VERISIGN DISTRIBUTED DENIAL OF SERVICE TRENDS REPORT

ISSUE 3 – 3RD QUARTER 2014

VERISIGN®

VERISIGN

CONTENTS

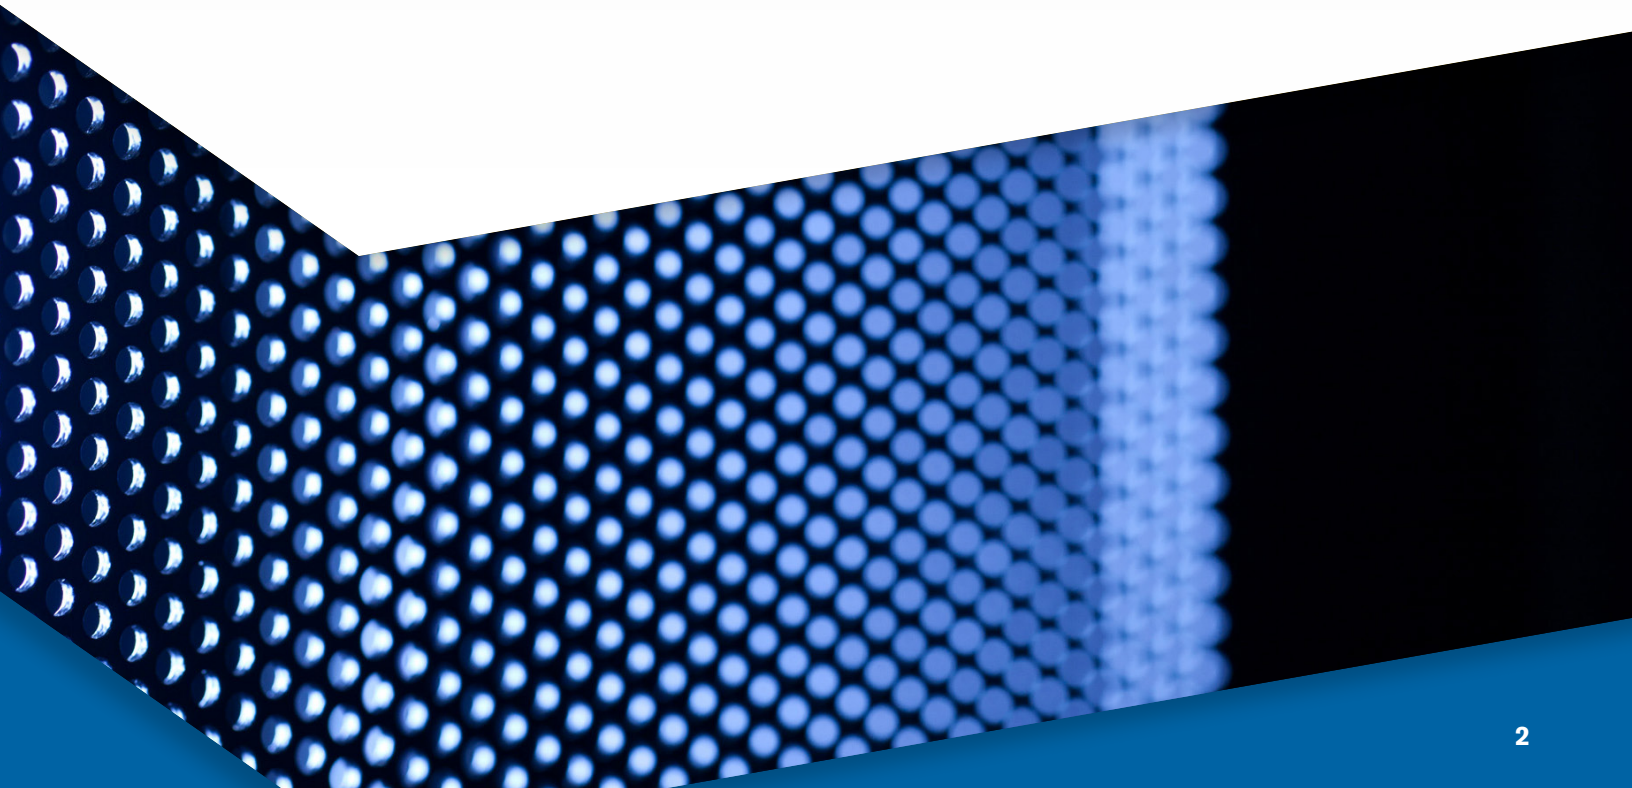

VERISIGN

>20 PERCENT
of attacks were greater than 10 Gbps

Attackers attempted an average of **3.3** separate attacks per customer in Q3

EXECUTIVE SUMMARY

This report contains the observations and insights derived from mitigations enacted on behalf of, and in cooperation with, customers of Verisign DDoS Protection Services and the security research of Verisign iDefense Security Intelligence Services. It represents a unique view into the attack trends unfolding online for the previous quarter, including attack statistics, DDoS malicious code analysis and behavioral trends.

SUMMARY OF FINDINGS: Vulnerabilities and Attack Trends in Q3 2014

For the period starting July 1, 2014 and ending Sept. 30, 2014, Verisign observed the following key trends:

- The number of attacks in the 10 Gbps and above category grew by 38 percent from Q2 to represent more than 20 percent of all attacks in Q3.
- Attackers were persistent in launching attacks against targeted customers, averaging more than three separate attempts per target.
- The most frequently targeted industry this quarter was Media and Entertainment, representing more than 50 percent of all mitigation activity.
- Due to the high frequency of massive attacks experienced in Q2, Verisign observed a decrease in average attack size by 48 percent quarter over quarter, but average attack size increased by 65 percent between Q1 and Q3 2014. Removing the very large attacks from the Q2 data set shows average attack size was 4.6 Gbps, which if compared to the Q3 average attack size, represents an increase of more than 40 percent.
