Intrusion Detection/Prevention Systems in the Cloud

Joseph Johann

ICTN6875

East Carolina University

Abstract

With more and more organizations moving all or part of their infrastructures to the cloud it makes sense that they would also have many of the same security services in the cloud that protect their on premise data. These services could include standard firewalls, intrusion detection/prevention systems (IDPS), antivirus solutions, data loss prevention (DLP) systems, and web application firewalls just to name a few, The concentration of this paper will be to discuss in detail intrusion detection/ prevention systems operating in the cloud to protect customer data as well as the advantages and disadvantages of doing so. It will also discuss the challenges faced by organizations in implementing and managing this service. Finally the costs will be compared to determine whether putting this security service in the cloud is cost effective.

The move to cloud services has made information available on-demand and more convenient to its users due to the ease of access usually through browsers. The cloud also usually provides redundancy and disaster recovery for organizations that couldn't provide it for themselves. With the gathering of all data into one location comes the increased risk of attacks to the service provider and the data they store. Modi et al (2013) stated in their article that the major security concern with the cloud after data security is intrusion detection and prevention. There are three main types of techniques for detecting attacks: anomaly detection, signature-based detection, and a hybrid version. Anomaly detection is used to detect attacks that are currently unknown to the research world. This type of detection analyzes what is considered to be the normal behavior of device or network. Anything that isn't considered normal behavior is flagged as an anomaly. This type of detection requires a period of time for the system to learn behavior. One disadvantage of this type of detection is that if any malicious traffic is captured during the learning period it is categorized as normal thus making your environment unsecure. Signature-based detection relies on characteristics of previously discovered attacks. The disadvantage of this type of detection is that there are many new exploits that haven't had signatures written for them yet this they won't get detected. Hybrid detection takes the best features of these two and combines them together to detect known and unknown attacks. The attacks on cloud services can vary with severity and the type of service they target including confidentiality, integrity, and availability.

One type of attack is the insider attack. This is where individuals with authorized access either abuse or disclose information to individuals who should not have access to the data. This can be very serious especially when it includes healthcare or financial information and can lead to identity theft. Another attack is the denial of service attack where someone tries to flood the

victim with a large number of packets in an effort to either make a service unavailable or to bypass security services that have been rendered useless. The next attack is the escalation of privileges attack whereby a user's password is compromised usually by sniffing the network. The attacker then attempts to gain privileged access such as root in an effort to gain control of the entire system. Port scanning can also be used on cloud services to determine which ports/services are listening and accepting connections. Once an open service is found the attacker can then try to manipulate the data stream into that port to make the service believe it is legitimate data. This can lead to the access and manipulation of data. Backdoor channels can also be set up within the environment to allow outside users access in usually through a compromised host. This type of access is normally preempted by an attack that successfully gains full access to a host. Finally, since many cloud infrastructures are hosted on virtual environments, an attack to the hypervisor can give access to any of the hosts which live on the system.

Intrusion detection/prevention systems are an effective way to mitigate threats both internally and externally. Some threats are malicious in nature while others are the result of user error such as mistyping a command or ip address. An intrusion detection system is commonly set up to only monitor traffic with very little intervention. They usually receive a copy of the traffic through a span port or network tap. This allows for the reporting of suspect incidents for investigation but allows the traffic to continue uninterrupted to its destination. The converse of detection is prevention. In intrusion prevention systems the traffic usually flows through the device allowing it to stop anything that it deems malicious. These types of systems require testing and fine tuning to ensure that normal traffic isn't being stopped. Like IDSs, IPSs need to be monitored and incidents need to be investigated in a timely manner.

There are several different types of IDPSs including network-based (NIDPS), host-based (HIDPS), and application-based (AIDPS). NIDPSs monitor network traffic for a particular segment or devices and analyze the traffic for malicious activity. HIDPSs concentrate on an individual system and monitor the behavior and state of that system. Anything that isn't considered normal is flagged for investigation. A NIDPS and a HIDPS can work together to create a better understanding of the traffic being collected. An AIDPS concentrates on the activity of an application by reviewing log files and system performance to determine deviations in normal activity. A distributed IDPS (DIDPS) is a combination of the previous types of systems dispersed throughout the network with the ability to communicate with each other or with a centralized system. This allows information to be shared giving the systems a better understanding of the environment as a whole. This can help the systems be more proactive by knowing the types of incidents that are happening on other parts of the network.

