Cisco 2Q11
Global Threat Report

Featuring Data from Cisco Security Intelligence Operations

cisco

Contents

Key Highlights	3
Introduction	4
Advanced Persistent Threats	5
Cisco ScanSafe: Web Malware Events	6
IPS Isn't Magic - But That's Okay	8
Cisco Intrusion Prevention System and Remote Management Services	10
Byting Back with Rapid Research	12
Cisco IronPort: Global Spam Trends	13
Conclusion	14

Key Highlights

- The rate of unique instances of malware more than doubled in the second quarter, from 105,536 unique instances of web malware encountered in March 2011 to 287,298 unique instances in June 2011.
- The average encounter rate in the second quarter was 335 encounters per enterprise per month, with the highest peaks in March (455) and April (453).
- From an encounter-per-seat perspective, companies with 5,001 to 10,000 employees and companies with more than 25,000 employees experienced significantly higher malware encounters compared to other size segments.
- Intrusion prevention and detection systems (IPS/IDS), as well as tools like NetFlow, can provide valuable ongoing alerting and forensics for early threat detection.
- Brute-force SQL login attempts increased significantly during the second quarter, coinciding with increased reports of SQL injection attacks—resulting in data breaches throughout the period.
- IPS event firings indicative of denial of service (DoS) attempts increased during the second quarter.
- Global spam volumes remained fairly steady throughout the first half of 2011, with a slight decrease observed in the second quarter.
- Phishing levels measured in proportion to all spam increased in the second quarter, reaching 4 percent of the total volume of spam in May 2011.

Introduction

The first half of 2011 (1H11) witnessed a seemingly nonstop array of data breaches directed at companies, and sometimes individuals, across many sectors. Equally as diverse as the targets were the motivators behind the attacks. In many of the breach incidents, customer data was stolen and publicly published. In some of those cases, the attackers claimed the motive was to shed light on security issues. But in other cases of stolen and published customer data, attackers claimed to be doing it for the "lulz."

Some incidents resulted in stolen or compromised intellectual property related to digital certificates and encryption technologies. In other incidents, attackers gained access to sensitive information but it was not immediately clear whether they had also stolen the information they accessed.

Advanced persistent threats (APTs) played a key role in many of the breaches. APTs are generally rootkit-enabled, exhibit no visible symptoms of infection, and often employ escalation of privilege and other forms of exploit to traverse the compromised network. Malware used in this type of attack can bypass signature detection and other standard forms of security protection. As a result, APTs are seldom passively discovered; instead, active and ongoing analysis of in-house security data sources and traffic analysis is required.

In this installment of the *Cisco® Global Threat Report,* we take a closer look at APTs and some of the methods that can be used to better identify an APT or intrusion event in your own network.

Contributors to the *2Q11 Cisco Global Threat Report* include:

Jay Chan
Gregg Conklin
Raymond Durant
John Klein
Mary Landesman
Armin Pelkmann
Shiva Persaud
Gavin Reid
Clad Skipper
Ashley Smith

All contents are Copyright © 2011 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

Advanced Persistent Threats

Prior to the Internet and today's borderless networks, if an individual wanted to steal corporate secrets they would first need to gain physical access to where the information was housed. Today, however, sensitive data is no longer limited to physical facilities, and attackers can gain access remotely and anonymously.

Malware has evolved along with the Internet and is now the tool of choice for would-be attackers. But the key lies in its ability to remain surreptitious: It must enable the attacker to remotely manipulate a system while remaining virtually invisible to standard defenses. This specialized class of malware, termed "advanced persistent threats" (APTs), presents a widely publicized yet little understood security challenge.

Detecting an APT is not an easy task. Given the way these threats operate, there is no "silver bullet" for identifying them in a network. "If we could identify APTs by a software signature, we wouldn't need to call them 'advanced persistent threats,'" explains Gavin Reid, manager of the Computer Security Incident Response Team (CSIRT) at Cisco. "If anyone attempts to sell your organization a hardware or software solution for APTs, they either don't understand APTs, don't really understand how computers work, or are lying—or possibly all three."

In large part due to the detection challenge, initially many questioned whether APTs even existed. This skepticism was finally put to rest in January 2010, when Google chief legal officer David Drummond announced that Google had experienced an APT on its own network and reported that "at least twenty other large companies" had been similarly targeted.

