Attivo NETWORKS®

Enhancing AWS Cloud Security with Deception Technology

As businesses adopt cloud computing at an increasing pace, they realize the benefits of on-demand, scalable infrastructure. They also find themselves relieved of some security concerns, such as managing the physical security of data centers. Unfortunately, the same threats that plague on-premise devices and systems are just as problematic in the cloud. AWS provides a range of security controls to help protect the confidentiality, integrity and availability of applications, data, and devices. While necessary, these controls are not sufficient to address sophisticated advanced persistent threats. AWS customers can, however, enhance the overall level of security protections by deploying active deception technologies in their AWS infrastructure.

This whitepaper provides an overview of AWS security controls and best practices, with particular emphasis on the responsibilities of customers. The controls and practices are assessed with respect to advanced persistent threats, and their limitations are discussed. Active deception technology is introduced and shown to complement other security controls and measures, especially with regards to protecting against advanced persistent threats. The whitepaper concludes with tips on deploying active deception technology to the AWS cloud.

AWS Security Controls and Best Practices

AWS security controls range from broad measures, such as virtual private clouds and network segments, to fine-grained access controls on storage and compute resources. Cloud administrators and information security professionals select and combine controls appropriate for their applications and business requirements. Commonly used controls include:

- Virtual Private Clouds (VPC), which isolate AWS resources in a virtual data center
- Network segments, for isolating network traffic
- Security Groups, to control the ingress and egress flow of traffic to servers
- Network access control lists (NACLs), to limit access to resources within network segments
- Identity management, for granting privileges to users and groups
- Configuration controls, to help ensure deployed devices meet configuration requirements

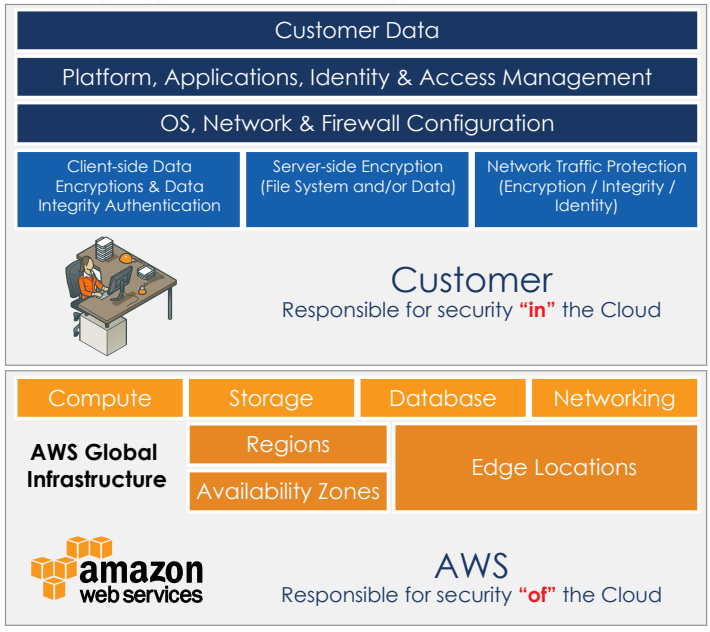

Figure 1. Amazon Shared Security Model

In addition to these security controls, Amazon publishes a set of security best practices that include:

- Categorizing and protecting assets on the AWS cloud
- Managing users, groups, and roles
- Maintaining OS-level security
- Securing data
- Monitoring and auditing

For organizations that need additional measures, Amazon recommends a threat protection layer that could include third-party firewalls, unified threat management systems, data loss prevention systems, and advanced persistent threat protection measures. Active deception technologies are an increasingly important type of advanced persistent threat protection.

Advanced Persistent Threats and the Need for Dynamic Deception

As the name implies, advanced persistent threats are not easily controlled by typical detection and blocking measures. These types of attacks employ multiple techniques over extended periods of time. They may start with social engineering attacks that lure victims into disclosing login credentials or downloading malicious software or with the use of stolen credentials. Regardless of how a particular attack starts, there are common characteristics of advanced persistent threats that make them particularly difficult to detect.

