October 2019

BEHAVIORAL ANALYTICS FOR FIREWALL POLICY RECOMMENDATION AND TUNING

Raghunath Kulkarni

Prapanch Ramamoorthy

BEHAVIORAL ANALYTICS FOR FIREWALL POLICY RECOMMENDATION AND TUNING

AUTHORS:
Raghunath Kulkarni
Prapanch Ramamoorthy

ABSTRACT

Techniques are described herein for optimizing rules created on a firewall / intrusion prevention system to ensure that they stay relevant and updated. This may be achieved by leveraging network observables generated at a management center. These techniques are useful because behavior analytics associated with the user access of network is crucial in creating these rules.

DETAILED DESCRIPTION

Creating rules on a firewall is a complex task. The number of network and threat parameters vary for each traffic pattern in the network. This makes network administration very challenging. Currently, the network administrator creates a rule based on known Layer 2 (L2) – Layer 7 (L7) parameters without any insights into user traffic profile. This may often lead to failed audits because of generic rules present in the configuration. This can also lead to security holes in the configuration that can be leveraged by adversaries to gain access and compromise the network.

Management centers may provide insights into data pertaining to applications (client and server), intrusion prevent system events, the operating system, Uniform Resource Locators (URLs), file downloads by specific users in the network, etc. However, there is no viable mechanism available to leverage this information to generate a recommendation for access rule creation.

At the data collection stage, the management center may build the profile for the user. The profile may include the network observables that are reported by the management center. The observables may include networks accessed, URLs accessed, applications used, files downloaded, and intrusion and malware events triggered. Data may be collected when a threat defense module is running in learning or inline mode.

A profile-to-rule recommendation may be provided. The existing rule base may investigate the existing rule base using the available data. For the user present in the rule, the observables may be examined to determine whether the existing rule is permitting/denying the traffic. If the permit rule is a generic rule, the generated recommendation may be specific to the networks/URLs/applications that have been historically accessed. Based on the insights generated for the user pertaining to intrusion prevention signatures and observed malware threats, the security association of intrusion prevention policies and malware policies may be recommended. It will be appreciated that the intrusion prevention policy may not be a generic intrusion prevention policy, and may instead be specific to the user traffic profile.

The user profile may be built based on network access. The user profile may be used to recommend rules for the network administrator. A user profile may be built based on source Internet Protocol (IP) address or username as the key. A profile associated with the user may include networks accessed, URLs accessed, applications used, files downloaded, and/or intrusion and malware events. Once the profile has been built for the user based on these parameters, the next step is to generate recommendations for rule updates/modification.

A rule recommendation may be generated by reading the existing rule base and mapping the user profile created with the rule base to recommend updates. The "scan rule base" module may read through all the existing rules and categorize them into different types based on the parameters used in their configuration. A rule may be created at the management center using one or more of the following parameters: Layer 3 (L3) parameters (e.g., source/destination network), Layer 4 (L4) parameters (e.g., source/destination ports), L2 parameters (e.g., source/destination Virtual Local Area Networks (VLANs)), zones (e.g., source/destination), URL (e.g., custom/categories), application detectors, and user information (e.g., based on user, group, or Security Group Tag (SGT)).

Based on the parameters that are used to configure the rules, they may be categorized so that it becomes easier or more efficient to map them against built-in user profiles. The user profile may be scanned to match the rules that are present in the configuration. Once a rule is found that matches the user profile, the parameters within the

user profile may be examined to identify the behavior pattern. Based on the behavior pattern for the user profile, a recommendation may be made for a change in a rule by restricting the rule to specific network access. Additionally, if the user profile has triggers for intrusion and malware, a recommendation may be made for association of the intrusion prevention / malware analysis profile. The recommendations may not be deployed/applied to the device, but a network administrator may use/leverage the recommendation.

Figure 1 below illustrates an example of a pre-existing ruleset on a device. Here, the Demilitarized Zone (DMZ) interface is connected to the firewall.

Rule number	Src network	Dst network	Dst port	URL	Application	Src zone	Dst zone	IPS	Action
1	Any	Any	Any	Any	email service	Inside	outside	No	Deny
2	10.1.1.0/24	Any	Any	www.example4.com	Any	Inside	Outside	Yes	Allow
3	Any	172.16.1.100	80	Any	Any	Any	Any	Yes	Allow
4	Any	Any	Any	Any	Any	Any	Any	No	Allow

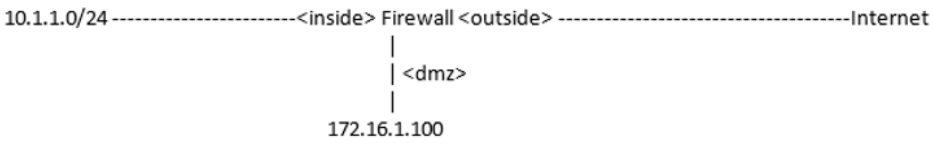

Figure 1

In one example scenario, a server hosts a webpage at http://example1.com. Rule 3 exists in order to allow external users on the Internet to connect to this web server on Transmission Control Protocol (TCP) port 80. However, on most occasions, customers tend to deploy rules which only specify layer 3 and layer 4 parameters, as is the case with rule 3. The solution described herein examines all traffic profiles destined to 172.16.1.100 and identifies all traffic limited to only http://example1.com. With this analysis, a recommendation may be made to specify a URL of example1.com to rule 3.