- The largest attacks observed this quarter targeted the E-Commerce industry – peaking at more than 90 Gbps.
- Verisign directly mitigated new types of UDP reflection attacks utilizing the Simple Service Discovery Protocol (SSDP / UDP port 1900).
- Network Time Protocol (NTP) continues to make up the majority of UDP-based reflective amplification attacks, with an observed shift to SSDP over the course of Q3.
- Verisign iDefense analysts reported that the “Shellshock” Bash vulnerability was leveraged globally to efficiently spawn DDoS botnets.
- An increase in the deployment of Linux DDoS malware shows an increased use of high-bandwidth servers in DDoS botnets.

VERISIGN-OBSERVED DDoS ATTACK TRENDS Q3 2014

Mitigations by Attack Size

Large-scale DDoS attack frequency has continued to trend upward as the number of attacks in the 10 Gbps and above category grew by 38 percent from Q2 to represent more than 20 percent of all attacks in Q3 (figure 1). On-premise mitigation defenses and devices are rendered ineffective the moment a DDoS attack exceeds an organization's upstream capacity. If they haven't already, organizations that focus on resiliency should consider deploying cloud-based or hybrid premise/cloud DDoS protection solutions to mitigate application-layer, multi-vector and volumetric attacks that exceed their available bandwidth with minimal increases in operational overhead.