Methods of Attack

Once attackers gain a foothold on the networks, they can probe servers and scan network traffic looking for exploitable vulnerabilities. They may, for example, download vulnerability scanners to look for known weaknesses in an application stack. In addition, an attacker might probe devices on the network cataloging software deployed to different servers and mapping the network configuration. It is sometimes difficult to distinguish attackers probing behavior from legitimate systems behavior. This is especially the case when attackers work slowly to minimize the risk of triggering intrusion detection alerts or leaving obvious indicators of malicious activity in system logs.

Attackers may take less than direct routes to their target. For example, an attacker seeking to extract a large volume of data could directly attack a database server. This may work if there is a significant weakness in the database software or its configuration that allows for rapid extraction of data. Databases that run on hardened operating systems and limit network traffic to a small set of application servers are not so easily compromised. Instead of targeting the database server directly, an attacker may first probe application servers or Web servers for vulnerable application programming interfaces (APIs). This type of attack that starts at the top of the application stack can succeed even if each layer of the stack is hardened and secured. The problem stems from the fact that to be of any use at all; each layer must communicate with other layers as well as end users.

Consider a simple Web application. The Web server may be configured to accept traffic only on ports 80 (HTTP) and 443 (HTTPS). The application server firewall rules accept incoming traffic only from the Web server. Similarly, the database server may be configured to only accept traffic from the application server on the port assigned to the database listener. This type of configuration limits the paths an attacker may take to reach the database, but it does provide a potential path nonetheless. If attackers can deliver malicious payload using application protocols or use stolen credentials, they may be able to breach multiple layers of defenses to reach their target.

Detecting and Analyzing Attack Patterns

Advanced persistent threats are tailored to the particulars of a target. Attackers may use a common set of tools, such as vulnerability scanners and command and control programs, but they are applied in different ways in different attacks. For example, rather than probe and attack starting at the presentation layer of an application stack and working down to the data services layer, an attacker may probe laterally across servers or other devices within a network segment. Also, attackers may use different tactics at various points of the attack. The sequence of steps in an attack, known as the kill chain, can vary as attackers seek to exploit vulnerabilities found at each stage of an attack.

The wide variety of ways an attack can proceed makes it difficult to discern advanced persistent threat activities. For example, most intrusion detection systems try to detect variations from known signatures and baseline patterns of activity. This is useful in detecting obvious anomalies, such as large downloads from a database server at unexpected times. More subtle attack activities require a different approach to detection.

Rather than try to distinguish malicious activities from legitimate operations on production servers, active deception technologies lure attackers into revealing their tactics using decoy devices. Early deception technology was virtually synonymous with honey pots, or servers designed to look like authentic targets but are actually decoys. Honey pots and uses of emulated servers are used to keep attackers away from production servers while providing an opportunity to collect information about attack patterns. Not surprisingly, attackers have devised ways to detect honey pots and emulated servers and avoid them.

More realistic deception technologies actively engage attackers and provide the type of experience an attacker might expect from actual production servers. This is particularly important for detecting attacks that begin with reconnaissance. The Attivo Networks BOTsink platform, for example, lures attackers by hosting network services on a variety of virtual devices, by placing deception credentials on endpoints and deception lures on its engagement servers. By deploying detection servers that look and function like production servers, attackers are tricked into applying their tactics while monitored by the BOTsink platform. Lateral movements within the network, known as east-west traffic, are easily detected even when the attacker uses slow, low volume, stealth attack patterns. This method has become particularly attractive for data center or other environments with large volumes of traffic whereas monitoring can be expensive to implement.

Active deception technologies complement other security measures that are designed to block malicious activity. Rather than just preventing the progression of an attack, active deception collects information about attacks and compromised devices, which can be used to identify previously unknown vulnerabilities in systems. Also, since active deception does not depend on predefined signatures, it is an effective means to counter zero day vulnerabilities, the use of stolen credentials, and the risk of phishing lures. Dynamic deception goes one step further in authenticity and refreshes deception lures, provides the ability to turn on and off services, and makes other dynamic changes for the highest degree of deception.

Deploying Dynamic Deception Technology in the AWS Cloud

Dynamic deception technologies from Attivo Networks® are designed to integrate seamlessly with AWS deployments and to scale with your needs. To maximize the benefits of deception technology, be sure to consider three key aspects of effective deployment:

- Comprehensive coverage
- Visibility
- Reporting and analysis
- Realistic, golden images

A properly deployed set of dynamic deception servers can provide a new level of protection not available from AWS security measures alone.