While specific attack details were never revealed, this candid disclosure by Google confirmed both the prevalence and pervasiveness of APTs. Today, the challenge isn't in proving that APTs exist—the challenge is to separate the APT from other malware and forensically identify it in a timely manner. As Reid explains, "There are no easy answers. With APTs, like any other tough security problem, the solutions may be complex, but the methodology is simple: Identify what your available options are, and then execute."

According to Reid, an organization's ability to detect and respond to APTs can improve when well-understood computer security incident response capabilities are deployed:

- The capacity to produce, collect, and query logs—the more the better, but at least the important ones—from a security perspective (e.g., host logs, proxies, and authentication and attribution logs).
- Some form of deep packet inspection that covers all the important "choke points" on your network.
- The ability to quickly query network connections or flows through NetFlow (or a similar service) across all network choke points.
- Development of trust-based relationships with other organizations to share intelligence on events. For instance, join an organization like the Forum of Incident Response Teams (FIRST.org), which helps facilitate this type of information sharing.
- Some degree of malware analysis (in-house or outside).

Reid also offered two examples of how his team has refined their approach to detecting APTs:

1. "Three years ago, we instituted a program to provide deeper analytic network and system forensics to a select group of employees that are likely targets for APTs—that is, individuals who have access to data that a criminal would want," says Reid. "For this group, we follow up and do more advanced watching and investigation."

2. "If you have the ability, capture and store all PDFs that come into your company over email, along with the associated email headers. On a regular basis, do some automated, additional checking beyond your company's antivirus solution to help detect PDFs that contain more than the content. Even though normal antivirus systems will not detect or stop these threats, often it is relatively easy to see they are modified by using simple string searches or running multiple antivirus scanners."

On a final note, Reid adds, "If you have something of interest and you're not seeing APT attacks in your organization, it is probably not that they are not occurring or that you're safe. It's more likely that you may need to rethink your detection capabilities."

Based on "Cisco CSIRT on Advanced Persistent Threat", by Gavin Reid, published March 2011.

(http://blogs.cisco.com/security/cisco-csirt-on-advanced-persistent-threat/)

Cisco ScanSafe: Web Malware Events

While APTs are frequently the result of directly targeted attacks, any malware encounter can lead to an advanced persistent threat. Attackers can—and do—segregate infected computers into interest areas and modify their methods accordingly. For example, after initial infection by a common downloader Trojan, subsequent information may be collected from infected machines to identify those systems more likely to lead to sensitive information. Subsequently, those "interesting" machines may be delivered an entirely different set of malware than would other "non-interesting" computers.

The majority of today's malware encounters occur via the web. During the first half of 2011, enterprise users experienced an average of 335 web malware encounters per month, with the highest peaks occurring in March (455) and April (453). This is shown in Figure 1.

Bear in mind that averages may not reflect real-world experience; the actual number of encounters per enterprise can range from a dozen per month to tens of thousands per month, depending on the number of employees, industry sector, and other factors.

Unique web malware encounters increased significantly throughout 1H11, from 72,294 unique encounters in January 2011 to 287,298 in June (Figure 2). Despite the increase in encounters, the number of unique malware hosts and unique IP addresses remained relatively consistent between March 2011 and June 2011 (Figure 3).

Companies in the Pharmaceutical and Chemical and the Energy and Oil sectors continued to be at highest risk of web malware throughout 1H11. Other higher risk verticals throughout the quarter included Transportation and Shipping, Agriculture and Mining, and Education. The median rate for all verticals is reflected as 100 percent – anything above 100 percent has a higher-than-median encounter rate and anything below 100 percent is below the median for all (Figure 4).