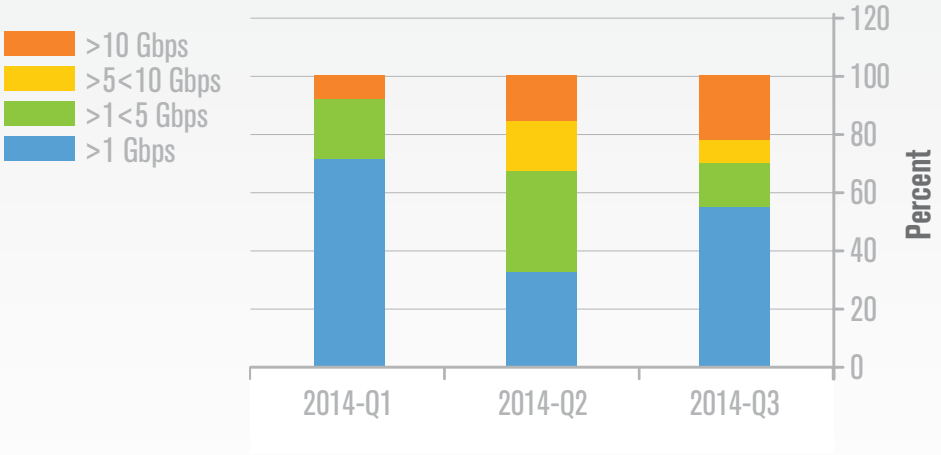

Figure 1: Increase in Attacks Greater Than 10 Gbps

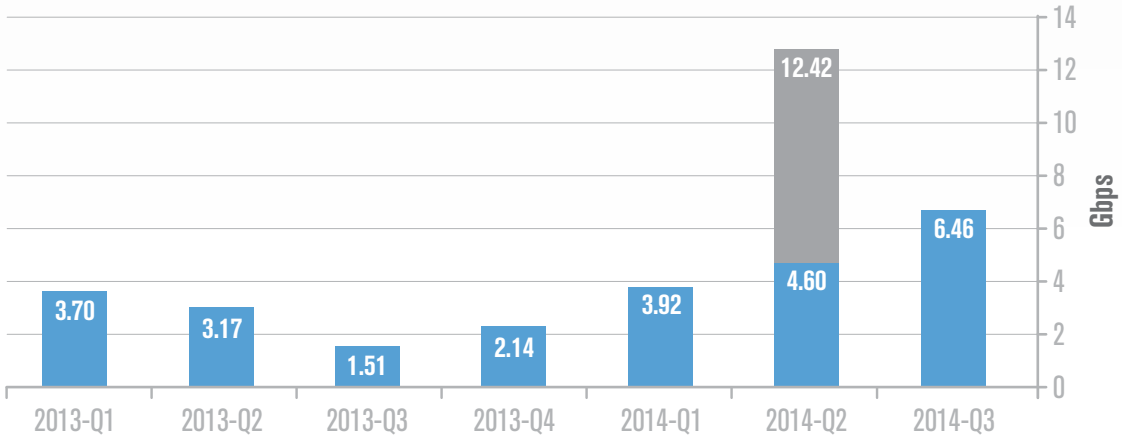

Figure 2: Average Attack Size by Quarter (in Gbps)

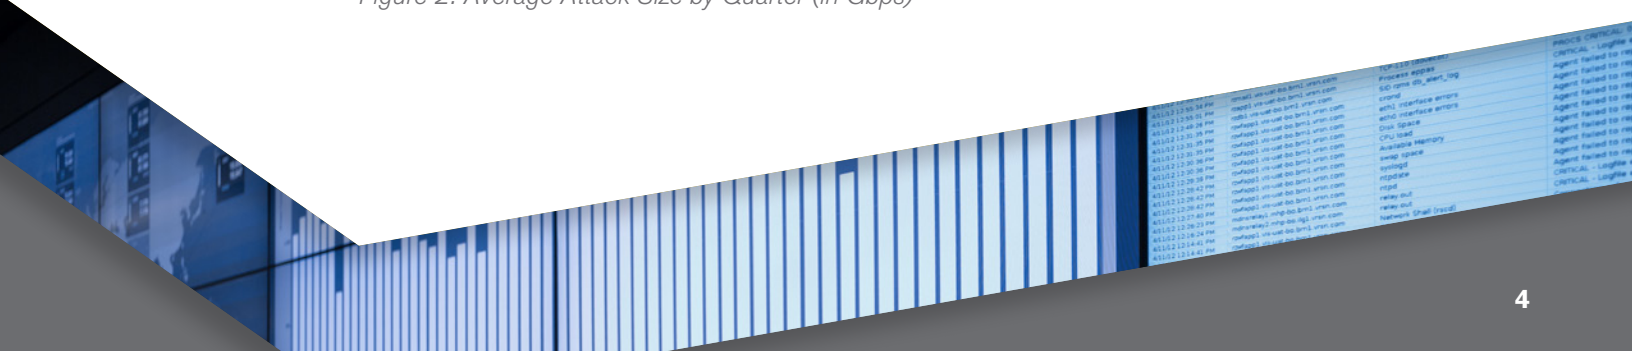

Attacks mitigated by Verisign in the third quarter averaged 6.46 Gbps (Figure 2), which represented a 48 percent drop in average attack size quarter over quarter, but a 65 percent increase in average attack size from Q1 2014. The exceptional increase in average attack size in Q2 (12.42 Gbps) was driven by multiple sustained volumetric attacks in the 200-300 Gbps range. Removing the very large attacks from the Q2 data set shows average attack size was 4.6 Gbps, which, if compared to the Q3 average attack size, represents an increase of more than 40 percent. The largest volumetric UDP-based DDoS attack mitigated by Verisign in Q3 was 90 Gbps; the largest TCP-based attack was more than 30 Gbps.

