the network security company™

Controlling Botnets with the Next-Generation Firewall

Introduction

The rise of botnets and modern malware is reshaping the threat landscape and forcing enterprises to reassess how they protect themselves. These modern threats have outpaced traditional anti-malware strategies and in the process, have established a foothold within the enterprise that criminals and nation-states can use to steal information and attack sensitive assets. This paper is broken into two sections: The first provides an overview of botnets and how they work, and a review of the industry responses to date. The second half introduces a model that enables enterprises to take an active role in protecting themselves from botnets based on the unique capabilities of the next-generation firewall.

What are botnets and what makes them so different?

Enterprise information security teams have been doing battle with various types of malware for more than two decades. However, all of this hard-earned experience does not mean that enterprises are necessarily winning the war. Recent studies have found reason for concern. A 2009 Cyveillance study found that on average, half of the active malware in the wild was not detected by industry anti-malware solutions, and a 2010 NSS Labs study found that only 53% of malware was detected on download. This is due in large part to the evolution of malware that either mutates or can be updated to avoid detection by traditional malware signatures. Bots (individual infected machines) and botnets (the broader network of bots working together) play a major role in this evolution of polymorphic malware, and are notoriously difficult for traditional antivirus/anti-malware solutions to detect. As such it's important to understand what it is that makes botnets so different from previous generations of malware, and how it impacts our efforts to control them.

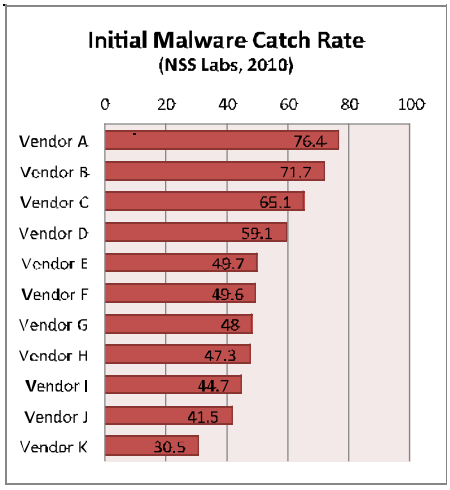

A potentially obvious but important distinction is evident in the name "botnet" itself. Literally, a botnet is a network of bots (infected machines). Unlike earlier types of malware, which were more or less a swarm of independent agents, bots are centrally coordinated and retain a communication channel to the outside world. In much the same way that the Internet changed what was possible in personal computing, botnets are changing what is possible in the world of malware. Now all malware of the same type can work together toward a common goal, with each infected machine growing the power and value of the overall botnet. The botnet can evolve to pursue new goals or adapt to changes in security measures. Some of the most important and unique functional traits of botnets are discussed in the following section:

- **Distributed and fault-tolerant** – Botnets are malware that take full advantage of the resiliency built in to the Internet itself. A botnet can have multiple control servers distributed all over the world, with multiple fallback options. Bots can also potentially leverage other infected bots as communication channels, providing with a near infinite number of communication paths to adapt to changing access options.

- **Multifunctional** – Updates from the command and control servers can also completely change the bots functionality. This enables a new economic approach for a botnet owner, who can now use portions of the botnet for a particular task such as collecting credit card numbers, while other segments of the botnet could be sending spam. The important point is that the infection is the most important step, because the functionality can always be changed later as needed.

- **Persistence and intelligence** – Given that bots are both hard to detect, and can easily change function, they are particularly well-suited for targeted and long-term intrusions into a network. Since bots are under the control of a remote human intelligence, a botnet is more like having a malicious hacker inside your network as opposed to a malicious executable. For instance, a bot can be used to learn more about the organization of the network, find targets to exploit and install additional backdoors into the network in case the bot is ever discovered.

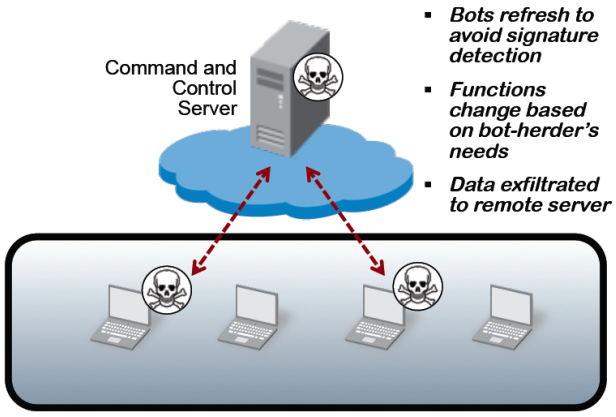

Threats to the Enterprise

Given their flexibility and ability to evade defenses, botnets present an enormous amount of risk to the enterprise. Botnets are virtually unlimited in terms of their functionality ranging from sending spam to the theft of classified information and trade secrets. The ultimate impact of a botnet is largely left up to the bot-herder and a node that was sending spam on one today could be stealing credit card information the next. Here we will summarize some of the most noteworthy classes of botnets and cover some of the more notorious examples of botnets.

Spamming Botnets

The largest botnets are often dedicated to sending spam. The premise is fairly straight-forward – the bot-herder (the person or persons who remotely control the botnet) attempts to infect as many endpoints as possible, which can then used without the owner's knowledge to send out hundreds of spam messages. The relative impact of this type of bot on the enterprise may seem low initially. An infected user sending spam could consume additional enterprise bandwidth and ultimately reduce the productivity of the user and even the network itself. The company could also easily become listed as an organization that sends spam, which could eventually lead to approved corporate emails being labeled as spam.