Figure 1 **Average Web Encounters per Enterprise, 1H11**

Source: Cisco ScanSafe

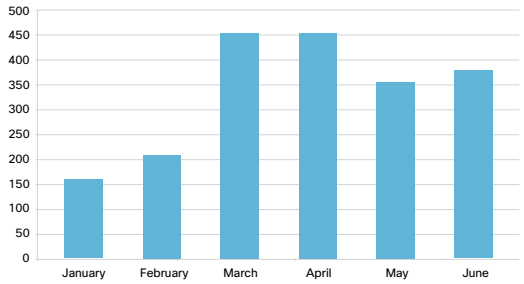

Figure 2 **Unique Web Malware Encounters, 1H11**

Source: Cisco ScanSafe

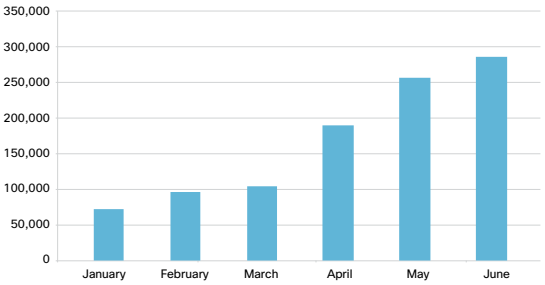

Figure 3 **Unique Malware Domains and IPs, 1H11**

Source: Cisco ScanSafe

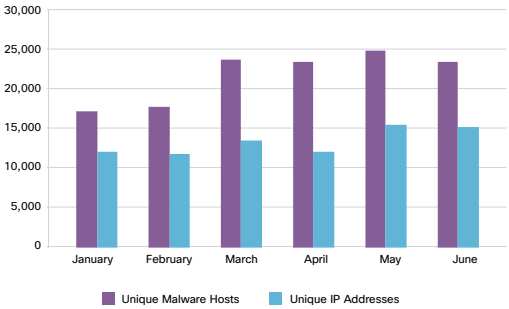

All contents are Copyright © 2011 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