The deployments of IDPSs in standard organizations are very common and have been designed to meet the challenges they are presented with although the technology and attack vectors are evolving daily. The cloud presents different challenges for these devices in both their utilization and deployment since its unique features and architectural issues combine with the standard issues found in most non-cloud systems. Patel et al (2013) state that the identification of the exact characteristics of a target environment is essential to establishing system requirements and system development. They list the characteristics of cloud computing systems as elasticity which involves the scalability of the systems, reliability as the ability to ensure the continued operations of the systems, quality of service (QoS) which is related to meeting the requirements of service level agreements (SLAs), agility and adaptability which refer to being able to handle changes in the systems related to workload and the environment itself, and

availability which establishes redundant services and data protection in the event of a failure. According to Patel et al (2013) there are no traditional IDPSs to meet these characteristics efficiently.

Along with the characteristics comes an array of challenges of deploying an IDPS in a cloud environment. It is very important to understand these challenges before deciding on the IDPS to deploy. First, since the monitored machines are virtual in nature, they can be dynamically added and removed which tends to change their security requirements. Traditional systems have security policies that represent the organization that is being protected and are managed by a security administrator responsible for that organization. In the cloud, there are many security administrators monitoring many unique environments which could lead to longer response times for incidents. These administrators don't have the intimate knowledge of the environments like local administrators do. There is also concern for the shared environment that houses the different customers. If the host of a malicious customer is added to the virtual environment, this can lead to insider attacks within the virtual space. Being that most of these environments are utilizing virtualization technology, any weaknesses in the hypervisor can lead to a compromise of any of the hosted sessions. There is also concern about insider attacks within the hosting companies. Since there are so many different organizations being hosted within the same virtual space there could be confusion as to the policies that apply to each. The cost of the data being transferred within the environment can also be a challenge since IDPS solutions often generate additional traffic. The monitoring of communications between virtual environments can present difficulties since most virtual environments utilize virtual switches within the software which don't presently have an effective way for monitoring devices to inject themselves between virtual hosts. Traffic is often rerouted outside the environment then back in to

accomplish this task which can waste valuable resources. Finally, most organizations need to provide some level of accountability as to the risk level of the hosted environments. Many cloud service providers combine their security measures in an effort to reduce their costs. This often leads to their inability to provide personalized security logs and audit data for each customer leaving them with reduced visibility into their risk.

Alsafi et al. (2012) discuss a method where the use of multithreading techniques provides a more efficient method for improving the performance of an IDPS in a cloud environment. This method enables the systems to handle a large number of data packet flows. It consists of a module that captures the packets then subsequently forwards them to an analysis module. This module processes the packets and filters out any bad packets which leaves the alerts to be analyzed. The reporting module then takes these alerts and creates a report which security analysts can use to determine the severity and accuracy of the data. The data can also be fed into a Security Event and Incident Management (SEIM) system which will collate the information with alerts from other systems to get a better idea of the details of the alert. The authors were able to test the single threaded system against the multithreaded system and received the results they had expected including improved processing and execution time. They concluded that more work needed to be done in the actual stopping of attacks rather than just the detection. Their recommendation for deploying an IDPS in the cloud is to place it behind the firewall protecting the cloud environment. This will allow the firewall to perform the basic blocking functions it is intended for thereby reducing the amount of traffic reaching the IDPS. Furthermore, the IDPS should be located within the virtual environment itself so as to be as close as possible to the resources it is trying to protect. They also recommend using the hybrid detection technique

discussed earlier which consists of anomaly and signature based detection as well as detection and prevention in the IDPS itself.

Jin et al. (2011) proposed an intrusion prevention system called VMFence that resides in a virtualized cloud computing environment. This system is used to monitor both network traffic flow but also file integrity monitoring in real time. File integrity monitoring gives the traditional IDPS extra information as to the changes that are taking place to files on the systems themselves. In testing the system it provided effective results with acceptable overhead. The detection processes are placed in privileged VMs so as to be able to monitor traffic to and from other VMs. A detection process will automatically start whenever a new VM is activated. The privileged VMs only control and manage the security of the cloud. No other user applications are running on it because privileged VMs need to be as secure as possible with the fewest number of vulnerability access points available to threats. The file integrity portion of VMFence plays a critical role in that any changes to the IDPS system itself can be detected. VMFence consists of five main components. The detection component is responsible for capturing all of the network packets that travel from system to system then to send them to other detection processes according to their MAC addresses. There are many detection processes and each has a task according to the rules of the corresponding VM. Each VM has default rules as well as customized rules. The policy updating component is used for intrusion response modifications. All of the alerts are collected from the detection processes by this component which also notifies the administrator. The frontend and backend communication component allows for channels of communications between these areas. When new policies are created for a particular VM, this information is then sent from the backend component to the frontend component with the backend component making the global decisions for the environment. Files that are read from

and written to in the backend are protected by the file integrity monitoring component. This component sends notifications if there are any modifications. Finally, the notification component collects basic information about the VMs servicing cloud users and sends it to them as well as notifies the cloud provider of alert information. The DARPA98 intrusion detection dataset is used in order to simulate a real network environment. VMFence has proven in tests to provide high availability, high performance, and real time monitoring in a cloud environment.