However, the risks of a bot-infected laptop can reach beyond the functionality of the botnet itself. An infected laptop can provide backdoors and entry-points into the enterprise network to spread and find additional targets to exploit. As such we should never become complacent about any botnet – put simply, there is no such thing as a benignly compromised host.

Example: Rustock

DDoS and Botnets

A slight twist on the spamming botnet model, is to use bots as part of a distributed denial-of-service attack (DDoS) which attempts to overwhelm a target system with traffic from a large number of endpoints. In these cases, the enterprise with the infected client is often not the target of the attack itself, instead simply using the compromised host to flood a remote target with traffic. Again, the bot-herder (the person in control of the botnet) attempts to leverage the massive scale of a botnet, and generate an amount of traffic that could overwhelm resources at the target. These DDoS attacks often target specific companies either for personal/political reasons or to extort payment from the target in return for halting the DDoS attack.

DDoS botnets represent a dual risk for the enterprise. The enterprise itself can potentially be the target of a DDoS attack resulting in downtime and lost productivity. Even if the enterprise is not the ultimate target, any infected end-users participating in the attack could consume valuable network resources, while inadvertently facilitating a crime.

Example: Skunkx

Financial Botnets

Financial botnets have had widespread coverage in the press, owing largely to the spectacular amounts of damage they have caused in the market. Banking botnets such as ZeuS and SpyEye are responsible for the direct theft of funds from all types of enterprises. These botnets are typically not as large and monolithic as the spamming botnets, which grow as large as possible for a single owner. Instead, banking botnets are often sold as kits allowing large numbers of attackers to license the code and set about building their own botnets and targets. Even with their smaller size, the impact of these botnets can be enormous. ZeuS botnets have repeatedly been able to steal millions of dollars from enterprises in very short periods of time. Other financial botnets focus on the theft of credit card information, or faking ACH bank transfers.

The impact of a financial breach can be enormous for an enterprise. The breach of customer credit card information can lead to serious financial, legal and brand damage for the enterprise. Additionally, if HR, finance or accounting groups are compromised, the enterprise could instantly lose money that could potentially never be recovered.

Examples: Zeus, SpyEye

Targeted Intrusions

Botnets are also a key component of targeted, sophisticated and ongoing attacks. These types of botnets are very different than their larger brothers. Instead of attempting to infect large numbers of users to launch actions of large scale, these smaller botnets aim to compromise specific high-value machines that can be used to further the penetration and surveillance into the target network. In these cases, an infected machine can be used to gain access to protected systems, and to establish a backdoor into the network in case any part of the intrusion is found out.

These threats are almost always unknown to antivirus vendors, and they represent one of the most dangerous threats to the enterprise in that they are specifically targeting the most high-value information in the enterprise such as research and development, source code, planning data, financial data, customer lists and the like. These threats represent some of the most dangerous botnets for an enterprise as they are premeditated and target the enterprise's most essential assets.

Example: Aurora

What is the Industry Doing About Botnets?

The explosion of botnets has certainly not gone unnoticed in the industry, and large companies have begun to team up with law enforcement to take action against some of the largest and most notorious botnets. In short, this goal is to separate the bots (the infected machines) from their brain (the command and control servers). If the bots can't get to their servers, they can't get new instructions, or upload data or any of the things that make botnets so unique and dangerous. The general approach can be thought of, and is often referred to as botnet "decapitation".

While this strategy may seem straight-forward, in reality it requires an enormous amount of investigation, expertise and coordination between both security researchers and law enforcement. Disabling a command and control server often requires both physically seizing the server, as well as taking ownership of the domain and/or IP range associated with the servers. This demands a very close coordination between technical teams, legal teams and law enforcement for any hope of disabling the botnet.

Making matters worse, a botnet typically does not rely on a single server, but rather has multiple command and control servers for redundancy purposes. Additionally each server is typically insulated by a variety of intermediaries to cloak the true location of the server. These intermediaries include peer-to-peer networks, blogs and social networking sites and even communications proxied through other infected bots. This means that simply finding the command and control servers is a considerable challenge, requiring time and deep technical investigations.

Making matters worse, most botnets are designed to withstand the loss of a command and control server, meaning that researchers must disable ALL of the command and control options at once. If any of the command and control servers are accessible, or any of the fallback options are successful, the bots will be able to get updates, and can quickly populate a completely new set of servers, and the botnet will quickly recover. It also means that the servers will need to be taken out almost simultaneously. As soon as a botnet owner sees he's under attack, he will immediately attempt to move to a new infrastructure. Thus even a single server remaining functional for even a small amount of time can give the bot owner the window he needs to update the bots and recover the entire botnet.

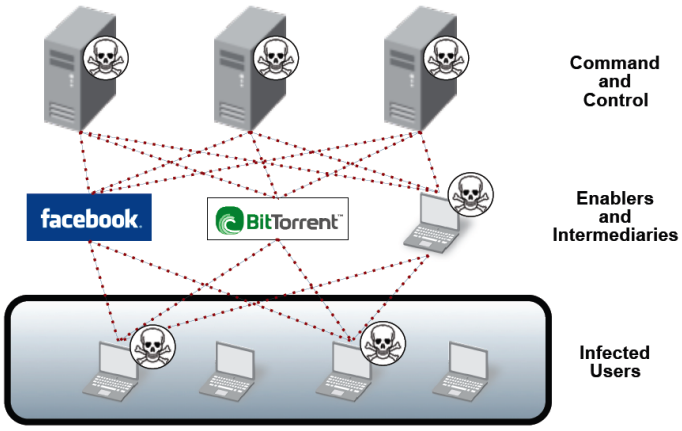

Intensive investigation required to map the distributed command and control infrastructure.