Figure 4 Vertical Risk, 1H11
Source: Cisco ScanSafe

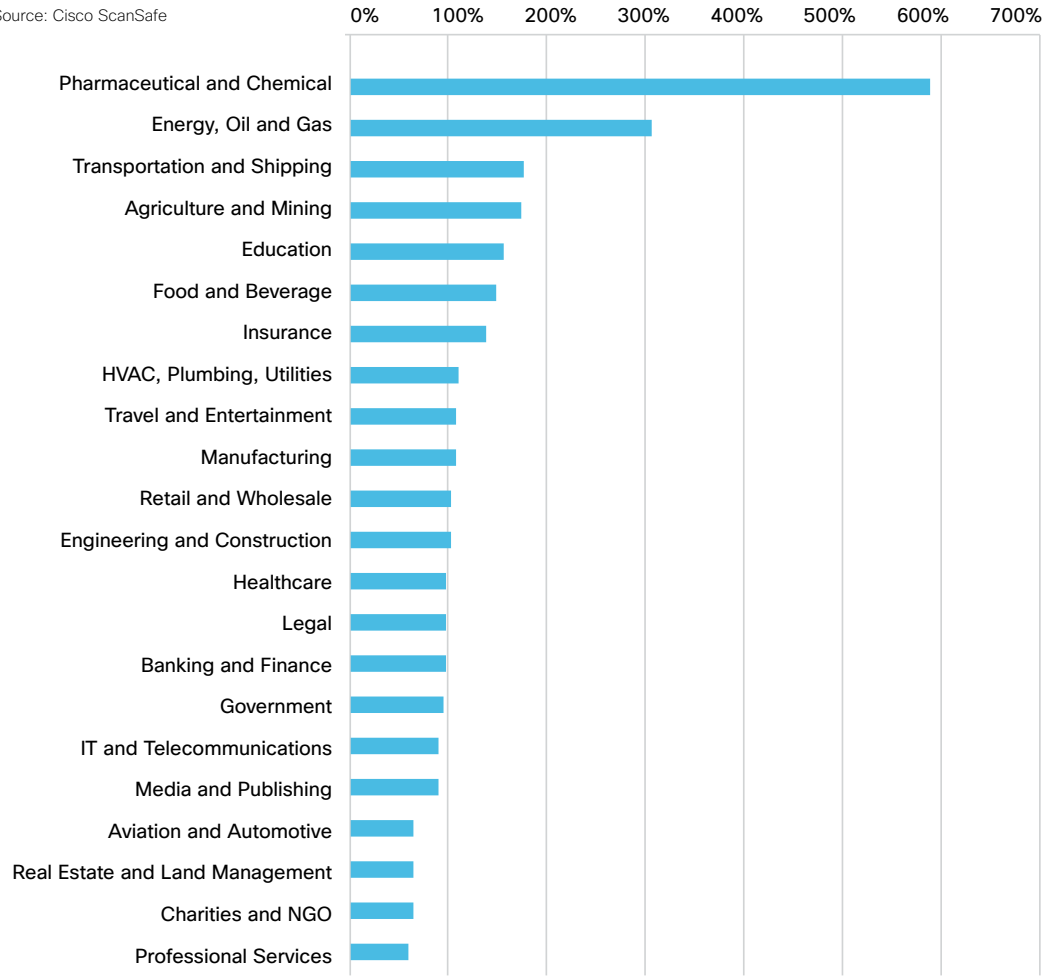

Some of these companies may have fewer than 200 employees, while others may have tens of thousands or even more. To help contextualize malware encounter risk, Figure 5 depicts the median encounter rate based on customer size.

Figure 5 Customer Size Encounter Risk, June 2011
Source: Cisco ScanSafe

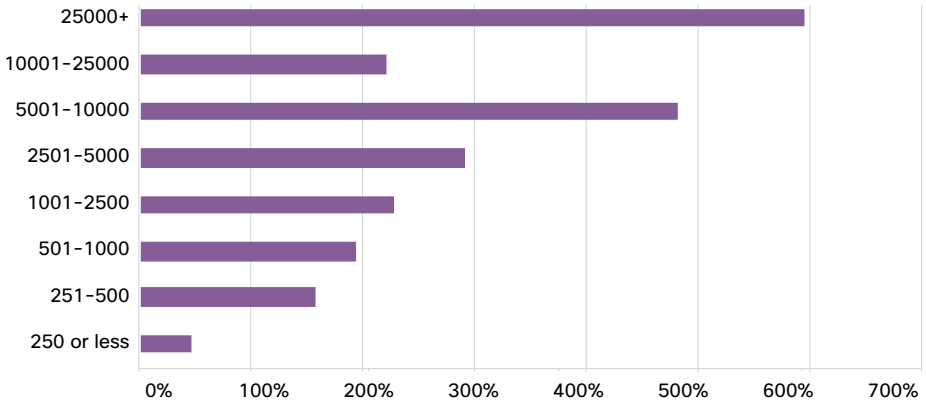

IPS Isn't Magic–But That's Okay

By Gavin Reid, Manager, Cisco CSIRT

Most CSIRT teams end up with responsibilities that can seem like counting sand on a beach or searching for the proverbial needle in a haystack. Many look for, or are sold, "magic" security products that claim to reduce alerts to only the important.

Magic never worked well for me so I'm not comfortable with relying on it. Instead, I would suggest making sure you're asking the right questions when monitoring. Start with what is possible—events you know you can take action on—and work out from there. You don't need a magic algorithm, just some dedication and common sense.

Make alerts "human-readable"

The first tip is assigning IDS location (locale, in Cisco IPS) variables. At Cisco we have IDS variables defined for anything meaningful. We use them in tuning and in custom IDS signatures. Most important, we use them to make alerts "human-readable." Here's an example:

```
Assign "locality" to the source ip and destination ip.
sigDetails=STOR command on dst ports 20 and 21" src=64.104.X.X srcDir=DC_OTHER_DC_NETS srcport=41507
dst=210.210.X.X dstDir=OUT dstport=21
```

If Cisco's monitoring team viewed the above alert, they would immediately see, without doing any host lookups, that a host in one of the company's data centers (DC_OTHER_DC_NETS) made an outbound FTP connection to an outside site (OUT). As Cisco investigates any outbound transfer from its data centers to the Internet, the monitoring team would immediately escalate this alert without the need for additional research.

