WHITE PAPER

Advanced Security Intelligence with Rap Sheet Analytics

Application and Threat Intelligence Powers ThreatArmor

Cybersecurity is a top priority for almost every large enterprise in the world today. With security breaches on the rise and the threat posed to companies large and small, network and security administrators are necessarily on the alert. The good news is the growing number of tools to address these risks. One organization deeply committed to improving the state of security is the Ixia Application and Threat Intelligence (ATI) Research Center.

Ixia's ATI Research Center has been performing advanced security research for well over a decade, providing intelligence updates in use by customers in every industry around the globe. Updated application signatures are leveraged by customers using Ixia's BreakingPoint™ and PerfectStorm™ products when validating performance of deep packet inspection (DPI) solutions and application servers. Customers of Ixia's Visibility products use these same signatures to recognize and pinpoint specific applications. Threat intelligence in the form of live malware, botnet patterns, new attacks, and location of bad actors is used to validate the protection of security appliances. Ixia's ThreatARMOR™ also makes the output of the ATI Research Center available to customers, improving the efficiency of their security infrastructure and operational teams by eliminating known bad threats and untrusted countries. In this paper, we will explore how Ixia collects, validates, and distributes real-time information on global threats through the ATI Research Center.

Threat intelligence is used to validate the protection of security appliances.

KEYSIGHT TECHNOLOGIES

Rap Sheets Describe Network Risks

Whenever ThreatARMOR blocks traffic to or from a known-bad site, a Rap Sheet is provided to explain why that IP address is considered bad. This helps customers better understand the risks facing their network and also avoid the risk of false positive. ThreatARMOR only blocks an IP address if the ATI Research Center has 100% certainty there is malicious or criminal activity at that site, and the Rap Sheet details the proof. The Rap Sheets themselves provide information such as the URL information of individual threats, the binary checksum of that malware, screen shots of a phishing page or malware installer, and the last date the individual piece of malicious activity was validated.

> “Real-time information on global threats.”