For all the challenges of this approach there have nevertheless been some very consistent and high-profile wins in the industry. In March of 2011, Microsoft working in concert with industry leader and the FBI were able to successfully take down the Rustock botnet, which had operated for more than 5 years and at the time was responsible for sending up to 60% of the world's spam.

Why the Top-Down Model Doesn't Protect the Enterprise

The model described above offers a credible industry response to some of the largest and most notorious botnets, and the companies leading these actions should rightly be lauded for their actions. However, it is equally important as security professionals to understand that these efforts, while good on a broad Internet-wide perspective, do very little to mitigate the threat that botnets pose to an individual enterprise.

The most obvious limitation is that the top-down model is incredibly intensive both in time and effort, and as such only the largest most notorious botnets are targeted. This means the process is particularly reactive, sometimes taking years to complete. Enterprise security needs are obviously far more immediate, and need to ensure that an intrusion or exploit does not succeed in the first place. In a very real sense relying on the industry to disable a botnet is akin to waiting for government to enact a law, when someone is breaking into your business right now.

Secondly, industry efforts tend to focus on the very largest botnets, and typically those that are the notorious spammers. While any bot on a network can introduce risk, its worth noting that these spamming botnets are not targeting the enterprise in particular - they simply are looking for more hosts so that they can send more spam. When an attacker wants to infiltrate an enterprise network or steal information, the malware is almost always far more customized, smaller and more difficult to detect. These smaller botnets are likely to be completely unknown to the industry at all, much less garner enough attention to be targeted by the industry for removal. Once again, the above industry response does very little to mitigate the risk posed to an enterprise.

How the Industry Targets Botnets	How Botnets Target the Enterprise
▪ Focus on the largets botnets	▪ Small, targeted, customized
▪ Focus on spammers	▪ Data theft and espionage
▪ Years required for remediation	▪ Requires immediate remediation

Some botnets will simply be impractical to attack at the command and control level based on how key botnet components are distributed around the world. Recall that one of the major challenges for researchers is the requirement to find and take control of all the command and control servers in a short window of time. The more distributed the botnet is, the more difficult it will be to decapitate. Even in the case of Rustock, authorities were somewhat fortunate in that almost all of the command and control servers were located within the United States, allowing federal law enforcement and court rulings to closely coordinate the all-important process of disabling all of the servers at once. Many botnets have servers all over the world, and will specifically function in areas that have very little enforcement for Internet crimes. This model directly mirrors the Internet itself, which from the beginning was designed to withstand the loss of any one site. In short, a distributed network is designed to withstand a decapitation attempt, so it will be very difficult to extend this model to botnets in general.

So while progress has been made fighting botnets at a global level, the simple truth is that the wins are more of an exception to the rule, and will do little to protect enterprises from the threats posed by botnets. This puts the responsibility for protecting the enterprise from botnets squarely on the shoulders of the enterprise itself. In the next section we will cover some of the new tools and techniques that can enterprises can use today to find these modern types of malware on an enterprise while also preventing the opportunities for infection and exploits.

Proposing a Model for Protecting the Enterprise from Botnet

In this section we introduce methods and best practices to control botnets and related modern malware. These recommendations are designed not to replace but to supplement the existing security strategies of an enterprise as part of a modern coordinated approach to defense in depth. Efforts have been made to discuss these practices as generically as possible. However, we have also included additional details specific to Palo Alto Networks solutions for the benefit of customers who looking to implement botnet protections today. The methodology addresses both techniques to limit the exposure to botnets as well as the detection and remediation of devices that may be already infected. The methodology also addresses the challenge that the most dangerous botnet to an enterprise may very well be unknown to the security community, and introduces techniques to identify these unknown threats in the enterprise network.

- The Role of the NGFW in Controlling Modern Malware
- Reduce the Attack Surface of the Network
- Investigate Unknown Traffic
- Control Malware-Enabling Applications
- Prevent the Use of Circumventors
- Protect Remote Users
- Finding Infected Hosts

Incorporating the Next-Generation Firewall

As malware evolves from individual end-points to coordinated networks, enterprises will similarly need to expand their analytical perspective to incorporate network level intelligence and controls. Network security allows us to focus on the very trait that distinguishes botnets from earlier forms of malware – its reliance on communication with a larger bot network. To twist John Gage's famous phrase, "the network is the computer", in a very real sense the threat has become a network. If our security measures don't operate at this same level, we run a very real risk of missing the forest for the trees.

Additionally, network security mechanisms provide an independent layer of monitoring and control, unlike the end-points themselves, which can be compromised by malware. Botnets and modern malware can include rootkits, whose purpose is to gain root access on the target machine in an effort to subvert antivirus protections or other security mechanisms on the machine. This creates a catch-22 for the security team since any security software running on a compromised host cannot truly be trusted. This certainly shouldn't imply that host-based security is obsolete, but rather simply to point out that the host layer certainly needs additional layers of defense in depth.

The points above could apply to network security in general. However, the next-generation firewall in particular provides arguably the most important addition to the fight against botnets – the reliable visibility and control of all traffic on the network regardless of evasive tactic. By understanding the full stack behavior of all traffic on the network, we can finely control the behaviors that are allowed in the corporate environment, while simultaneously eliminating the shadows that botnets rely on to remain hidden. Bots quite simply must talk in order to function, and finding these telltale communications is becoming a critical component of controlling botnets and the threats they pose to the enterprise. A next-generation firewall performs a true classification of traffic based not simply on signatures, but an ongoing process of application analysis, decryption, decoding and heuristics to progressively peel back the layers of a traffic stream to determine its true identity. This ability to pinpoint and analyze even unknown traffic without regard to port or encryption is defining characteristic of a true next-generation firewall, and as we shall see this ability is invaluable in the fight against botnets.