Making the alert easily understandable is also very useful in IPS custom signature making. For example, a network management locale would allow security professionals to instantly tune management systems that may legitimately perform discoveries from IDS signatures that look for one-to-many connections (worm-like scan activity). Below you can see us adding some systems to the locale variable:

```
xxx-dc-nms-1# conf t
xxx-dc-nms-1(config)#
service event-action-rules rules0
variables MGT_SYSTEMS address 10.6.30.5,10.6.30.6,10.30.6.7,10.50.1.5,10.50.1.6,10.50.1.7
```

And here we use a filter to tune those management systems from multiple IDS signatures:

```
filters insert drop_mgt_system_alerts
signature-id-range 4003,3030,2100,2152
attacker-address-range $MGT_SYSTEMS
victim-address-range $IN
actions-to-remove produce-alert|produce-verbose
```


All contents are Copyright © 2011 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

When we find new management systems—usually through detection, but sometimes IT lets us know—it is easy to update the variable and, in turn, all the IPS signatures that use that variable. So, if our monitoring team sees a one-to-many scan coming from MGT_SYSTEMS, they know it's expected.

Nothing covered here is very glamorous or difficult. Certainly, none of it is magic—or even perfect. But all of it can help to effectively reduce risk with IDS.

Adapted from "IPS Isn't Magic—But That's Okay", by Gavin Reid, published March 2011.

(http://blogs.cisco.com/security/ips_isnt_magic_but_thats_okay/)

Baseline to Detect Mass Outbreaks

There is another real-world, zero-day outbreak detection method that can help you understand whether what's happening on your network is or is not "normal": baselining. Consider this the "no-magic" zero-day mass outbreak detection method.

Baselining can be applied to any type of IDS. Security professionals should chart the infected host count per detection vector, establish thresholds, and then trend. When the thresholds are breached, it is a great indication of a mass outbreak.

Another type of baselining that can enable quick outbreak detection is recording the number of IP addresses found per run of each malware report, and then looking for deviations from what is expected.

Cisco Intrusion Prevention System and Remote Management Services

Ongoing data analysis can help you baseline what is normal for your enterprise, an important first step in readily identifying new or previously unseen incidents. Figures 6 and 7 show Intrusion Prevention System event firings observed by Cisco Remote Management Services (RMS) and Cisco Intrusion Prevention System (IPS) from April 1, 2011, through June 30, 2011.

Figure 6 **Top 10 Signature Firings, 2Q11**
Source: Cisco RMS

Signature	Events
Generic SQL Injection	64.21%
Malformed SIP Packet	10.04%
Cisco Unified Videoconferencing Remote Command Injection	6.95%
TCP Hijack	2.27%
Cisco CDS Internet Streamer Web Server Vulnerability	1.67%
Web View Script Injection Vulnerability	1.62%
TCP SYN/FIN Packet	1.56%
Gbot Command and Control over HTTP	1.48%
Impossible IP Packet	1.33%
SNMP Community Name Brute-Force Attempt	1.04%

Figure 7 **Top 25 Port Activity, 2Q11**
Source: Cisco RMS

Port	Percent
80	72.27%
5060	16.11%
443	1.85%
161	1.67%
40436	1.57%
25	0.83%
22	0.82%
4500	0.80%
455	0.67%
20	0.63%
1935	0.47%
554	0.39%
1433	0.30%
8080	0.19%
7654	0.13%
1836	0.11%
3985	0.09%
7777	0.06%
500	0.06%
8000	0.05%
1583	0.05%
7717	0.04%
42810	0.04%
50354	0.04%
53	0.04%

All contents are Copyright © 2011 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

As seen in Figure 8, denial of service (DoS) attacks had a steady presence throughout 1H11, with the most significant peaks occurring in May and June 2011. While once largely prank-related, DoS attacks are increasingly politically and financially motivated.

Brute-force SQL server login attempts also increased during the second quarter, correlating with increases in SQL injection attacks during the same period (Figure 9).

