Risk Rebalancing Approach ("RRB")

Once one of the attackers is identified by using any one of the above described algorithms, the risk factor of the remaining clients are updated using *Risk Rebalancing* (RRB) technique in order to speed up identification of the remaining attackers. Each alert attribution is stored in the SDN controller that contains the list of clients and the amount of risk factor attributed to each client due to that particular alert. Whenever an attacker is identified, the list of alerts is searched, and the set of alerts that involved the attacker are collected. The accumulated risk for each client due to each alert in that list is removed because of the insight that the alert is attributable to the now discovered attacker and not the other clients. Thus, legitimate clients have their risk lowered leading to faster identification of the other attackers.

4.4 End-to-End Workflow

We detail the end-to-end workflow of TOPHAT in the context of the SDN-based system:

1. **Initial:** The SDN switch at ingress node forwards each new client's request to the SDN controller as the flow table will be initially empty. TOPHAT, which is installed as an application over the SDN controller, stores the associated risk and the server allocated at each layer for all the clients.
2. **Server Assignment:** TOPHAT assigns each client to a particular server at layer 1 as described in Sects. 4.1 and 4.2. Then the corresponding flow rules are installed at the SDN switch in layer 1 and subsequent layers. The initial risk factor of all the clients are set to 0. (T_S: Time for server assignment).
3. **Connection Establishment:** Each client establishes a connection with the servers at layer 1 using TCP 3-way handshake. At this point, all the clients except the attackers can access the servers in subsequent layers using their respective access privilege. (T_C: Time to establish connection).
4. **Attacker Exploration:** In order to get access to the subsequent layers, the attackers have to explore the layer 1 server for vulnerabilities and then exploit a vulnerability. Let T_x denote the time to exploit a server at a particular layer. T_x varies across different layers and across different attackers.
5. **Alert Generation:** The attacker continues to compromise the servers at subsequent layers until an IDS detects a malicious action (e.g., port scan, known CVE, etc.) or alert correlation from multiple IDS alerts generates a strong alert. Let T_A be the time to generate a strong alert.
6. **Connection Termination:** The strong alert is sent to the SDN controller, which initiates the shuffling by disconnecting the clients from the servers in layer 1 (except those in the safe set for LRA) and reassigning them.
7. **Risk Updation:** The risk factor of the clients are updated according to either the Uniform or the LRA scheme. Let T_{RA} be the time to update risk values.
8. **Attacker Identification:** After the risk updation, the probability of each client is calculated using Eq. (1). The clients with a probability $P(C_i = A) \geq \tau$ (τ: user settable threshold) are identified as attackers and isolated.
9. **Risk Rebalance:** After the attacker is identified, TOPHAT rebalances the risk factor of all the remaining clients (Sect. 4.3).
10. **Server Reset:** TOPHAT instructs the SDN controller to reset all the active servers in the network by broadcasting a control message which ensures that the attackers need to exploit it again, in order to re-initiate the MSA. (T_R: Time to reset a server).
11. **Connection Re-establishment:** All the clients including the attackers will re-initiate connections to the servers in layer 1 and the steps repeat.

4.5 Multiple Attackers

Multiple simultaneous attackers can be handled by TOPHAT, without any modification. We model multiple attackers as each having independent, random times to exploit (T_X), where a successful exploit results in an alert being generated. If one attacker is more aggressive (smaller T_X), then alerts will be generated due to this attacker and this attacker will be identified by TOPHAT before moving

on to the next attacker. This essentially makes the process of identifying multiple attackers sequential. If on the other hand, there are multiple attackers with similar T_X values, then it will be a matter of chance which attacker gets identified first. But the risk factor of the other attackers will be retained in TopHat, thereby helping in the convergence time for the subsequent attackers.

False Positives and Mitigation: It is possible for TopHat to generate false positives with multiple attackers present that have similar T_X. For example, if there are four clients C1-C4, of which C2 and C4 are malicious and two servers S1 and S2. In the first round, C1 and C2 are assigned to S1 and C3 and C4 to S2. C2 alerts resulting in reshuffling. In the next round, C1 and C4 happen to be assigned to S1 and C2 and C3 to S2. Now C4 alerts and as a result, the legitimate client C1 is falsely flagged. This is a relatively rare occurrence and we show the false positive rate in Experiment 3 (it is below 5% even in the most pathological case).

5   Experimentation

5.1   Model System

The SDN Network used for evaluation is described in the Fig. 2. All the experiments are evaluated using the default values given in Tables 1 and 2 unless otherwise specified. As shown in Fig. 2, for the sake of simplicity we consider that each server in layer i has a stove piped connection or one-one connection (represented by different colors) to any server in layer i + 1 in order to avoid mixing of network flows at later stages. The experiment 4 shows the convergence for non-stove piped case. In LRA approach, the server 3 is considered to be safe server and the clients in active set S_A are shuffled between the server 1 and server 2.