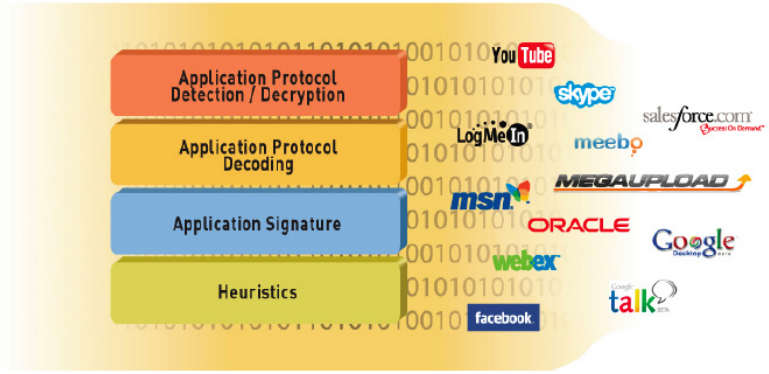

Additionally, a true next-generation provides a fully integrated approach to threat prevention. The distinction in this case is the true coordination of multiple security disciplines as opposed to simply co-locating them on the same box. For example the application identity, malware detection, intrusion prevention, URL filtering, file type controls and inspection of the content should all be integrated into a unified context. This integration provides a far more intelligent and definitive understanding of botnets than any individual technology can provide by themselves. This collective intelligence is needed in order to see and understand the telltale signs of unknown threats.

Preventing Botnet Infection

One of the most important steps that an enterprise can take to control modern malware is to reduce the vectors of infection and eliminate the ability for the bots to hide. Today the majority of the infection vectors used by botnets are virtually unchecked, and botnet traffic is typically small enough that it can easily blend into the background of "normal" network traffic. By regaining full visibility and control of exactly what traffic is allowed into the network and why, security teams can go a long way to satisfying both of these goals.

Reduce the Attack Surface

An important first step for the enterprise is to return to a positive control model. Positive control simply means specifically allowing the traffic you want as opposed to blocking every individual thing that you don't. The notion of positive control has long been one of the defining characteristics of network firewalls that separates them from other types of network security. For example, if you want to use Telnet, then you open port 23 to allow Telnet without necessarily allowing all other types of traffic. Unfortunately, traditional firewalls have progressively lost the ability to enforce positive control in any reliable way, as applications have learned to use non-standard ports such as using commonly open ports (port 80, 443, 53) or simply hop between any available ports.

Enforcing positive control is an essential tool in the fight against malware as it provides an easy way to greatly reduce the attack surface of the enterprise and reduce overall risk. The simple truth is that the number and diversity of applications has exploded, and almost all of them can introduce some level of risk. The rise of web 2.0, widgets and a variety of scripting options have empowered individuals to develop powerful applications or services, most of which are designed to connect or be used with other applications and sites. Making matters worse, very few of the applications that are written on a daily basis have any real value to an enterprise. By incorporating a positive control model, security teams can focus on enabling the approved applications, as opposed to constantly trying to stay up to speed with all of the applications that they want to block. This approach can immediately preclude large numbers of applications from ever touching the network, while dramatically reducing the number of vectors that botnets can use to get in or out of the network.

Tips

What is a True Next- Generation Firewall?

Today, Palo Alto Networks' next-generation firewall is the only firewall to truly enable the positive control of modern application traffic. This required the firewall to evolve to a more sophisticated and reliable method of understanding traffic than simply opening or closing a port. Unlike other firewall vendors, Palo Alto Networks actually started from scratch and developed a new classification intelligence for the firewall to classify all traffic at the application level instead of relying only on port and protocol.

However even with a next-generation firewall, enforcing positive control is not as easy as simply flipping a switch. Some applications may be used by staff and have values that are not readily apparent. As such, IT and security teams should plan to consult with a variety of groups within the organization in order to establish an appropriate set of approved applications and use roles.

Additionally, some applications will have both business and personal uses such as Facebook, which can be used by employees for personal correspondence as well as by the company to stay in touch with customers and prospects. In these cases, policy should be further refined to allow only certain users to access an application or limit the use of an application to certain approved features. This notion of securely enabling applications will be covered in later sections of this document.

Summary – Reduce the Attack Surface

- Establish policies of approved applications and uses based on company needs and culture
  - Establish a baseline of what is on the network – applications, protocols, etc
  - What applications are in use?
  - What applications are required for the business and who needs to use them?
  - What dual-use or personal applications does the enterprise want to allow
- Enforce positive control of all traffic
  - Prevent the unnecessary or high risk traffic
  - Regardless of port evasion or encryption

Investigate Unknowns

Once the enterprise has regained the ability to accurately classify the approved traffic on the network, we then have a base from which to investigate any remaining unknown traffic in the network. The presence and behavior of unknown traffic is a critical clue in the identification of botnets, which will often present as "unknown" traffic due to their use of proprietary encryption and unique behavior.