Figure 8 **DoS Event Firings, 1H11**
Source: Cisco IPS

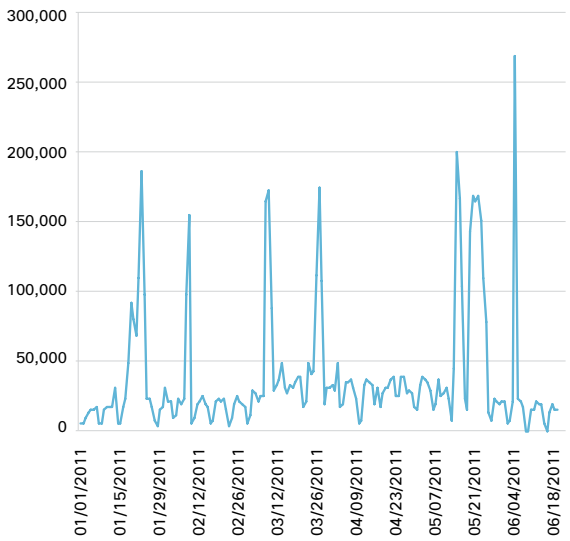

Figure 9 **Brute-Force SQL Login Attempts, Sensor Count 2Q11**
Source: Cisco IPS

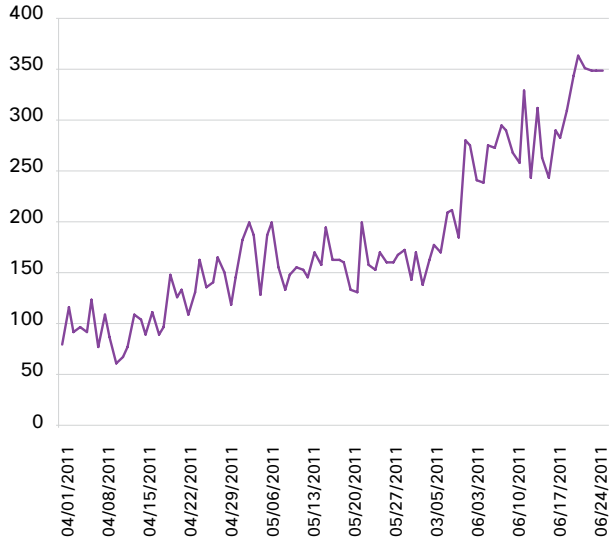

Using NetFlow for Incident Response

By collecting and storing flow records in a searchable database, security professionals can improve their ability to spot intrusions and other potentially dangerous activity. The following examples illustrate how NetFlow can be used to support incident response.

Identifying compromised machines. By querying NetFlow, administrators can determine where an attack originated and what other machines may have been impacted. For example, if botnet activity is detected, administrators can query a NetFlow database for all connections to the IP address and port of the malicious server.

Policy-based alerts or reporting. Administrators can verify that connections destined for areas within their enterprise network are in accordance with company network and security policies. This can help ensure employees are not doing things such as web surfing from a data center system.

Evaluating firewall access control lists. For example, if a network has a web server and a DNS server in a DMZ and administrators have applied access control lists to block all other traffic, they can set up alerts for any traffic not on Ports 80 or 53.

Detect covert channels and/or web-based uploads. This can be useful even in areas where data is encrypted. You can query for web traffic where the ratio of upload to download doesn't match expected behavior. For example, if a user connects to a web server and uploads 20 MB of data while downloading 200K, the user is probably uploading files to the web server or tunneling traffic.

Excerpted from "NetFlow for Incident Response", by Gavin Reid, published January 2011
(http://blogs.cisco.com/security/netflow-for-incident-response/#utm_source=rss&utm_medium=rss&utm_campaign=netflow-for-incident-response)

Byting Back with Rapid Research

By Shiva Persaud

Often the questions that surface when investigating security incidents cannot be answered with information that is readily available. Fortunately, the community has produced a rich set of tools to facilitate finding the proverbial needle in a haystack.

When it comes to analyzing traffic captures, I turn to the Wireshark suite because of its rich set of protocol dissectors and flexible command-line tools. TShark makes it easy to search through large amounts of traffic captures to find exactly what you are looking for. For example, the following command identifies RPC traffic containing shellcode used by a Conficker variant.

```
$tshark -r traffic_sample.pcap tcp contains \
e8:ff:ff:ff:ff:c2:5f:8d:4f:10:80:31:c4:41:66:81 and tcp.dstport eq 445
```

To learn how a given protocol is dissected, I prefer to browse Packet Details Markup Language (PDML) output from TShark because of the large amount of information available in the output. Through reading PDLM, I inadvertently learn more about the protocol. Following is the command that pipes the PDML output for a traffic capture containing a DNS query to vim:

```
$tshark -r dns.pcap -T pdml dns | vim -
```

Here is a PDML snippet from the command above:

```
<field name="dns.qry.name" showname="Name: cisco.com" size="11" pos="54"
show="cisco.com" value="05636973636f03636f6d00"/>
```