Mitigations by Industry

DDoS attacks are a global threat and not limited to any specific industry, as illustrated in Figure 3. This comparative data can be helpful in prioritizing security expenditures based upon the observed exposure of your industry to this threat. Industries with the highest risk are generally those who are either active politically, or will suffer significant financial loss from downtime. That said, as Verisign has observed over the past decade, a target can become a target for an ever-expanding array of reasons, and every organization should consider its risk and potential exposure in this context.

Media and Entertainment customers continue to experience the largest volume of attacks, peaking in size at just over 20 Gbps in Q3, (Figure 4) which is more than enough to overwhelm most on-premise mitigation capabilities. The E-Commerce industry, while attacked less frequently, was targeted with the largest attack of the quarter, reaching over 90 Gbps (Figure 4). This attack was a pulsing UDP flood employed in short bursts of 30 minutes or less. It consisted primarily of NTP reflective amplification attack traffic. This activity was aimed at disrupting the critical online commerce capability of the customer and was successfully mitigated by Verisign.

With the 2014 holiday season in full swing, the E-Commerce and Financial industries must be particularly vigilant and prepared for increasing DDoS attacks during their peak revenue and customer interaction season. Historically, Verisign has seen an increase in DDoS activity against these verticals during the holiday season and anticipates that this trend will continue.

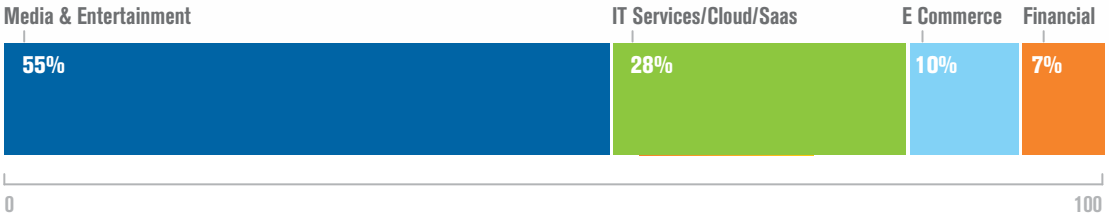

Figure 3: Percentage of Attacks by Industry

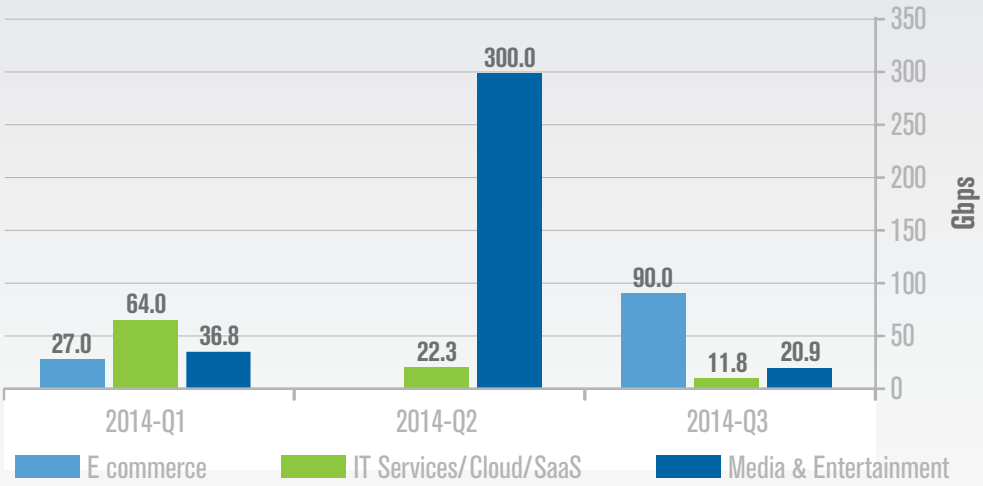

Figure 4: Peak Attack Size by Top Industries

Mitigations by Attack Frequency

Over the course of 2014, Verisign has observed a consistent increase in the number of attacks per customer, including attacks that changed tactics mid-stream. Q3 saw the largest increase in attack frequency, rising to an average of more than three attacks per targeted customer, a figure that rose 60 percent higher than Q2 (Figure 5). The increase in attack frequency, like the increase in attack size, may be attributed to maturation of attackers, easier access to ready-made DDoS botnets and toolkits, and adversary observation of attack impact on their targets. As attackers continue to evolve and become more sophisticated, Verisign expects to see this trend continue into the foreseeable future.