(receive_time in last-hour) AND (app eq unknown-udp)

Receive Time	Application	Name
04/29 10:32:08	unknown-udp	Bot: Mariposa Command and Control
04/29 10:17:24	unknown-udp	Bot: Mariposa Command and Control
04/29 10:17:24	unknown-udp	Bot: Mariposa Command and Control
04/29 10:16:50	unknown-udp	Bot: Mariposa Command and Control
04/29 10:02:09	unknown-udp	Bot: Mariposa Command and Control
04/29 09:47:23	unknown-udp	Bot: Mariposa Command and Control
04/29 09:47:23	unknown-udp	Bot: Mariposa Command and Control
04/29 09:46:52	unknown-udp	Bot: Mariposa Command and Control

Investigate any unknown traffic

The next-generation firewall should provide this ability to find and analyze this unknown traffic in the network. Unknown traffic regularly sent by the same client machine should be investigated to determine if the individual is using a legitimate application that is not recognized versus a potential botnet infection. Teams can also investigate where the traffic is going. Does it reach out to websites known to serve malware or to social networking sites? Does it transmit on a noticeable schedule? Does someone attempt to download or upload files to an unknown URL? All of these behaviors can identify the presence of client machine that is infected by a bot. By using the next-generation firewall to definitely identify the approved traffic on the network, the "unknown" traffic should become increasingly rare, and thus enabling potentially malicious botnet traffic to be found and analyzed quickly.

This methodology has already proven effective in identifying previously unknown botnets in an enterprise. In fact, one of the first discoveries of the Mariposa botnet in the wild was found by Palo Alto Networks customer where the customer performed just this methodology of progressively investigating and classifying unknown traffic in the network. The analysis of unknown traffic can be automated by using the Behavioral Botnet Report, which is discussed in detail in following sections.

Tips

Tracking the Unknown

Palo Alto Networks firewall makes it easy to investigate and control unknown traffic. Unknown traffic can be tracked, analyzed and controlled just like any other application including where and how much traffic is being sent, the URL categories involved, threat or malware signatures that are triggered and if any file transfers were attempted. Users can also capture a PCAP of the unknown traffic for detailed analysis or to be passed on to the Palo Alto Networks Threat Research Team. The Behavioral Botnet also automatically correlates common botnet behaviors in order to identify machines that are likely infected by a bot. This capability is covered in detail in later sections.

Palo Alto Networks also enables the enterprise to properly classify any custom applications to ensure that approved applications do not present as unknown. The Palo Alto Networks firewall allows users to write their own App-IDs for these custom applications. This allows teams to eliminate unknowns over time with the majority of traffic being well understood and controlled.

Summary – Investigate the Unknowns

- Investigate "unknown" traffic
  - Track source and destination, volumes of traffic
  - Correlate against URL, IPS, malware and file transfer records
  - Define Custom App-IDs as needed for any internal or custom applications
  - Capture PCAPs for any unrecognized, publicly available applications and deliver to Palo Alto Networks
  - Investigate "unknown" traffic for potential unauthorized user behavior or potential botnet behavior

Control the Enabling Applications

Applications are an indispensible part of the botnet lifecycle, and are key to both the initial infection stage as well as the ongoing command and control of the botnet. This association of malware and applications is nothing new. In the past, the de facto enabling application for malware was corporate email. Viruses and email simply went hand in hand from a security perspective. However, even though email is still used by attackers, it has gradually lost its luster for attackers as these applications are typically heavily secured by the enterprise, and email messages can be analyzed at considerable depth while the message is at rest on the email server. However bot-herders have now shifted much of their attention to softer target applications that interact with users in real-time and offer far more flexibility than corporate email.

All applications are not created equal, and botnet owners have gravitated to applications that support the bot's ultimate aim to facilitate social engineering while remaining hidden. Social networking and personal use applications meet both of these criteria, and have become some of the most common sources for malware infection and the subsequent command and control. This includes applications such as social networking apps themselves, web-based email, instant messaging applications, peer-to-peer networks and a variety of file transfer applications. These applications are fundamentally designed to easily share information in a variety of ways, and users often bring a more cavalier attitude to these applications because they may be accustomed to using them outside of the office. This provides an intelligent attacker with a multitude of infection options to pursue and develop.

Social applications also present an ideal environment for social engineering, enabling an attacker to impersonate a friend or colleague in order to lure an unsuspecting victim into clicking on a dangerous link. For all of its sophistication, malware infection continues to rely on enticing an unsuspecting user into an ill-advised click. Instead of opening an email attachment, the click may be a link in a tweet or a link on a Facebook page that appears to be from a friend. Cross-site scripting can populate dangerous links among friends and sniffing technologies such as FireSheep allow hackers to take over social networking accounts.

This connection between social networking and malware has been observed in the real world. Research from the Information Warfare Monitor and Shadowserver Foundation has provided compelling evidence for the role of social networking applications in the botnet lifecycle. In the 2010 paper Shadows in the Cloud, the group tracked a very targeted and persistent intrusion into a network using a customized botnet. In their analysis, the group found that the bot-infected machines rarely, if ever, communicated directly with the command and control servers. Instead the initial malware traffic from the infected host would go to popular blogs, Google Groups, Twitter accounts, and Yahoo! Mail accounts, which allowed the malware communications to blend in with "normal" traffic. This illustrates the key lesson that botnets will often attempt to blend in with what is considered normal but low value traffic in the network. How often would a security admin investigate what appears to be a user simply posting something to an innocuous blog?

The Move to SSL by Default

The reliance on the Internet and cloud computing is leading a widespread adoption of SSL for a variety of technologies and industries. Social networking sites are also inadvertently making it easier for malware to remain hidden by moving to the default use of SSL to protect user communications. This is a needed improvement given that hackers can eavesdrop and hijack unprotected HTTP sessions. Tools such as FireSheep have made this process simple for anyone and threaten the notion of privacy on the web. On the other hand, most enterprises lack the ability to dynamically look within SSL encrypted communications, thus making the social networking traffic more or less invisible to the enterprise. This represents a net loss for the enterprise security team - the user gets improved privacy for their social traffic, but in the process it establishes an invisible communication infrastructure for the same sites and applications favored by modern malware. This move to SSL by default can realistically make social applications as valuable to attackers as P2P applications such as BitTorrent have been for the past several years. The pattern is quite the same – an encrypted communication channel driven the personal motivations of an end-user.