I now know that Wireshark calls the DNS name field dns.qry.name. I can hone in on the host that was queried by running:

```
$tshark -r dns.pcap -T fields -e dns.qry.name cisco.com
```

It is possible to use display filters to create traffic capture files which contain only the traffic you are interested in. The following script takes a packet capture (pcap) filename as input and overwrites that file with a new pcap that contains only TCP traffic:

```
$cat tcp_only.sh
#!/bin/bash

tshark -r ${1} -w tcp_${1} tcp
mv tcp_${1} ${1}
```

Of all the scripts I have written that wrap around TShark, the one I use the most splits a traffic capture file into several smaller files that each contain only one TCP stream. I'll leave this as an exercise.

The information you need to draw conclusions when doing security research isn't out of reach. With the right tools, you will find those nuggets in no time.

Happy hunting! ~ Shiva

All contents are Copyright © 2011 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

Cisco IronPort: Global Spam Trends

The 2011 takedown of segments of Rustock, combined with multiple spam botnet takedowns in 2010, continues to have positive impact on overall spam volume. Figure 10 reflects global spam volume as reported through Cisco SenderBase Network participants.

Figure 10 **Global Spam Volume, 1H11**

Source: Cisco IronPort

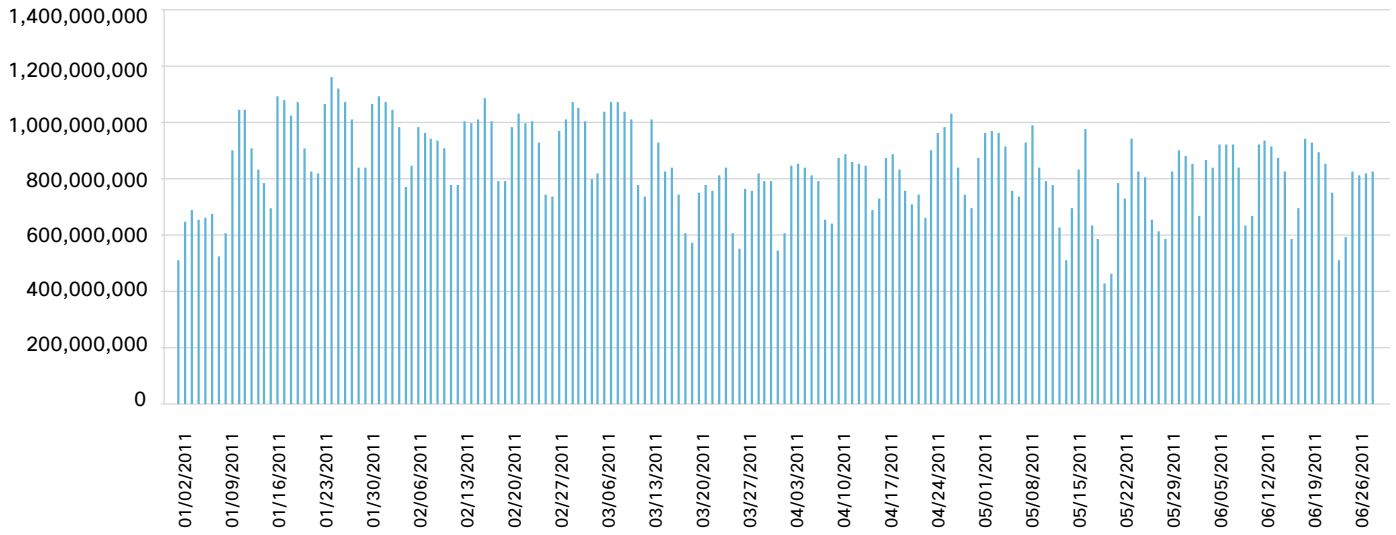

As seen in Figure 11, although Spam remained fairly steady and even exhibited a slight decrease during the second quarter, phishing attacks increased during the same period.