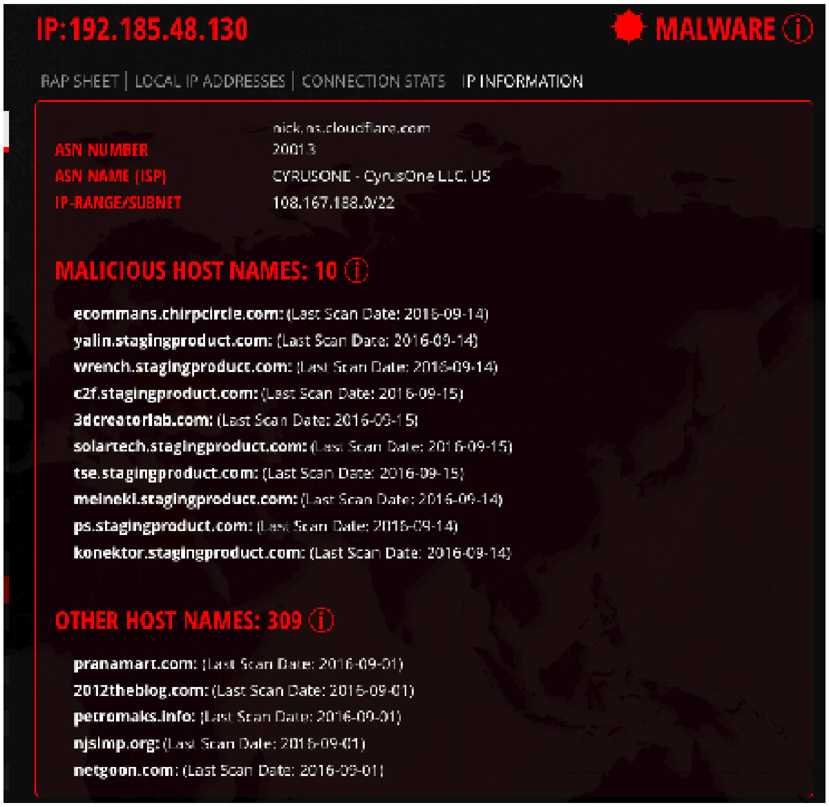
MALWARE
RAP SHEET | LOCAL IP ADDRESSES | CONNECTION STATS IP INFORMATION

IP:192.185.48.130

nick.ns.cloudflare.com

ASN NUMBER 20013

ASN NAME (ISP) CYRUSONE - CyrusOne LLC, US

IP-RANGE/SUBNET 108.167.188.0/22

MALICIOUS HOST NAMES: 10

ecommans.chirpcircle.com: (Last Scan Date: 2016-09-14)
yalin.stagingproduct.com: (Last Scan Date: 2016-09-14)
wrench.stagingproduct.com: (Last Scan Date: 2016-09-14)
c2f.stagingproduct.com: (Last Scan Date: 2016-09-15)
3dcreatorlab.com: (Last Scan Date: 2016-09-15)
solartech.stagingproduct.com: (Last Scan Date: 2016-09-15)
tse.stagingproduct.com: (Last Scan Date: 2016-09-15)
meinekl.stagingproduct.com: (Last Scan Date: 2016-09-14)
ps.stagingproduct.com: (Last Scan Date: 2016-09-14)
konektor.stagingproduct.com: (Last Scan Date: 2016-09-14)

OTHER HOST NAMES: 309

pranamart.com: (Last Scan Date: 2016-09-01)
2012theblog.com: (Last Scan Date: 2016-09-01)
petromaks.info: (Last Scan Date: 2016-09-01)
njsimp.org: (Last Scan Date: 2016-09-01)
netgoon.com: (Last Scan Date: 2016-09-01)

Fig. 1: Rap Sheet example of malicious activity

Additional DNS-related information has recently been added to the Rap Sheets to help users better understand the activity at a particular site, especially multi-hosting sites.