Tips

Controlling the Enabling Applications

Palo Alto Networks provides the tools and techniques to control and secure the use of these malware-enabling applications. This ability to securely enable any application is a critical requirement for enterprises where simply blocking all access to blogs, webmail, IM and social networking apps would be both impractical and unduly constrain the enterprise's ability to communicate and stay connected with the outside world.

Control Applications Based on User - A first step of securely enabling an application is to limit access to the app by user or user group that have an approved need for the app. For example, access to Facebook and its underlying applications can be limited to sales and marketing teams who are responsible for maintaining the company's online identity, while other employees are not allowed or have tighter restrictions. This again can significantly reduce the attack surface of the enterprise and reduce the risk of an infection.

Limit Applications to Specific Features - Security teams also will have the option to allow certain applications, while preventing the use of specific features that introduce unnecessary risk. For example, the company could allow access to a social networking application but disable the posting functionality or prevent the application from downloading files or other risky behaviors such as tunneling other applications or sharing the user's desktop. This step can significantly limit the ability for a malware payload to be transferred to a specific target.

Prevent Drive-by-Downloads - Often, the target end-user is not even aware that they are downloading files in the first place. An infected webpage can easily cause the user to automatically download a file from the site in the background. This is commonly referred to as a drive-by download, and can occur even on perfectly valid web pages that have been compromised. The Drive-By-Download protection feature in Palo Alto Networks protects against this type of infection by prompting users to verify that they really intended to download a file and that files are not pulled down without the user's knowledge. This functionality is critical for social networking attacks where links may redirect to sites that are serving up exploits which are then used to download malware or droppers in the background.

Selectively Decrypt SSL - Next, the enterprise must establish a strategy for dealing with SSL encrypted traffic. The Palo Alto Networks next-generation firewall offers on-box SSL decryption, which like all Palo Alto Networks features, can be leveraged based on application or application type. This allows staff to specifically target social networking applications for SSL decryption and content inspection. Conversely teams can also use URL filtering categories to avoid decrypting sensitive personal information such as traffic destined for financial or healthcare sites.

Key Recommendations – Control the Enabling Applications

- Prevent use of known "bad" applications
  - P2P
  - Limit application usage to users/groups who have a need
- Prevent the use of dangerous features
  - File transfer
  - Desktop sharing
  - Tunneling of other applications
- Prevent Drive-by-Downloads and train users to use the feature
- Selectively decrypt SSL based on application and URL category
  - Decrypt social networking, webmail, Instant Message
  - Do not decrypt traffic to/from health care or financial sites
- Inspect and enforce all allowed risky application traffic
  - Intrusion and Threat Prevention
  - Malware Protection
  - URL Filtering

Prevent the Use of Circumventors

The previous section focused on the common end-user and web 2.0 applications that can be co-opted by malware for use against the enterprise. However, there is a second class of applications that are proactively designed to pass through traditional network security without interruption. This includes a variety of remote desktop technologies, proxies and purpose-built circumventing applications, which will also require tight control. Some of these applications have valid enterprise uses, while others are a sure sign of unapproved and dangerous behavior. In all cases they will require tight control by IT to prevent the opening of unmanaged threat vectors into the enterprise.

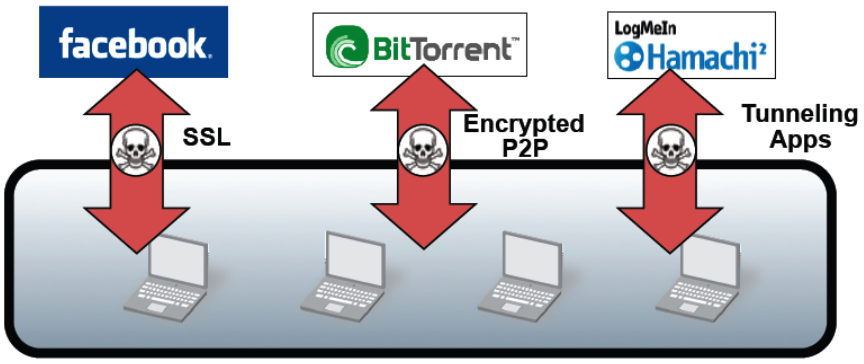

Remote desktop technologies have become wildly popular both with end-users as well as IT and support teams. Even most web-conferencing applications have added the ability to give a remote user control over a user's machine. Such technologies introduce two risks. First, when a user connects to his remote PC, he is free to surf to any destination and use any application without that traffic being inspected by the firewall. In addition to circumventing policy, the remote desktop opens an unmanaged threat vector by allowing a user to remotely undertake all kinds of risky behavior and then have the results tunneled back to his machine inside the enterprise. Secondly, remote desktop technologies provide the risk of an outside user gaining full access to a machine inside the trusted enterprise network. This type of remote control is one of the ultimate goals of malware in the first place, and as such it obviously creates a dangerous opening to launch an intrusion.

We also see common enterprise applications such as SSH which will have valid uses within the enterprise but can easily create unintentional exposures if misused or used by unauthorized or untrained users. For example, most enterprises will use SSH to manage systems and applications in a corporate datacenter. However, if SSH is used to open tunnels into the datacenter, it can open direct, unmanaged access into the enterprise's most critical assets. These applications will need tight control in the enterprise with only select individuals approved and any tunneling features closely regulated.