Figure 11 **Percent of Phishing in Spam Volume, 1H11**

Source: Cisco IronPort (Spam Traps / User Submissions)

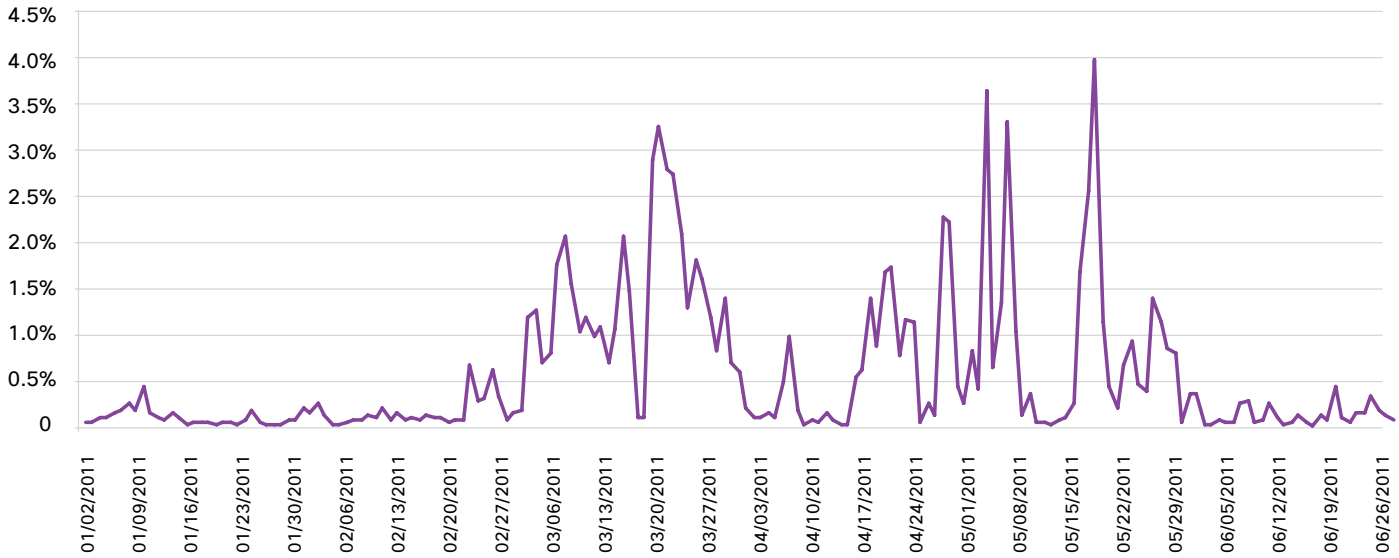

Conclusion

Cybercriminals are launching more targeted, sustained, and hard-to-detect attacks. But organizations are not defenseless against these intrusions. While there is no magic bullet, many approaches to monitoring, detection, and incident response are readily available—and often free. As discussed in this report, security professionals should consider embracing strategies such as:

- Using NetFlow to support incident response by identifying zero-day malware that has bypassed typical security controls; exposing compromised machines; verifying that connections destined for areas within the enterprise network are expected in accordance with company network and security policies; evaluating firewall access control lists; and detecting covert channels and/or web-based uploads.
- Taking an analytical approach to detecting APTs and deploying well-understood computer security incident responses. These include the ability to produce, collect, and query logs; some form of deep packet inspection to cover key network "choke points"; the ability to quickly query network connections or flows through NetFlow or similar services; the development of trust-based, intelligence-sharing relationships with other organizations; and malware analysis.
- Assigning IDS location variables to make alerts more "human-readable," so that security teams can instantly identify and escalate an incident without needing to first decipher the alert.
- Baselining to detect anomalous events. Approaches include charting infected host count per detection vector, establishing thresholds and trending, or recording the number of IP addresses found per run of each malware report and then looking for deviations from what is expected.
- Collaborating and sharing knowledge. Develop trust-based relationships with other organizations to share intelligence on events. This is a long process that you will have to purposely resource and tend. A great start would be joining an organization like FIRST.org.

Regardless of the motivation of attackers—whether it's to steal data, prove a point, or grab a laugh—breaches are costly and the number of incidents continues to increase. Combined, the above approaches can help security teams more quickly identify and remediate intrusions on their own networks, and help avoid potential losses.

All contents are Copyright © 2011 Cisco Systems, Inc. All rights reserved. This document is Cisco Public Information.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters Cisco
Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. C02-681613-00 7/11