Active defence through deceptive IPS

Authors
Apostolis Machas, MSc (Royal Holloway, 2016)
Peter Komisarczuk, ISG, Royal Holloway

Abstract

Modern security mechanisms such as Unified Threat Management (UTM), Next-Generation Firewalls and Security Information and Event Management (SIEM) have become more sophisticated over recent years, promising advanced security features and immediate mitigation of the most advanced threats. While this appears promising, in practice even this cutting-edge technology often fails to protect modern organisations as they are being targeted by attacks that were previously unknown to the security industry. Most security mechanisms are based on a database of previously known attack artefacts (signatures) and they will fail on slightly modified or new attacks.
The need for threat intelligence is in complete contrast with the way current security solutions are responding to the threats they identify, as they immediately block them without attempting to acquire any further information. In this report, we present and evaluate a security mechanism that operates as an intrusion prevention system which uses honeypots to deceive an attacker, prevent a security breach and which allows the potential acquisition of intelligence on each intrusion attempt. [a]

[a] This article is published online by Computer Weekly as part of the 2017 Royal Holloway information security thesis series `http://www.computerweekly.com/ehandbook/Active-defence-through-deceptive-IPS`. It is based on an MSc dissertation written as part of the MSc in Information Security at the ISG, Royal Holloway, University of London. The full thesis is published on the ISG's website at `https://www.royalholloway.ac.uk/isg/`.

Introduction

There seems to be a constant increase in the occurrence of security breaches for large and small organizations. As companies become more and more dependent on digital systems for their stable operation, the impact that a security incident may have on their business significantly increases. Common security mechanisms and techniques have proved effective to some extent against low-level attackers. However, they may still fail against advanced and persistent attackers. One of the reasons for failure is that most security mechanisms are effective only against known attacks and remain vulnerable to new or modified attacks. The security community needs to identify modifications of already known attacks, and remain informed about new vulnerabilities, techniques and tools.

As most common security mechanisms will block an attack at its early stage, security analysts are not able to obtain any further information about the attack's next steps. This information could be used to identify new vulnerabilities and to update the deployed security mechanisms. There are vast industry networks to gather intrusion data, where a variety of sensors are deployed. There are also other initiatives that share information about attacks which help update intrusion detection signatures and advise cyber security professionals. Some organisations take part in the deployment of sensors and make a decision on the trade-off between the value of information being collected and the level of risk that the organisation is willing to face.

In this article we present and evaluate a security mechanism, developed as a masters level project proof of concept, that acts as an intrusion prevention system (IPS). It combines two popular security mechanisms, honeypots and intrusion detection and prevention systems. This mechanism uses an intrusion detection engine and attempts to detect an attack in its early stages and redirect it into a honeynet. The honeynet will then allow the attacker to carry on the attack in an environment where they

can be observed, providing valuable information to the security community while access to the original protected network will be blocked. If the honeynet appears similar to the original network, attackers will falsely believe that they have successfully infiltrated the organisation. The amount of information gained is dependent on the sophistication of the attacker and the realistic nature of the honeynet. The system developed can be described as an advanced active defence security mechanism that can produce cyber threat intelligence feeds.

Cyber threat intelligence

While the immediate identification of a security breach is important for an organization, it still does not prevent the occurrence of it. In order for the security analysts to prevent a security breach, they have to win a race against the attackers as they detect and respond to the attack before it succeeds. It is an arms race and the attackers' techniques are becoming more and more sophisticated. It can be argued that the attackers are currently in the lead but with wider industry cooperation this could be reduced or reversed.

Most information security solutions focus on the denial of access to sensitive information or services to an attacker. Firewalls, antivirus software, intrusion prevention systems and various other security mechanisms usually operate on a "Yes or No" rule set allowing or preventing access to the protected resources. While this approach has proved effective against common attacks and unskilled individuals, it can fail to protect an organisation against determined sophisticated attackers.

Intrusion Prevention System (IPS)

A security mechanism, either hardware or network based, that attempts to detect and prevent intrusions by analysing and responding to network and/or system events.