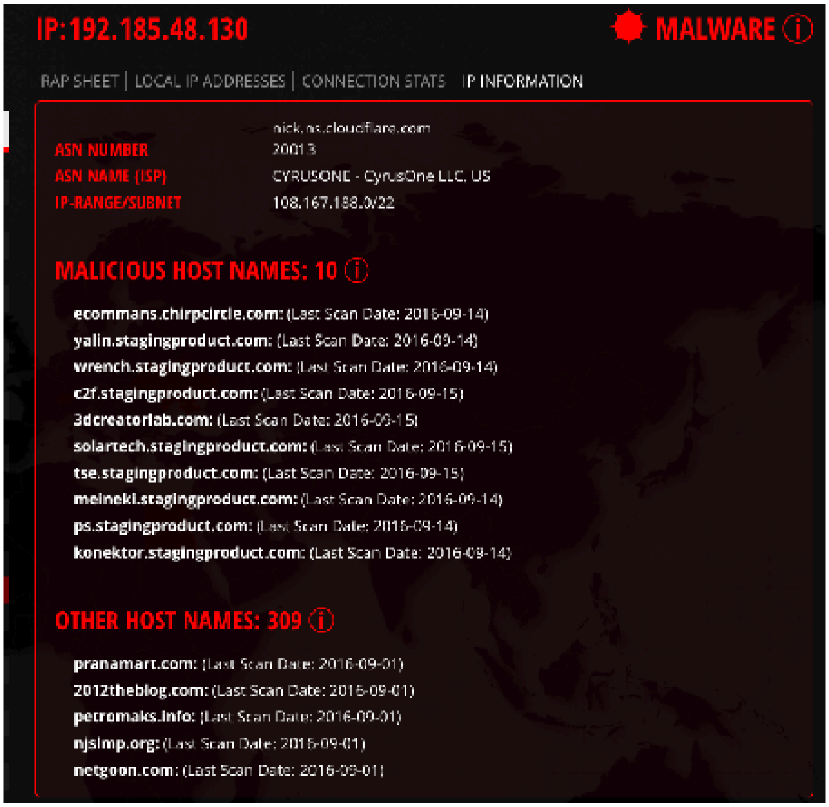
RAP SHEET | LOCAL IP ADDRESSES | CONNECTION STATS IP INFORMATION

IP:192.185.48.130 MALWARE

nick.ns.cloudflare.com

ASN NUMBER 20013

ASN NAME (ISP) CYRUSONE - CyrusOne LLC. US

IP-RANGE/SUBNET 108.167.188.0/22

MALICIOUS HOST NAMES: 10

ecommans.chirpcircle.com: (Last Scan Date: 2016-09-14)
yalin.stagingproduct.com: (Last Scan Date: 2016-09-14)
wrench.stagingproduct.com: (Last Scan Date: 2016-09-14)
c2f.stagingproduct.com: (Last Scan Date: 2016-09-15)
3dcreatorlab.com: (Last Scan Date: 2016-09-15)
solartech.stagingproduct.com: (Last Scan Date: 2016-09-15)
tse.stagingproduct.com: (Last Scan Date: 2016-09-15)
meineki.stagingproduct.com: (Last Scan Date: 2016-09-14)
ps.stagingproduct.com: (Last Scan Date: 2016-09-14)
konektor.stagingproduct.com: (Last Scan Date: 2016-09-14)

OTHER HOST NAMES: 309

pranamart.com: (Last Scan Date: 2016-09-01)
2012theblog.com: (Last Scan Date: 2016-09-01)
petromaks.info: (Last Scan Date: 2016-09-01)
njsimp.org: (Last Scan Date: 2016-09-01)
netgoon.com: (Last Scan Date: 2016-09-01)

Fig 2: Rap Sheet example of multiple malicious domains at an IP address

A lot of work goes on behind the scenes to collect, analyze, validate, and distribute that Rap Sheet information, with raw input data coming from a variety of sources.

Input Data for Known-Bad Sites

The raw data for suspicious sites can enter the system from several different streams, including:

- Source feeds
- Internet scanning
- Honeypots
- SPAM
- Binary Analysis

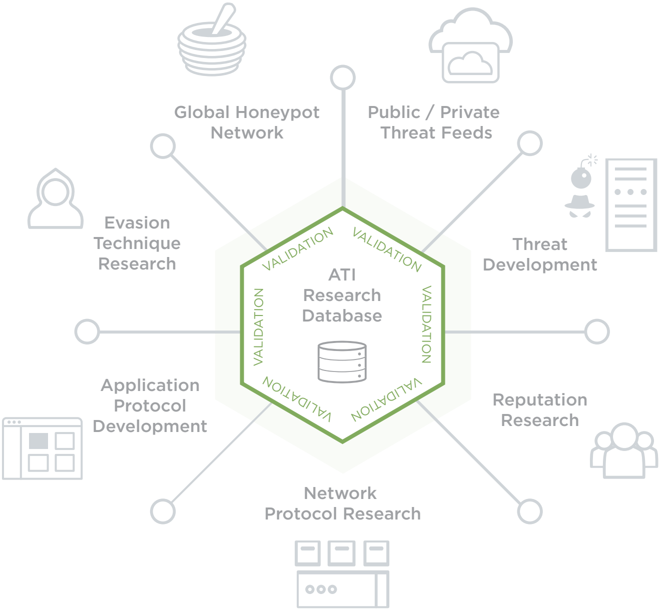
Global Honeypot Network
Public / Private Threat Feeds
Evasion Technique Research
Threat Development
ATI Research Database
VALIDATION
Application Protocol Development
Reputation Research
Network Protocol Research

Fig 3: ATI Research Center

Source Feeds are collected from different public and private streams, including commercial Threat Intelligence feeds which some users and security vendors apply without further validation. The ATI Research Center also collects feed data from the open-source community and various security partnerships. Of course, the data from these feeds are just considered suspects – they aren’t treated as criminal until the ATI Research Center has validated each one individually.