For the same rule, an intrusion prevention policy may be detected and applied with a default set of signatures enabled. With the aforementioned analysis concluding that this is primarily Hypertext Transfer Protocol (HTTP) 1.1 traffic destined to the server at 172.16.1.100 running a specific version of software, a recommendation of a customized intrusion prevention policy may be made that only has rules enabled for HTTP and the specific version of software, and has all the other rules disabled. Existing approaches use

intrusion prevention rule recommendations but those are system-wide rule recommendations affecting all traffic. The techniques described herein provide a per-rule and per-endpoint traffic profile based rule recommendations including custom intrusion prevention rules for each host.

In another example scenario, the subnetwork under consideration may be 10.1.1.0/24. One specific host in this subnetwork, 10.1.1.1, may generate network traffic to obtain access to URLs "example2.com" and "example3.com," to use a given application to place a video call, and to generate an intrusion prevention signature alert while accessing a server 172.16.1.100. If the host 10.1.1.1 attempts to access the URL "example3.com" this is allowed by rule number 4, which does not affect the intrusion prevention policy.

Based on the generated data, a profile may be built for host 10.1.1.1 that takes into consideration the network activity for the host 10.1.1.1. In this example, the host only accessed URLs "example2.com" and "example3.com," the given application was used by the host, and an intrusion prevention event was generated by the host when access occurred. An example profile may appear as follows:

10.1.1.1
{
URL: example2.com, example3.com
Application: [given application]
Intrusion prevention event generated: Yes, (signature id: 12345)
}

In the existing rule lookup phase, it may be determined that there are two rules which are relevant to the host 10.1.1.1 from the present rule set. Rule 1 applies if the host attempts to access an email service, rule 2 applies if the accessed URL is www.example4.com, rule 3 applies if the host is trying to access 172.16.1.100, and rule 4 applies for all the network traffic access. Rule 4 is generic and does not have an associated intrusion prevention policy.

In the generate recommendation phase, website access for the host should be limited to "example2.com" and "example3.com" rather than generic rule 4 which provides open access. The given application is being used by the host, and hence a specific rule is

to be created which allows only this application. Since the intrusion prevention event is generated by the host, the intrusion prevention system may be associated with all rules where host 10.1.1.1 is relevant (in this case rules 1 and 4).

Figure 2 below illustrates an example of a ruleset on a device in accordance with the techniques described herein.

Rule number	Src network	Dst network	Dst port	URL	Application	Src zone	Dst zone	IPS	Action
1	10.1.1.1	Any	Any	example1.com, example2.com, example3.com	Any	Inside	Outside, DMZ	Yes	Allow
2	10.1.1.1	Any	Any	Any	given application	Inside	Outside	Yes	Allow
3	Any	172.16.1.100	80	example1.com	Any	Any	Any	Yes	Allow
4	10.1.1.1	Any	Any	Any	Any	Any	Any	No	Deny
5	Any	Any	Any	Any	email service	Inside	Outside	No	Deny
6	10.1.1.0/24	Any	Any	example4.com	Any	Inside	Outside	Yes	Allow
7	Any	Any	Any	Any	Any	Any	Any	No	Allow

Figure 2

Comparing Figures 1 and 2, it is clear that the rule set has become more specific based on traffic profiles for the actual endpoints, be they clients or servers. Only recommendations are shown which are based on Layer 3 and Layer 4 parameters, URLs accessed, applications used, and intrusion prevention. Other parameters may also be used, such as user identities and group, file-based policies, VLANs, etc.

Techniques described herein may involve building a traffic profile for every endpoint, using the profile to compare against the existing rule set, and generating the recommended rule set based on the aforementioned analysis. Individually, there are systems that build traffic profiles and there are separate systems for rule optimization. However, traffic profiles are not associated with rule recommendations beyond Layer 3 and Layer 4 parameters on a system. In addition, the solution may be extended to improve the security posture by recommending more specific policies based on the results of the behavioral analysis of users. For example, the system may recommend intrusion rules that cover only the nature of applications accessed by users, characteristics of the endpoint(s) used by the user, etc.

Rule explosion (e.g., increasing the number of rules from four to seven, as shown above) for increasing numbers of endpoints is not a concern. Multiple endpoints may have

similar traffic profiles and may fall under the same group of users. Such endpoints may be categorized as being associated with the same application/operating system. Thus, the generated recommendation may be optimized.

In summary, techniques are described herein for optimizing rules created on a firewall / intrusion prevention system to ensure that they stay relevant and updated. This may be achieved by leveraging network observables generated at a management center. These techniques are useful because behavior analytics associated with the user access of network is crucial in creating these rules.