Cyber threat intelligence (CTI) is widely used by the IT security community to accurately detect and respond to emerging threats in a timely manner. Even though attacks are constantly evolving, attackers are frequently using pieces of known malware and tools that in some part can be detected. Moreover, exploit kits, malware and botnets being sold as-a-service are flooding the black hat market as their effectiveness is proven. Security vendors, based on the events collected by sensors during security incidents, publish intelligence feeds to their customers that will help them to identify and respond to a potential attack.

Sophisticated attackers are utilizing techniques such as specifically crafted phishing emails, exploiting weak authentication credentials and social engineering, targeting directly the end-points behind the security perimeter of the organisation. These can successfully evade the conventional security mechanisms and require a more intelligent approach and a further understanding of the attackers' background than typical CTI can provide. Thus, the defending organizations require advanced mechanisms that will enable them to increase the timeframe of the investigators to respond, disrupt the attacker, and block the attack before it succeeds.

(Pro)Active cyber defence & deception

While the term "active cyber defence" (ACD) has been discussed widely by information security researchers, government agencies and organisations, the scope of offensive operations that can take place beyond the defender's network is unclear. The U.S Defence Advanced Research Project Agency (DARPA) describes its ACD research program as a strictly defensive program that nonetheless involves direct engagement with the attackers.

On the other hand, Ernst & Young describes "active defence" as a continuous process that does not interact with the attacker's network. Instead it focuses on the organisation's defence from carefully designed threat scenarios and the continuous "hunting" for attackers that have infiltrated the corpo-

rate network. In this report, we follow the definition of ACD as a real-time defensive strategy, tightly connected with CTI, that focuses on the detection, analysis and mitigation of cyber-attacks by utilizing techniques and mechanisms that can disrupt or deceive the attackers before they manage to achieve their goals.

Whether offensive or not, ACD strongly relies on cyber threat intelligence for the detection of sophisticated attacks while its operations usually produce vital intelligence for the mitigation of future cyber-threats. Deception, as a non-offensive implementation of ACD, can significantly stall and confuse the attackers during an intrusion attempt and usually exceeds a common honeypot deployment. Its main focus is the protection of the organisation's network and services against sophisticated attackers and threats. As a result, it has to be carefully designed to deceive and detect even the most skilled attackers.

Honeypot

A system or a service that has been deliberately appears vulnerable, but is strictly constrained and monitored, so that once attacked, it will collect significant information about attacker's actions and tools. Based on the level of interaction a honeypot provides to an attacker it can be categorized as a low, medium or high interaction honeypot.

Deception appears to be one of the most promising defensive techniques not only due to its effectiveness against advanced threats but also due to its capability to produce intelligence in a more detailed and accurate way than common security mechanisms do.

Experimental solution

Our research focused on the creation and test (a basic proof of concept) of an intrusion detection engine to redirect malicious traffic to honeypots without being limited to specific attacks, systems or services. The proposed system can be classified as an advanced active defence mechanism that attempts to benefit from the advantages that both honeypots and intrusion detection systems provide. However, it should be noted that the aim of this research is not to provide an alternative to the existing intrusion detection and prevention systems or honeypots but to emphasize the potential benefits when deployed in collaboration.

Our mechanism focuses on creating an attack-driven deception that will delay an attack as much as possible, giving more time for the security analyst to respond to and prevent the attack. Once an attack is detected the malicious packet will be dropped, and all future traffic originating from the same attacker will be routed to a honeynet. Traffic originating from legitimate users will continue to be served by the original production system. All traffic to the honeynet will be monitored but not dropped, ensuring the transparency of our mechanism and increasing the time cost to the attacker. The use of honeypots and appropriate monitoring mechanisms provides information about the attacker methods, motives and tools that can be used to further tune the organisation's security mechanisms and assist the security analysts in the correct classification of the incident.

Snort

One of the most popular and configurable open-source intrusion detection and prevention systems. It can be installed on a gateway or on a single host and it can operate either in a passive or active mode. In its passive mode, it will only alert on detection of a malicious event without taking any further action, while in its active mode it will block or drop malicious traffic.