Additionally, a variety of web-proxies and encrypted tunneling applications have evolved whose primary goal is to provide secure and anonymous communication across firewalls and other security infrastructure. Proxy technologies such as CGIProxy or PHProxy provide a relatively easy way for users to surf securely without enterprise control and have been found in more than 75% of enterprise networks. Applications such as UltraSurf, Hamachi and Tor are purpose-built to traverse security infrastructure and are regularly updated in order to remain undetected. These applications have very few if any valid reasons for use within the enterprise, and their presence generally indicates an intentional attempt to avoid enterprise security. These tools not only pass traffic without being inspected, but they also tend to be used for high-risk behaviors such as file sharing or expressly blocked content and sites, which in turn carry a significantly higher risk of malware infection. As a result, these applications should be blocked in almost all cases.

Tips

Controlling Circumventors

Remote Desktop

Palo Alto Networks allows the enterprise to control a variety of remote desktop technologies automatically and by policy. As with all applications, the remote desktop technologies are grouped together and can be easily controlled by creating an application filter. As new remote desktop applications are released into the market, Palo Alto Networks introduces new App-IDs that map to the appropriate filter categories, meaning that users are automatically protected from new applications. Additionally, remote desktop applications are prime candidates for user-based controls. For example, IT and support teams that rely on remote desktop for their work can be allowed to use these applications, while other employees can be denied.

SSH

Palo Alto Networks allows the enterprise to control SSH based on policy. This of course includes the ability to limit SSH usage to the few IT users that have a real need. Just as importantly, Palo Alto Networks allows teams to selectively disable the tunneling capability of SSH without blocking SSH altogether. This enables the enterprise to use the tools they need, but without the risk of creating open tunnels into the enterprise.

Encrypted Proxies

Encrypted proxies represent some of the most challenging applications in the world in terms of control and blocking, and is also where the quality of application research teams can be seen clearly. Applications such as UltraSurf, Hamachi and Tor were originally designed to avoid censorship controls, but has also been linked with the hacking community for their ability to circumvent traditional security. To make matters worse, these applications are regularly updated in order to avoid detection by security technologies. It is incumbent on the firewall vendor to continually track these applications and update security measures accordingly. To date, Palo Alto Networks is the only next-generation firewall to detect and block these applications at all. Furthermore Palo Alto Networks researchers continually track changes to these applications and update App-IDs as needed to continue to prevent the use of these front-line hacking tools.

Key Recommendations – Control Circumventors

- Limit Remote Desktop Usage
  - IT only
- Securely Enable SSH
  - Allow but prevent SSH tunneling
- Block use of unapproved proxies
- Block Encrypted Tunnels
  - UltraSurf
  - Hamachi
  - Update App-IDs weekly

Protecting Remote Users

Thus far we have implied a fairly traditional network topology with a clear separation between the inside and outside of the network. However, enterprise computing has evolved to reach well beyond the traditional physical boundaries of the enterprise. Users take their laptops home with them and expect to be able to connect and work literally from everywhere. This creates an imbalance in security posture in that users expect to be able to do the same work from any location, while the lion's share of security infrastructure (firewalls, IPS, etc.) are applied only when the user is inside the traditional physical perimeter of the enterprise. To make matters worse, a user's browsing and application behaviors often tend to be riskier when outside the office than when inside, as users unconsciously revert back to their personal behaviors they use at home. This behavior greatly increases the likelihood of clicking on a dangerous link or visiting a site that serves up a drive-by download.

Tips

GlobalProtect

Palo Alto Networks' GlobalProtect focuses on bridging the gap between traditional enterprise security and the comparative lack of security when the user roams beyond the physical boundaries of the enterprise. In short, GlobalProtect allows a remote or traveling user to remain connected to the enterprise firewalls. This ensures that all user traffic receives the same application-based firewalling, IPS, malware prevention, URL filtering and botnet detection capabilities. The GlobalProtect solution is transparent to the user and leverages all of an enterprise's Palo Alto Networks firewalls to ensure strong performance and usability regardless of where the user travels. GlobalProtect also ensures access to all next-generation firewall features such as the drive-by download protection to protect users from infection regardless of whether they are in the office or in the coffee shop. The end-result is consistency both in terms of policy enforcement and threat prevention no matter the user's location.

Key Recommendations – Protect Remote Users

- Enforce full enterprise firewalling and threat prevention regardless of user location
  - Include GlobalProtect to protect remote users
- Enforce Drive-by-Download protections
- Enforce customized policies based on user location
  - Not allowed to download files from secure systems when remote

Finding Infected Hosts

In spite of the security team's best efforts at prevention, enterprise machines will inevitably be infected with malware. This could be via an unknown type of malware, an unknown vector or by physical connections such as a USB drive. Malware has proven time and again that it is possible to infect even the most heavily secured systems in the world. As a result, it is important for teams to assume users are infected and develop the skills needed to find the infected hosts in the network. This can be a challenging task given that the malware may have already avoided traditional malware signatures and may already have root access on the infected machine.

To pinpoint these infected hosts, we again must shift our attention from malware signatures to instead analyze behaviors that are observed in the network. For all of their secrecy and ingenuity, botnets need to communicate in order to function, and they also need to make themselves difficult to find and difficult to trace. These basic requirements create patterns that we can use to identify bot traffic or behaviors that stand out from the normal end-user traffic, even if the bot is completely unknown in the industry.