Internet Scanning is used to essentially walk the Internet looking for sites of interest. This is done in an automated fashion, running 24/7, to identify suspicious sites. Sometimes the system immediately finds proof of malicious activity, such as malware binaries. Those sites are automatically scanned, cataloged, and added to the Rap Sheet database. Other sites may be found not to have malicious content but to have known vulnerabilities such as unpatched web server software. There is a high likelihood that such vulnerable sites will soon be compromised, so they are put into a special queue for frequent re-scanning.

Honeypots are a staple of threat intelligence collection. The ATI Research Center has a global network of sites with services such as SSH, FTP, RDP, VoIP, and HTTP available. Any connection attempts to these services are logged and cross-referenced with data collected from other points in the network. This identifies Internet Scanning sites looking for vulnerable servers. They may change locations frequently, so it’s important to have a large and global network to detect their activity.

SPAM is another resource leveraged by the ATI Research Center to collect inputs for analysis and validation. Ixia uses a variety of both commercial and internal feeds for collecting SPAM information that is then fed into a machine learning engine. SPAM email often contains embedded hyperlinks which direct an unsuspecting user to malicious web sites masquerading as legitimate ("click here to update your banking password!") and the link when accessed will often push malware onto the user's system. Ixia follows these links and analyzes any discovered binaries for malicious content so users will be protected from such phishing sites. If they are discovered to be phishing sites, then screen shots of the offending site are captured and stored for use in Rap Sheets. Inspection of phishing sites is one of the harder tasks to automate, so the ATI Research Team invests manual effort in double-checking these sites.

Every blocked IP address is individualLY validated.

Binary Analysis is often the final and key step in determining whether a site is engaged in malicious activity. ThreatARMOR is very selective in blocking sites, only preventing connections to sites with 100% proof of malicious activity. The ATI Research Center performs dozens of static and dynamic (or "sandbox") analyses of each target binary. Static analysis is used to analyze the file itself, looking for known signatures within the binary of known malware samples. Sandboxing inspects the binary as it executes, paying particular attention to any system and network calls made by the binary. Access or modification of system files, network connections to known bad sites, and installation of additional files and binaries are typical indicators of malicious behavior. If the binary launches an installer, screen captures are recorded for use in the Rap Sheets. Of course, when a binary is found to be malicious, we examine and record any network connections it initiates, as the destinations of those connections are immediately fed into the candidate list of potentially bad sites.

Regardless of the source of the candidate IP address, no IP address is added to the ATI database until it has been individually validated. This means there are no confidence scores in the ThreatARMOR because there is no tuning required. No site is blocked without 100% assurance. And because the same ATI feed is used in the Ixia Perfect Storm and BreakingPoint products, IP addresses flagged as malicious in those products are known to be malicious with 100% assurance.

SORTING OUT CONNECTED-WORLD MALICE

No longer "corner cases," many web sites host both valid and malicious pages, spread across multiple domains. This is common in content delivery networks (CDNs) and hosting environments, where multiple domains are hosted at a single IP address. For some sites like this, such as Dropbox, Amazon, and Azure, ThreatARMOR allows IP addresses even though malicious activity may have been found there. ThreatARMOR has this ability to block connections to malicious domains while allowing access to others at the same IP address because the ATI Research Center tracks malicious activity on a per-domain basis.

ThreatARMOR was never designed to replace a firewall or antivirus system—those tools still handle the remaining traffic from these sites, but do so with higher performance because of their reduced load. For mutlihosting sites only, ThreatARMOR will inspect the domain in an HTTP query and block the connection if it is to a known-bad domain. In some cases, if a particular IP address is heavily infected with bad domains, all domains at that IP are blocked because of the clear evidence that infections and malware distribution are rampant at that site.

ThreatARMOR does not rely on signatures.