As a proof of concept we implement our system using Snort, a popular open-source signature based network intrusion detection and prevention system, and we evaluate it in a small simulated corporate network. In our implementation, Snort is implemented as a centralized network intrusion detection system able to detect an attack during the reconnaissance step (port, network or vulnerability scans) and on detection redirect the attacker to the decoy honeynet. Monitoring the whole network traffic, instead of a single instance provides a higher level of visibility enabling also the detection of worms

or network scans. Moreover, the proposed system introduces minimal overhead to a traditional Snort implementation and does not limit the number or type of the systems protected and the honeypots that reside in the honeynet.

In our experimental network (Figure 1), the IDS (intrusion detection system) acts as the main network gateway. It monitors and swiftly redirects any malicious network traffic from its DMZ (demilitarised zone) interface to its honeynet interface. The network isolation enables the network administrator and the security analysts to safely access the management interface of the IDS and the firewall and view all the alerts created by both the IDS and the honeypots. Finally, the firewall prohibits any malicious traffic originating from the compromised honeypots targeting the production or the management network. Further rules can be applied on the honeynet firewall, and other security mechanisms can be installed to restrict the honeypots' access to the Internet.

In order to enable Snort to redirect an attacker to the honeynet we modified its source code. Our aim was to introduce as little overhead as possible to the original implementation and to redirect each attacker instantly when a detection rule is triggered. One of the most popular techniques to extend Snort's capabilities is based on its alerts. However, there have been multiple reports of undesirable delays on Snort's alerting engine introduced by this technique. Driven by our aim to create a lightweight active defence mechanism, we decided to perform our modifications directly on Snort's main detection engine and to achieve the redirection of the attackers based on their IP address, using policy routing.

Our approach does not restrict the size of the honeynet or the type of the honeypots that will be used. However it does require the honeynet to have the same address space as the network being protected. The honeypots used should be as close to the original systems as possible but if desired, the honeynet can also contain more systems than the actual network does. This approach has been widely used for intrusion detection, as each connection attempt to an unused (unpublished) IP address should be considered suspicious and redirected to the honeynet.

The aim of the developed honeynet is not only to deceive potential attackers but to collect as much information as possible about their tools and systems as well as their skills and motives. Thus, in every honeypot located in our honeynet, we have installed a local instance of Snort running as a host-based IDS, and p0f, a powerful passive fingerprinting tool, so that we can detect known attacks and collect as much information as possible about the attacker's system.

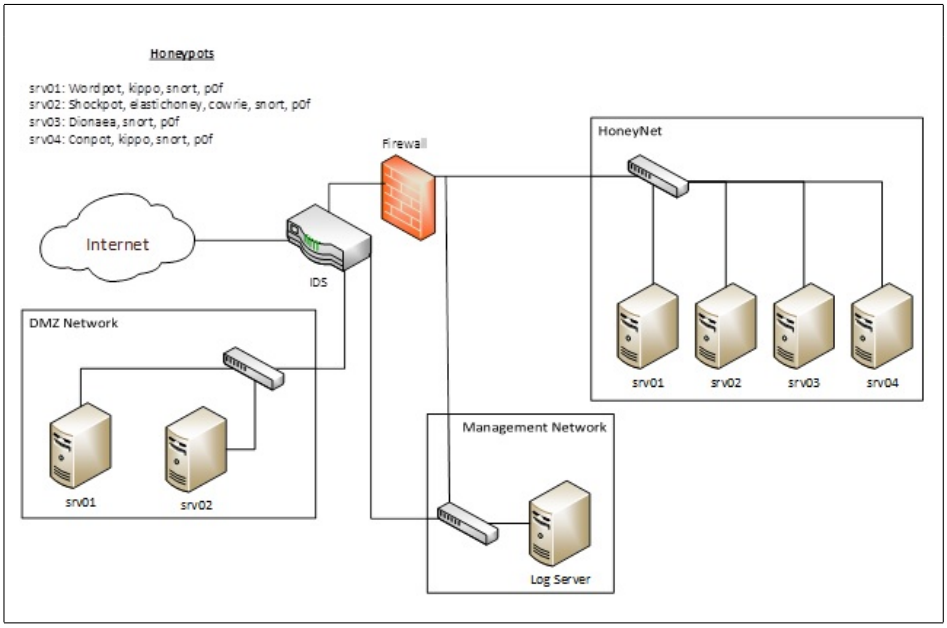

Figure 1: *Experimental network.*

Evaluation

We compared our system with a default Snort installation operating as an intrusion prevention system. Both systems use the same signature database but they react differently when a high severity rule ("DROP" or "SDROP") is triggered. Our system adds a permanent redirection rule to a honeynet each time one of these rules gets triggered whilst a conventional IPS will only drop the malicious packet without taking any further action.