Fig. 2. SDN model system considered for the evaluation.

Table 1. Default network and time parameter values.

Notation	Meaning	Default value
N_C	No. of clients	1000
N_A	No. of attackers	4
N_S	No. of servers at layer 1	3
N_L	No. of layers	4
L_{alert}	Strong alert layer number	2
I_R	Ratio of clients moved from active set to safe set	0.25
T_S	Server allocation time	1 ms
T_C	Connection establishment time	30 ms
T_A	Alert generation time	1 ms
T_{RA}	Risk assignment + attacker identification time	1 ms
T_R	Server reset time	45 s [11]
T_X	Attacker exploit time	Normal distribution

Table 2. Default attacker exploit time T_X for 4 attackers.

Attacker no.	T_X at layer 1		T_X at layer 2	
	Mean (s)	Variance (s)	Mean (s)	Variance (s)
1	20	5	30	5
2	40	5	60	5
3	10	2	15	2
4	80	5	80	5

Simulation Environment: Along with TOPHAT, the SDN environment is modeled using the network and time parameters in C++[3]. Each event in the SDN environment is represented by a corresponding time component as described in Sect. 4.4. The clients are assigned to the available servers using uniform random distribution and the attacker's exploit time is modeled based on normal distribution as in Table 2. For each attacker, the exploit time varies by mean across each layer and varies by variance across different iterations or shuffles. For some experiments, where multiple simulations can be aggregated, we take the median of 20 runs to provide data smoothness with respect to the random attack times.

Evaluation Parameters: For simplification, we assume all the clients send requests to the servers at layer 1 at the same time. All the experiments described below are evaluated using the following parameters:

1. **Experiment Time:** The time at which particular event like server assignment or alert generation happens.

---

[3] https://github.rcac.purdue.edu/DependableComputingSystemsLab/TopHat.

2. **Convergence Time:** The time at which single attacker or all the attackers are found.
3. **Probability of Attacker found (P_A):** The probability of client being identified correctly as an attacker given by Eq. 1.
4. **Percentage of Failed Transactions (PFT):** The number of client disruptions during the time of attacker identification. It is a function of time given by

$$PFT(t) = \frac{No.\ of\ Failed\ Transactions}{Total\ No.\ of\ Transactions} \qquad (2)$$

where we model client transactions as continuous time event for simplicity. We aggregate PFT across clients and all time to compute a **cumulative PFT** for the purpose of comparing per-simulation metrics. Note that the PFT is per-client, and not all clients are disrupted simultaneously during a shuffle event in TopHat.

5.2 Experiment 1: Convergence over Time

The experiment 1 demonstrate TopHat's operation in the time domain for both single and multiple attackers. During each attack, the two primary metrics (PFT and P_A) are collected based on the experiment time at which an alert is generated. Default values are used for all parameters except N_A. Figure 3 shows the results from our simulation and the results are explained in the next sections.

Convergence: In case of single attacker, the convergence is given directly by P_A and for two attacker's case, it is given by average probability. At each alert generation, the probability is updated, and the value for the attacker increases as shown in the figures. The uniform algorithm converges more quickly in both cases primarily because it has 3 servers to use for risk attribution while LRA reserves a server for the safe pool and uses only 2 servers for risk attribution. The step function increases as the number of ties are broken, and the attacker is repeatedly involved in high-risk attribution events.

In the case of multiple attackers, one attacker has a faster exploit time than the other. Since shuffles occur on the fast attacker's alert, the slow attacker is statistically unlikely to ever generate an alert until the fast attacker has been disrupted. This causes a time-domain crowding of alerts early in the simulation until the first attacker is identified, and then the alerts become more spaced out as opposed to an independent case where the alerts would be interleaved. Upon close inspection, one of the LRA's potential weaknesses can be seen in that it is using more shuffles to identify the attacker in the two attacker case. Furthermore, because the slow attacker has low risk, she can be placed in the safe zone, and it is more likely that a slow attacker can generate an alert inside of the safe area–something that does not happen in this experiment, but will in a later experiment.

PFT: The PFT shows how clients are impacted through time. In all cases, the width of the PFT bar represents the reset time for cleaning impacted servers

(a) Single Attacker   (b) Two Attackers

Fig. 3. The convergence of TOPHAT is shown for the single attacker and two attacker cases with both the uniform and low-risk assignment (LRA) algorithms.

in the system T_R. For the uniform algorithm, all clients are re-assigned and all servers on the attack path are cleaned, resulting in outages for all of the clients, hence the peak is always at 100. In the LRA case, only those clients remaining in the active set are impacted for each attack. This results in a decaying PFT over time as the low risk clients are assigned to the safe server at the rate I_R. Consequently, system operators have a choice between faster convergence and attacker identification (the Uniform variant) or slower convergence with better client access (the LRA variant).