Comprehensive Coverage

With the Attivo ThreatDefend™ deception platform, instances can be deployed across multiple network segments, availability zones and regions to provide comprehensive coverage of your AWS infrastructure. It is important to remember that advanced persistent threats are by their nature highly adaptive. Patterns and tactics used in one attack may be altered in subsequent attacks. An attacker that exploits a north-south path through the application stack in one case might use a more lateral, east-west type of attack pattern in later attacks.

Attacks may be tailored to the type of instances in network segments or virtual private clouds. For example, an attacker may use fuzzers in an application server network segment to search for vulnerable APIs while employing more specialized tools to compromise database servers in the data services tier. It is important to have active deception operating in all network segments. It is not safe to assume that monitoring and active deception in one network segment is sufficient to protect other network segments.

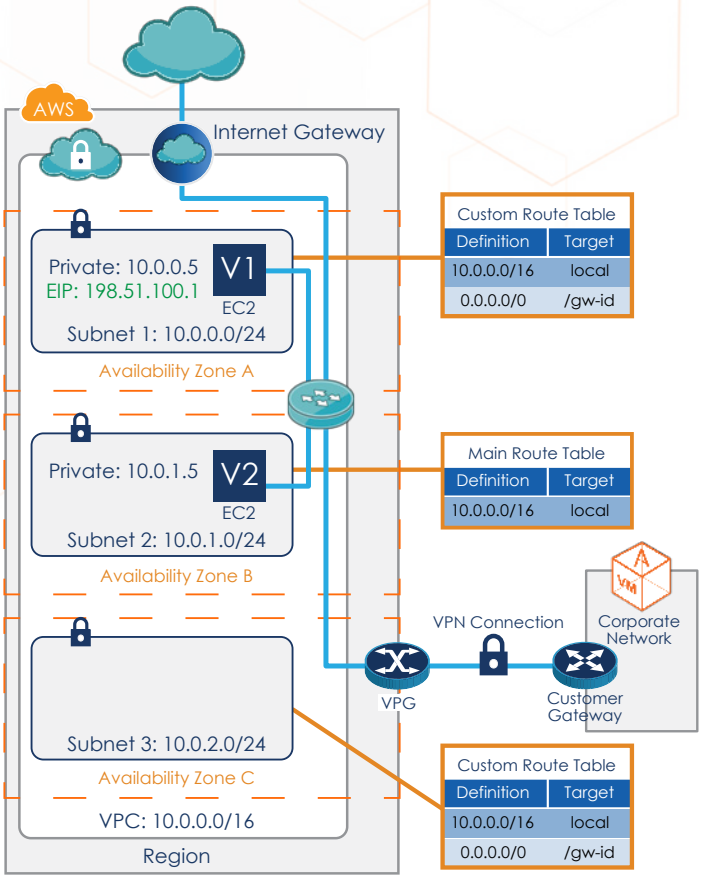

Patient, methodical, attackers will probe for openings in your defenses and exploit vulnerabilities in potentially unanticipated ways. Comprehensive coverage by dynamic deception is essential to countering this threat.

Visibility

The Attivo ThreatDefend™ deception platform provides visibility into inside-the-network threats across enterprise, public, private, and cloud environments. The combination of the the Attivo ThreatDefend deception server and ThreatStrike End-point Deception Suite provide the deception platform to detect attacks that come in from the use of stolen credentials, attacker reconnaissance, and ransomware or phishing attacks. Additionally, insider threats both intentional and accidental, along with unauthorized movement from 3rd parties can be detected by the Attivo deception platform providing visibility into attacks that in most cases cannot be picked up by traditional detection methods.