We performed a series of attacks and monitor the results from the viewpoint of both the attacker and the defender. We separated our attacks into two phases, the initial service and vulnerability discovery phase and the final exploitation phase, during which the attacker attempts to exploit any vulnerability discovered during the first phase. We consider that our mechanism to be successful if:

- The attacker:
  - Fails to compromise the original systems.
  - Fails to identify any vulnerabilities on the original systems.
  - Fails to detect the existence of an intrusion prevention system.
- The defender:
  - Has gained some time to analyse the malicious traffic and respond.
  - Can easily distinguish between malicious and legitimate looking traffic.

Discovery phase

The graph in Figure 2 shows the number of hosts, their vulnerabilities and the accessible services/ports that were reported by several popular network and vulnerability scanners during the discovery phase. The grey bars indicate the actual number of systems, ports and vulnerabilities in the DMZ, the blue bars the number seen through application of the IPS and the yellow bars indicate the systems, ports and services seen in the honeynet for redirected traffic. The conventional IPS is successful in hiding a number of vulnerabilities. However, no matter how strictly it was tuned the standard attacker tools were still able to discover some of them. In addition, the conventional IPS failed to hide most of the services/ports that were available on the protected DMZ network, providing the attacker with a useful set of information.

The scans using the modified Snort and honeynet remained realistic, revealing the four hosts in the honeynet (as shown in the network diagram in Figure 1), 15 ports and 25 vulnerabilities. However, contrary to the default IPS, the services and vulnerabilities reported were not those corresponding to the protected systems but to those emulated by the systems located in the honeynet. The attacker has been detected early during the scanning, and was immediately redirected to the honeynet without being able to identify any actual hosts, ports or vulnerabilities. Of course an attacker that changes IP address of their attack machines would not be redirected until Snort detected an intrusion and added that address to the policy routing, likewise a denial of service could be envisaged where an attacker spoofs scans from innocent hosts forcing them to be redirect to the honeynet, so this is by no means a fool proof solution.

From the defender's viewpoint, in both cases it was easy to detect that an attacker is scanning the protected network. However, in the tests the number of alerts produced by the default IPS did not allowed the defender to isolate events and investigate them further as the monitoring console became unresponsive due to the number of alerts. On the other hand, the deployment of redirects and honeynet mechanism produced a minimal number of alerts as the attacker was detected and all the alerts produced by the honeynet were reported to the honeynet system.

