February 2020

Classification of source IP addresses for dynamic access control for cloud resources

Dhandapani Shanmugam

Classification of source IP addresses for dynamic access control for cloud resources

ABSTRACT

Current security and access control technologies permit controlling authorization and authentication based on network addresses. However, such technologies do not address problems resulting from administrator error and oversight. Inadvertent mistakes and/or oversight can result in cloud resources being erroneously exposed to external access by malicious actors. This disclosure describes the application of a trained machine learning model to determine a credibility score for the source IP address of a request received at a cloud resource. The outcome of the credibility evaluation is used to enact access controls for the corresponding source IP address.

KEYWORDS

- Cloud platform
- Cloud resources
- Cloud security
- Access control
- Edge system
- Router
- Source IP address
- Credibility score
- IP characteristics
- IP attributes

BACKGROUND

Cloud platform operators need to implement appropriate access controls to manage access to cloud resources allocated to customers. Current technologies permit controlling authorization and authentication based on Internet Protocol (IP) addresses. Such IP address based controls may operate at the application and/or network levels. In addition, firewalls and DMZs use IP addresses to control network traffic passing to and from the cloud resources.

However, these technologies do not address issues resulting from administrator error, such as inaccurate specification of authorization options, lack of adherence to standards and best practices, etc. Moreover, humans tasked with managing access control may sometimes forget to perform one or more essential steps in the task. For instance, a person can forget to whitelist trusted IP addresses, turn on authentication, configure applications, etc. Inadvertent mistakes and/or oversight can result in resources on the cloud platform being erroneously exposed to external access. In such cases, malicious actors may be able to gain access to the data, applications, or other resources on the cloud platform.

DESCRIPTION

This disclosure describes the use of a trained machine learning model implemented on systems at the edge of a cloud platform to enhance security. For example, such edge systems can include destination routers. The model is applied to analyze a given source IP address in terms of various attributes, such as location, frequency of access, time of day, historical score, etc. Based on these attributes, a credibility score for the source IP address is determined. If the score is above a high threshold value, the source IP address is classified as good, and if the score is below a low threshold value, the source IP address is classified as bad. For cases where score falls

between the two threshold values, the source IP is marked for additional probing since a classification for the source IP address is not ascertained with sufficient certainty.

The above described credibility classification is performed for the IP address of every external and internal system that attempts to access resources on the cloud platform. A previously unseen IP address starts with the minimum credibility score of 0. The score is updated based on usage patterns of the IP address for requesting access to cloud resources. For instance, daemons or bots may receive positive scores only after receiving an explicit endorsement from the edge system in the cloud network. The outcome of the credibility evaluation is used to enact access controls for the corresponding source IP address. If the source IP address is classified as good, access to the requested cloud resources is granted, and if the source IP address is determined to be bad, access is denied.

A public IP address within the cloud platform is associated with the location and time zone of the source IP addresses. Requests from source IPs that deviate from previously associated locations and/or time zones are probed to determine credibility. Such probes can involve questions related to the nature of the resource request, such as application name, purpose, etc. For instance, such questions can be presented to the human users at a device associated with the source IP addresses, e.g., in the form of CAPTCHAs. The responses to the probe are recorded and, in turn, used to classify the source IP as good or bad. Once the source IP is classified, access to the requested cloud resource is granted or denied as appropriate.

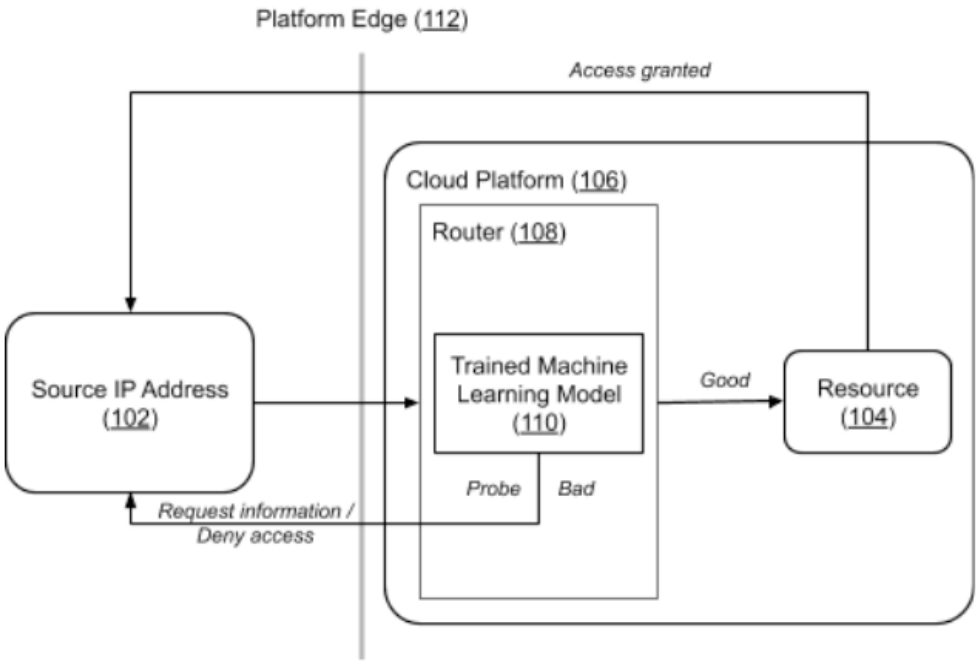

Fig. 1: Granting access to cloud resources based on dynamic credibility assessment

Fig. 1 illustrates a device at a source IP address (102) requesting access to a resource (104) hosted within a cloud platform (106). The request goes through the router (108) operating at the edge (112) of the cloud platform. A trained machine learning model (110) implemented within the router is used to classify the source IP address. If the classification indicates that the source IP address is considered good, access to the resource is granted. On the other hand, if the source IP is determined to be bad, access is denied. In case of uncertainty, the source IP address is probed for additional information to help determine its credibility for the resource request in question. In this manner, the behavior of the cloud platform in granting or denying access is different based on classification of the source IP address. Thus, the cloud public IP address that is associated with the resource exhibits different behavior for different requesting source IP addresses, in essence acting like a chameleon.

The credibility classification described in this disclosure permits differential access to cloud resources for different source IP addresses and corresponding end users. Moreover, the

described techniques operate automatically and dynamically, thus guarding against inadvertent mistakes or oversight from human administrators charged with managing access to cloud resources. For instance, the operation described above can be applied to flag malicious accesses to publicly available cloud resources in cases where the resource is publicly accessible because of misconfiguration.

The high and low threshold values corresponding to the good and bad classification for the source IP addresses can be specified by the system, set dynamically during operation, and/or configured by the cloud administrators. The described techniques can be implemented within any cloud platform, associated edge devices and applications.

CONCLUSION

This disclosure describes the application of a trained machine learning model to determine a credibility score for the source IP address of a request received at a cloud resource. The outcome of the credibility evaluation is used to enact access controls for the corresponding source IP address. The credibility classification permits differential access to cloud resources for different source IP addresses and corresponding end users. Moreover, the described techniques operate automatically and dynamically, thus guarding against human errors in configuration of a cloud resource. The described techniques can be implemented within any cloud platform and associated edge devices and applications.

REFERENCES

1. Lund, Peter K., Scott M. Petry, Craig S. Croteau, Kenneth K. Okumura, and Dorion A. Carroll. "Electronic message source reputation information system." U.S. Patent 7,668,951, issued February 23, 2010.