For multiple attackers, in the Uniform case the PFT follows the single attacker profile, but it is repeated for the second attacker with a higher width due to T_X. For the LRA case, since less shuffling servers are available, it takes more alerts to converge and thus more shuffles, and more period of high PFT. The impact of all the parameters on the convergence time and PFT is explained in detail in Sect. 5.3 in [9].

5.3 Experiment 2: Attacker Effort

In this experiment, we demonstrate how TOPHAT, by utilizing MTD, is able to increase the total attack effort that must be expended to compromise the protected system. We measure attacker effort as the number of times a server must be compromised, at any layer, by any attacker. This includes the effort spent exploiting servers that have been reset. We also measure the number of shuffles or alerts generated in the system, and this metric covers the number of trials an attacker has at penetrating a system for which the exploit is not known.

Figure 4 shows the effort in these two metrics. In Fig. 4a, the total exploits goes up with the number of attackers. This process is not linear, however, because many attackers will be reset even when they do not generate an alert themselves due to the moving target nature of TOPHAT. Each attacker may penetrate layer 1 and be shuffled before making an attempt on layer 2, for example. Consequently,

(a) Total effort for all attackers   (b) Number of tries at alert stage

Fig. 4. Attacker efforts until identified

TOPHAT is able to make it much more difficult to attack the system when multiple attackers are present, even if the attacker identification takes some time. In Fig. 4b, the total number of resets are shown. This scales roughly linearly with the number of attackers because there is a lower limit to this number until the attackers can be found, as described in Sect. 4.1. of note here, however, is that there are a limited number of exploit attempts allowed at layer 2 before the attackers are identified and a layer 1 patch can be created. These 5–100 alerts will attribute the attacker, and upstream compromises (at layer 1) can be patched as a result, a key benefit of TOPHAT.

5.4   Experiment 3: Effect of Risk Re-balancing

For this experiment, we evaluate the impact of the risk re-balancing (RRB) technique (Sect. 4.3) on the convergence time and the false positives. We stress the system by having multiple attackers with the same distribution for T_X. Without RRB, when an attacker is identified, the risk for *all* other clients is reset to zero. With RRB, when an attacker is identified, only the legitimate clients that had been mixed in with the identified attacker have their risk reduced, not reset to zero.

Figure 5a shows the impact of RRB on both the Uniform and the LRA algorithms. In both cases, the use of the RRB speeds up convergence as expected. The number of false positives is higher for LRA. This is because the placement of many clients in the safe zone and subsequent alerts from that zone can degrade the process of identification of the attackers. TOPHAT is still able to provide low false positive rates (less than 0.5%) for small numbers of attackers relative to the total number of clients (10), even in this challenging scenario of similarly aggressive attackers.

5.5   Experiment 4: Effect of Number of Server Replicas

For this experiment, we use a system with 5 layers having [5, 4, 3, 2, 2] replicas in the layers, starting from layer 1. The inter-layer connections are uniformly

(a) Impact of maintaining state (w/ RRB) on convergence and false positive rates.

(b) As the alert becomes deeper in the system, it requires more shuffles and thus increases the time to converge.

Fig. 5. Experiments 3 (left) and 4 (right)

balanced as much as possible. We evaluate the impact of alert depth on the risk attribution. Figure 5b shows the impact of the alert layer on the convergence speed of both algorithms. The number of replicas decreases as one goes further inside the system. This is not uncommon because the number of requests that touch servers deep inside the enterprise typically decrease. We expect that alerts deep in the system will provide less discriminating information about the attackers because the shuffling can occur with coarser granularity, thus lumping more number of clients (legitimate with a few attacking) together on the same server. In the case of LRA, the safe zone is on a single stove-piped layer while the other shuffling servers are all connected into the multi-layer system. As the layer deepens, it is similar to reducing N_S because the size of the alert group increases and the number of groups N_G decreases. Therefore there is a logarithmic increase in the convergence time as the depth of the alert layer increases.

6   Conclusion

In this paper, we presented TOPHAT, a solution to the problem of attributing an alert to an attacker in a multi-layered system. The problem is challenging due to the mixing of multiple flows at servers inside the periphery of the system. TOPHAT utilizes moving target defense techniques, namely shuffling, implemented on top of a software defined network infrastructure. We provided two algorithms for shuffling, one that focuses on convergence speed and another that focuses on improving client connectivity during attacks. Further, we show that TOPHAT increases the attackers effort by requiring multiple re-exploiting of the target systems. We evaluate TOPHAT using the metrics of time to detect and isolate the attackers and the impact on the legitimate clients in the system. Using our system, network administrators can begin to attribute alerts and attacks, to external flows so that they may be blocked or studied for further defense improvement.