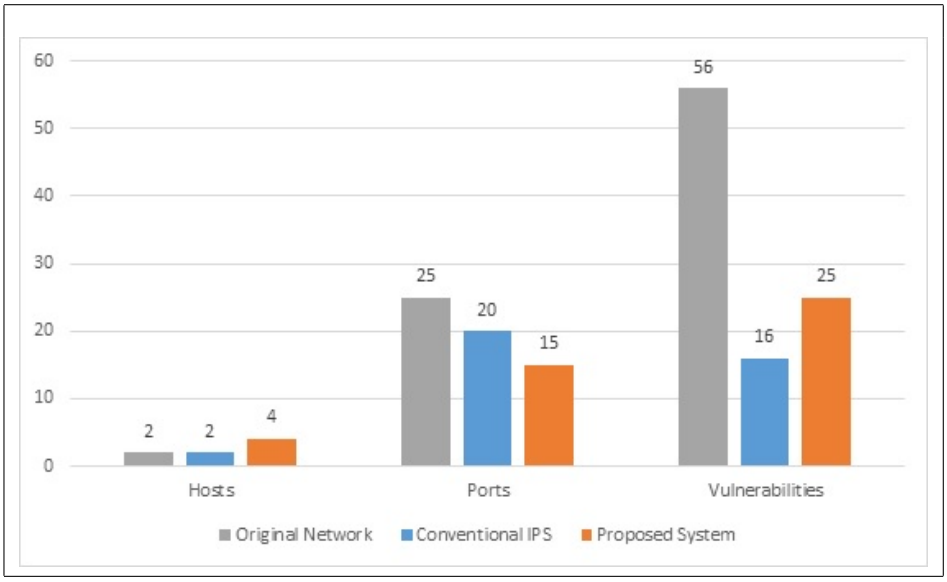

Figure 2: *The number of hosts, ports, and vulnerabilities reported by network and vulnerability scanners during the discovery phase.*

Exploitation phase

In the exploitation phase, we attempted to exploit the vulnerabilities discovered during the scanning phase using publicly available tools and known exploits. While these were known standard attacks and so should have been blocked by the default IPS, the standard attacker tools still managed to gain access to the protected systems in the DMZ using three different known exploits.

Using the modified Snort and the honeynet the attacker was not able to exploit the DMZ. The honeypots used in our proof of concept implementation were medium and low interaction honeypots and the attacker failed to exploit any of the reported vulnerabilities. Thus, the protected systems remained intact during the attack and the alerts produced by the honeypot agents (circa 40,000 alerts) contained significant information based on which the defender would be able to identify the tools and the exploits the attacker had used during the attempt to compromise the network.

Conclusion

Through this proof of concept project we have identified some of the benefits of deploying a redirect and honeynet. During the evaluation we experienced a failure of one of the most popular intrusion detection and prevention systems against a noisy and not very skilful attacker using standard open tools. While the IPS managed to identify and block some of the attacker's intrusive actions, it failed to protect the vulnerable system located in the DMZ network. It can be argued that a better intrusion signature database or configuration of the IPS should have prevented the attacker's access to the vulnerable machines in the DMZ. However, this solution is limited by the security analyst's knowledge and skills and would still fail again against previously unknown attacks.

Moreover, blocking a malicious packet or an attacker alerts the attackers that a security mechanism is blocking their probes. This information can be exploited by the attackers. They could use it to perform a denial of service attack by spoofing IP addresses, or to confuse the security analysts and hide their intrusive actions by producing a large number of false positive alerts. In addition, attackers may try and evade the IPS. On the other hand, the redirect system proposed in this dissertation has shown its effectiveness when using the same intrusion detection signature database .

Our mechanism relies and acts on the detection of the intrusion during its early phase preventing fur-

ther escalation of the attack on the production systems. Unlike blocking mechanisms, the proposed system does not reveal itself to the attackers immediately. This buys time and forces the attacker to attempt other attacks to gain access. Furthermore, our prevention mechanism produced a significantly lower number of alerts that did not reduced our systems performance and could be more easily managed by the security analysts. On the other hand, the honeypots used produced a large number of alerts that contained more information about the attacker's tools and techniques.

Unfortunately, our mechanism relies on the correct classification of an attack by an intrusion detection system, and as a result, it remains vulnerable to bad configuration that can result in false positive events. A false positive alert on our mechanism will result in denying legitimate users access to the systems and redirecting them to the honeynet. However, in that case it will still provide the security administrator with enough information to fine-tune the intrusion detection system.

Biographies

Apostolis Machas holds an M.Sc degree in Information Security from Royal Holloway, University of London and a B.Sc degree in Computer Science from Athens University of Economics and Businesses. He started his career as a Blue team member of a SOC team and has worked in both commercial and military organisations during his career. Currently, Apostolis is working as a Security Advisory Consultant in a world-leading network and security organisation. His interests are but not limited to IoT security, Web Application and Network attacks while during his previous research he focused on access control mechanisms and user's privacy on Cloud systems.

Peter Komisarczuk is a member of the Information Security Group at Royal Holloway, University of London, where he is Programme Director for the MSc Information Security (Distance Learning). He is a chartered engineer and has a PhD (Surrey), researches in networks and security and has worked in industry in various R&D roles at Ericsson, Fujitsu and Nortel Networks.