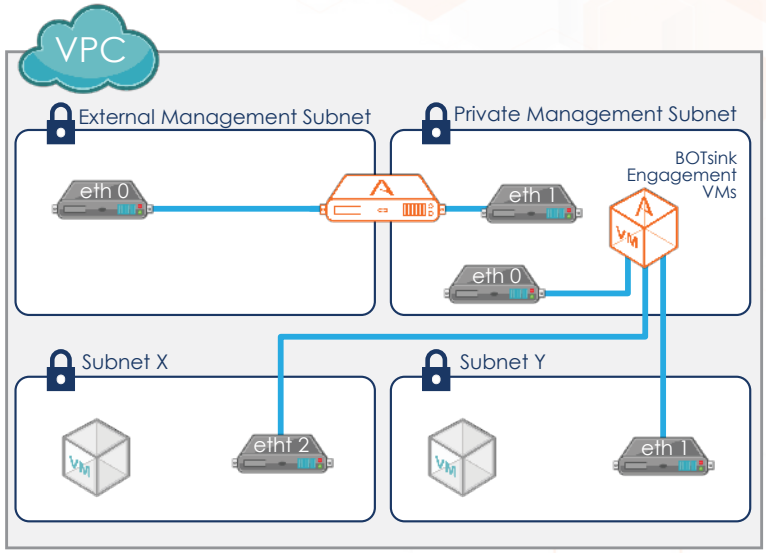

Figure 4. BOTsink deception platform detects intrusion, engages the attacker, and generates attack forensics.

Reporting and Analysis

The Attivo Networks dynamic deception technologies include easy to use reporting and analysis tools and a centralized threat intelligence dashboard. These allow cloud administrators and information security professionals visibility into activity on their networks, to assess threats in real time and to collect forensic data needed to shut down attacks and improve the overall security of one's cloud infrastructure.

Realistic Golden Machine Images

One of the weaknesses of traditional honey pots or deception that uses emulation is that they do not always appear to be actual production servers. With Attivo Networks technology, customers can deploy dynamic deception virtual machines with a variety of choices in operating systems, protocols, and services. For the greatest level of authenticity the same golden images used for production servers can also be loaded as the ThreatDefend image. This reduces the chances that an attacker would bypass an otherwise tempting target because the server looks suspiciously like a honey pot or an emulated server.

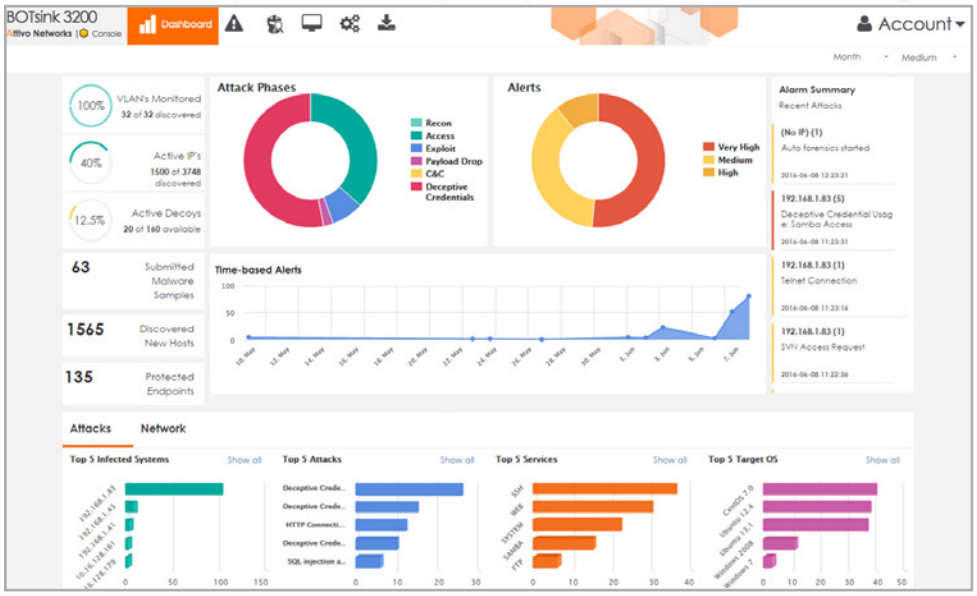

Figure 5. The Attivo Networks Threat Intelligence Dashboard provides a single point of access to information about the state of deception servers.

Conclusion

Advanced persistent threats pose some of the greatest challenges in information security. Traditional block and tackle security measures are insufficient to protect against patient, methodical, and determined attackers. Fortunately, recent advances through dynamic deception offer significant protections to the increasingly costly consequences of advanced persistent threats.

About Attivo Networks

Attivo Networks® provides the real-time detection and analysis of inside-the-network threats. The Attivo ThreatDefend™ Deception and Response Platform detects stolen credentials, ransomware, and targeted attacks within user networks, data centers, clouds, SCADA, and IoT environments by deceiving an attacker into revealing themselves. Comprehensive attack analysis and actionable alerts empower accelerated incident response. www.attivonetworks.com