References

1. Alserhani, F., Akhlaq, M., Awan, I.U., Cullen, A.J., Mirchandani, P.: MARS: multi-stage attack recognition system. In: 2010 24th IEEE International Conference on Advanced Information Networking and Applications (AINA), pp. 753–759. IEEE (2010)
2. Baba, T., Matsuda, S.: Tracing network attacks to their sources. IEEE Internet Comput. **6**(2), 20–26 (2002)
3. Clark, D.D., Landau, S.: The problem isn't attribution: it's multi-stage attacks. In: Proceedings of the Re-architecting the Internet Workshop, p. 11. ACM (2010)
4. Clark, D.D., Landau, S.: Untangling attribution. Harv. Nat. Secur. J. **2**, 323 (2011)
5. Dawkins, J., Hale, J.: A systematic approach to multi-stage network attack analysis. In: Proceedings of Second IEEE International Information Assurance Workshop, pp. 48–56. IEEE (2004)
6. Feamster, N., Rexford, J., Zegura, E.: The road to SDN: an intellectual history of programmable networks. ACM SIGCOMM Comput. Commun. Rev. **44**(2), 87–98 (2014)
7. Jafarian, J.H., Al-Shaer, E., Duan, Q.: Openflow random host mutation: transparent moving target defense using software defined networking. In: Proceedings of the First Workshop on Hot Topics in Software Defined Networks, pp. 127–132. ACM (2012)
8. Kampanakis, P., Perros, H., Beyene, T.: SDN-based solutions for moving target defense network protection. In: 2014 IEEE 15th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM), pp. 1–6. IEEE (2014)
9. Kannan, S., Wood, P., Deatrick, L., Beane, P., Chaterji, S., Bagchi, S.: TopHat: topology-based host-level attribution for multi-stage attacks in enterprise systems using software defined networks. Technical report, CERIAS Tech Report TR 2017-4 (2017). https://www.cerias.purdue.edu/apps/reports_and_papers/
10. MacFarland, D.C., Shue, C.A.: The SDN shuffle: creating a moving-target defense using host-based software-defined networking. In: Proceedings of the Second ACM Workshop on Moving Target Defense, pp. 37–41. ACM (2015)
11. Mao, M., Humphrey, M.: A performance study on the VM startup time in the cloud. In: 2012 IEEE 5th International Conference on Cloud Computing (CLOUD), pp. 423–430. IEEE (2012)
12. McKeown, N.: Software-defined networking. INFOCOM Keynote Talk **17**(2), 30–32 (2009)
13. Medved, J., Varga, J., Tkacik, A., Gray, K.: OpenDaylight: towards a model-driven SDN controller architecture. In: 2014 IEEE 15th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM), pp. 1–6. IEEE (2014)
14. Modelo-Howard, G., Sweval, J., Bagchi, S.: Secure configuration of intrusion detection sensors for changing enterprise systems. In: Rajarajan, M., Piper, F., Wang, H., Kesidis, G. (eds.) SecureComm 2011. LNICST, vol. 96, pp. 39–58. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31909-9_3
15. Savage, S., Wetherall, D., Karlin, A., Anderson, T.: Practical network support for IP traceback. In: ACM SIGCOMM Computer Communication Review, vol. 30, no. 4, pp. 295–306. ACM (2000)
16. Strayer, W.T., Jones, C.E., Schwartz, B.I., Mikkelson, J., Livadas, C.: Architecture for multi-stage network attack traceback. In: The IEEE Conference on Local Computer Networks 30th Anniversary, pp. 8–pp. IEEE (2005)

17. Sultan, F., Srinivasan, K., Iyer, D., Iftode, L.: Migratory TCP: connection migration for service continuity in the internet. In: Proceedings of 22nd International Conference on Distributed Computing Systems, pp. 469–470. IEEE (2002)
18. Wang, R., Butnariu, D., Rexford, J., et al.: Openflow-based server load balancing gone wild. Hot-ICE **11**, 12 (2011)
19. Xu, Z., Wu, Z., Li, Z., Jee, K., Rhee, J., Xiao, X., Xu, F., Wang, H., Jiang, G.: High fidelity data reduction for big data security dependency analyses. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 504–516. ACM (2016)
20. Zhu, Y., Bettati, R.: Unmixing mix traffic. In: Danezis, G., Martin, D. (eds.) PET 2005. LNCS, vol. 3856, pp. 110–127. Springer, Heidelberg (2006). https://doi.org/10.1007/11767831_8