Detection of Command-and-Control Traffic

One of the major advantages of a next-generation firewall is the ability to classify potentially complex streams of traffic at the application level. This includes the ability to progressively scan within traffic to peel back protocols running within protocols until the true underlying application is identified. This expertise in identifying complex traffic is very valuable when identifying the unique command and control traffic of particular botnets. For all intents and purposes a botnet is an application and its unique traffic can be identified by Palo Alto Networks. The detection of C&C traffic is an integral component of the threat prevention module and is regularly updated along with other threat and content updates.

Leveraging IPS to Detect Botnets

Palo Alto Networks also includes a variety of additional techniques to identify potentially polymorphic malware based on specific components within the malware. For example, SpyEye a very popular and growing banking botnet reserves space so that the malware can constantly change its size, and therefore change its signature. However, SpyEye periodically downloads an encrypted configuration file to update the bot. Researchers at Palo Alto Networks were able to break into this configuration file and were able to find a unique pattern across all configuration files, that enables the IPS module to identify the presence of the bot, even if the malware itself is not recognized. This is just one example where IPS and malware detection can intersect to find a modern threat.

Behavioral Botnet Report

While the investigative techniques described earlier are often crucial, many enterprises simply don't have the time for many manual investigations. The Palo Alto Networks Behavioral Botnet Report automates this process of tracking and correlating the behaviors that indicate the presence of a bot. This intelligence looks for a variety of characteristics, which are briefly summarized below:

- Unknown TCP/UDP – As seen earlier in the paper, botnet traffic is regularly encrypted and unknown. Since Palo Alto Networks identifies all traffic tracking unknown TCP and UDP traffic can be a perfect starting point for finding bot-infected machines. The report allows staff to track unknown traffic by sessions, destinations and bytes.

- Presence of Dynamic DNS – Malware will often use dynamic DNS in order to make botnet communications more difficult to track. By bouncing traffic between multiple infected hosts with an ever-changing list of IP addresses, it can become very difficult to track the path of the bot and its true source and destination.

- Activity on Known Malware Sites – As part of the URL filtering solution, Palo Alto Networks constantly tracks sites that have hosted malware whether intentionally or unintentionally. Palo Alto Networks can track if a user is repeatedly visiting one of these sites and attempting to download files.

- Visiting Recently Registered Domains – Botnets are constantly moving around in order to avoid detection and to recover as servers are discovered or disabled. As a result, botnets will often have to use new domains to support the command and control infrastructure. A user repeatedly visiting a newly registered domain will certainly not be conclusive, but may help to provide corroborating evidence of an infection.

- Browsing to IP domains Instead of URL – In a similar vein, bots will often use hard-coded IP addresses or known IP ranges in order to communicate as opposed to users which typically prefer to use URLs. As with tracking newly registered domains, tracking connections using IP domains can sometimes indicate the presence of a bot at work as opposed to a human.

- IRC traffic – IRC traffic is one of the most well-known communication methods for botnets, and provides an additional strong piece of correlating data for finding a bot.

The Behavioral Botnet Report takes all of the factors above and automatically correlates them to find hosts that are likely infected with a bot. When run, the report provides specific directory user names of the users or machines that are likely infected along with what behaviors contributed to the analysis. Each user is also provided a score based on how many of the factors listed above were correlated, allowing staff to focus on the devices that are the most likely to be infected.

Summary of Best Practices

Establish Positive Control and Investigate the Unknowns

- Establish policies of approved applications and uses based on company needs and culture
- Define Custom App-IDs as needed for any internal or unrecognized approved applications
- Investigate “unknown” traffic for potential botnet behavior

Control the Enabling Applications

- Prevent use of known “bad” applications
  - P2P
- Limit application usage to users/groups who have a need
- Selectively decrypt SSL based on application and URL category
  - Decrypt social networking, webmail, Instant Message
  - Do not decrypt traffic to/from health care or financial sites
- Prevent the use of dangerous features
  - Posting on social networking sites
  - File transfer
  - Desktop sharing
- Prevent Drive-by-Downloads and train users to use the feature
- Inspect and enforce all allowed risky application traffic
  - IPS
  - Malware
  - URL Filtering

Control Circumventors

- Limit Remote Desktop Usage
  - IT only
- Securely Enable SSH
  - Allow but prevent SSH tunneling
- Block use of unapproved proxies
- Block Encrypted Tunnels
  - UltraSurf
  - Hamachi
  - Update App-IDs weekly

Protect Remote Users

- Enforce full enterprise firewalling and threat prevention regardless of user location
  - Include GlobalProtect to protect remote users
- Enforce Drive-by-Download protections
- Enforce customized policies based on user location
  - Not allowed to download files from secure systems when remote

Finding Infected Hosts

- Enforce full enterprise firewalling and threat prevention regardless of user location
  - Include GlobalProtect to protect remote users
- Enforce Drive-by-Download protections
- Enforce customized policies based on user location
  - Not allowed to download files from secure systems when remote

Summary

Botnets have transformed the world of malware and how modern networks are attacked. These threats are experts at remaining hidden from traditional security, while exhibiting an intelligence, resiliency and scale that had never been seen in malware before. Controlling these threats will require multiple security disciplines. While no single solution will solve the problem of botnets on its own, the next-generation firewall provides the unique visibility and control and true integration of threat prevention disciplines needed to find and stop these threats today and into the future. These methods will continue to evolve as new techniques are developed, and malware evolves. Please contact us with any comments or questions about how we can help control botnets in your environment and improve your network security as a whole.

paloalto NETWORKS
the network security company™

3300 Olcott Street
Santa Clara, CA 95054

Main: +1.408.573.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com