To purchase many of the name brand IDPS appliances it can cost anywhere from $15,000 to $50,000 or higher which is much more than small or medium companies can afford. The cloud offers an alternative utilizing Security-as-a-Service (SaaS). This not only can be less expensive but can deliver in areas where hardware –based security products have limitations. For example, many organizations don't have the expertise in house to manage the systems and analyze the alerts. Cloud solutions can provide the right protection you need without the hardware investment. It doesn't require the training of your staff or additional personnel. Also if your requirements grow there is no need to invest in another expensive in house solution. You can simply upgrade your current solution to the next platform with a minimal monthly cost increase. Alternatively, the cost of IDPS appliances that are purchased are normally a small portion of the costs compared to the annual maintenance fees and support fees. Cloud based systems usually also provide disaster recovery built-in with the solution. Brick and mortar operations usually rent expensive locations which are never fully utilized. With the rapid development of new technologies, IDPS architectures can become obsolete in a very short period of time. There are also expenses related to power and air conditioning that are not needed in cloud environments. Cloud systems also allow some organizations to pass of the responsibility to 3rd parties when it comes to meeting regulatory requirements which can reduce the in house

expertise needed to perform these audits. However, organizations also lose control over their data and the security systems that protect it.

With the growing trend to move everything from email to data into the cloud, there is still a requirement to protect that information and guard it against malicious attacks. Just the very nature of cloud environments makes them an attractive target to hackers looking to expose personal, medical, and financial data. The Security-as-a-Service (SaaS) trend has moved many security functions into the cloud where they are needed to provide the same type of protection as those in standard network infrastructures. Intrusion detection and prevention systems are part of the base security measures needed. With the use of virtualization in cloud environments and the hosting of virtual machines containing data from many customers, it is imperative to be able to detect any malicious activity leaving one environment and trying to access others. The IDPS needs to be combined with other defense mechanisms to enhance the overall posture of the environment and ensure the data is protected fully protected from all variations of attacks. As attacks become more sophisticated the IDPSs in the cloud need to evolve along with them.

References

Modi, C., Patel, D., Borisaniya, B., Patel, H., Patel, A., Rajarajan, M. (2013). A Survey of Intrusion Detection Techniques in Cloud. *Journal of Network and Computer Applications*. June 2013, pp 42-57.

Patel, A., Taghavi, M., Bakhtiyari, K., Junior, J.C. (2013). An Intrusion Detection and Prevention System in Cloud Computing: A Systematic Review. *Journal of Network and Computer Applications*. January 2013, Vol. 36, Issue 1, pp 25-41.

Jin, H., Xiang, G., Zou, D., Wu, S., Zhao, F., Li, M., Zheng, W. (2011). A VMM-based Intrusion Prevention System in Cloud Computing Environment. *The Journal of Supercomputing*. December 2013, Vol. 66, Issue 3, pp 1133-1151.

Alsafi, H.M., Abduallah, W.M., Pathan, A.K. (2012). IDPS: An Integrated Intrusion Handling Model for Cloud Computing Environment. *International Journal of Computing and Information Technology*. Vol. 4, Issue 1, pp 1-16.

Modi, C., Patel, D., Borisaniya, B., Patel, A., Rajarajan, M. (2012). A Survey on Security Issues and Solutions at Different Layers of Cloud Computing. *The Journal of Supercomputing*. February 2013, Vol. 63, Issue 2, pp 561-592.

Pandian, V.A., Kumar, T.G. (2014). A Novel Cloud Based NIDPS for Smartphones. *Recent Trends in Computer Networks and Distributed Systems Security.* Vol. 420, pp. 473-484.

Fernandes, D.A.B., Soares, L.F.B., Gomes, J.V., Freire, M.M., Inacio, P.R.M. (2014). Security Issues in Cloud Environments: A Survey. *International Journal of Information Security*. April 2014, Vol. 13, Issue 2, pp. 113-170.

Ramachandran, M. (2011). Component-Based Development for Cloud Computing Architectures. *Cloud Computing for Enterprise Architectures, Computer Communications, and Networks*. pp 91-114.

Hill, R., Hirsh, L., Lake, P., Moshiri, S. (2013). Cloud Security and Governance. *Guide to Cloud Computing, Computer Communications, and Networks*. pp 223-239.

Soares, L.F.B., Fernandes, D.A.B, Gomes, J.V., Freire, M.M., Inacio, P.R.M. Cloud Security: State of the Art. *Security, Privacy, and Trust in Cloud Systems*. pp 3-44.

Lee, H., Kim, J., Lee, Y., Won, D. (2012). Security Issues and Threats According to the Attribute of Cloud Computing. *Computer Applications for Security, Control, and System Engineering*. Vol. 339, pp 101-108.