The ATI team continually re-scans each IP address in the blocking database at least once per day, oftentimes even more frequently. Sites are aggressively removed from the database when they are found to have been cleaned up. The threshold for removal may be adjusted based on various factors, but the typical time for removal is about three days from last detection. Of course, sites which have been observed to distribute malware tend to be re-infected frequently, so even after sites have been removed from the blocking database they are re-scanned frequently for subsequent reinfection. Some sites are hijacked or broken into and used for malicious purposes, others are owned by bad actors and taken out of commission when a particular malware campaign stops. In either case, sites will often cycle in and out of the blocking database.

All of this research is used to not only block malicious connections attempting to enter or leave a network, but also present useful information to users about what threats are confronting their network and why they are being blocked. This information is presented to users in the form of Rap Sheets and record any network connections it initiates, as the destinations of those connections are immediately fed into the candidate list of potentially bad sites.

The Five Categories found in Rap Sheets

Rap Sheets are arranged in the following categories:

Malware

This category is for malicious software in many different forms. Many malware instances morph or are regenerated daily so they can avoid signature-based detection systems. ThreatARMOR blocks all downloads from known-bad sites and does not rely on signatures.

Phishing

Phishing is an attack where seemingly valid emails are sent to unsuspecting users in an attempt to trick the user into clicking on an embedded link in the email. Phishing emails typically appear to be from valid senders, such as an employer or the user's bank. When the recipient follows a link, they are then prompted to enter sensitive information ("Enter your Social Security Number to claim your bank fee refund") or malware is pushed down to their system. By blocking phishing sites in the database, ThreatARMOR protects users from phishing attacks.

Botnet

A botnet is an orchestrated army of infected hosts that unknowingly participate in malicious campaigns under the control of a botnet herder or controller. The botnet controllers communicate with these infected hosts using "Command and Control (C2)" connections, which can be sent over many different common protocols such as IRC, HTTP, DNS, and others, and are typically encrypted to avoid detection. Botnet controllers send commands to infected hosts directing them to leak sensitive data, download additional malware, or attack other targets in a DDoS attack. By blocking all connections with all botnet controllers in the ATI database, ThreatARMOR prevents these C2 connections and prevents the infected hosts from receiving additional commands or leaking sensitive corporate or personal data. ThreatARMOR also identifies which internal hosts have been attempting to communicate with a botnet controller and need to be cleaned up.

Exploit

Hackers often conduct large-scale reconnaissance, looking for hosts with exposed vulnerabilities. Many services have to be advertised to the Internet in order to conduct business – HTTP/HTTPS, SSH, VoIP, RDP, VPN and others. Vulnerabilities are discovered over time in the software packages which advertise these services, and hackers are constantly scanning for sites which run vulnerable services. The ATI Research Center has deployed a global network of honeypots which advertise common services and monitor for incoming connections. Sites which attempt to connect to these services are cataloged and analyzed, and they are added to the blocking database if they are found to be performing systematic scans.

Hijacked

Hijacked IP ranges are stolen from their legitimate owners, typically by corrupting the routing tables of Internet backbone routers. Once hijacked, they are used for malicious purposes such as phishing and malware distribution. This technique often evades security solutions which rely on IP reputation and URL filtering, since the domain and IP addresses may have years of valid use before being hijacked. The ATI Research Center continuously tracks the millions of hijacked IP addresses on the Internet and ThreatARMOR blocks all connections to and from them, since those connections by definition are not with the valid network owner.

Summary

The ATI Research Center is constantly evolving its detection and analysis capabilities, as well as expanding the network of sensors it uses to detect malicious hosts on the Internet. ATI's database is updated in real time as new threats are detected or cleaned sites are removed, and ThreatARMOR devices update from the ATI cloud as frequently as every five minutes. The detection, collection, and analysis capabilities of the ATI Research Center are used to enhance many Ixia products, including ThreatARMOR, BreakingPoint, and the ATI Processor. As used in ThreatARMOR, they keep users safe from a variety of attacks and free up valuable resources on firewalls, SIEMs, IPS/IDS systems, and the teams that run them, making networks more secure and operations teams more efficient.