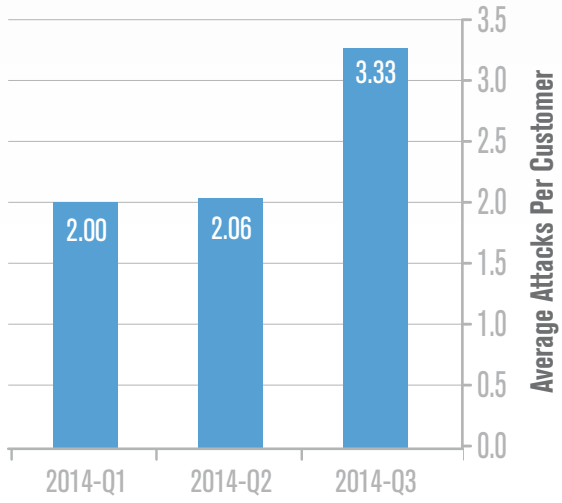

Figure 5: Average Number of Attacks per Targeted Customer (by quarter)

VERISIGN

Largest SSDP attacks in Q3

15 GBPS

and

4.58 MPPS

Feature:
NEW PROTOCOL USED FOR REFLECTION ATTACKS

The Simple Service Discovery Protocol (SSDP)

During this quarter, in addition to NTP-based attacks, we observed SSDP being exploited in UDP-based DDoS amplification attacks. The largest SSDP-based attacks mitigated by Verisign in Q3 targeted the IT Services industry and peaked at just under 15 Gbps and 4.58 million packets per second (Mpps). Verisign was able to quickly detect and mitigate the SSDP attacks encountered this quarter using its monitoring service and globally connected DDoS mitigation platform.

Christian Rossow published a paper on Feb. 22, 2014, Amplification Hell: Revisiting Network Protocols for DDoS Abuse,[1] that identified the bandwidth amplification factor. He identified SSDP as having a bandwidth application factor of as much as 30.8, meaning that a search command on this protocol will return a response 30.8 times the size of the request. The US CERT[2] team issued an alert based on the initial research on Jan. 17, 2014.

Though the amplification it generates is smaller than that possible with DNS or NTP reflection attacks, SSDP attacks still have the capability to overwhelm organizations that are using traditional security appliances to protect their assets. Consistent with other reflective amplification attacks, malicious actors will spoof the source IP when making an SSDP request to target a victim.

For most organizations, SSDP implementations should not need to be open to the Internet. In this case, ingress queries from the Internet targeting this protocol can be blocked at the network edge to protect from this particular vector. Verisign recommends an audit of internal assets, including outbound network flows to ensure that your organization is not being unknowingly leveraged in SSDP-based DDoS attacks.

WHAT IS SSDP?

SSDP is a network protocol used for the advertisement and discovery of network services and presence information, and is most commonly used as the basis of the discovery protocol for Universal Plug-and-Play. Implementations send and receive information using the UDP on port number 1900. SSDP is abused like many other UDP-based protocols, mainly due to its connectionless state, which allows source IP address spoofing, and due to the amplification factor in the response. According to ShadowServer data[3], there are over 15 million devices on the Internet that have SSDP enabled and may be vulnerable to being used in a DDoS attack. Similar to any reflective attack, an attacker must spoof the source IP address of the request to match that of the intended target to cause all vulnerable devices to flood the target with SSDP responses.

1 Christian Rossow. Amplification Hell: Revisiting Network Protocols for DDoS abuse
2 www.us-cert.gov/ncas/alerts/TA14-017A
3 https://ssdpscan.shadowserver.org/

GLOBAL OBSERVATIONS

Verisign Observes DBOT Linux DDoS Malware

Verisign iDefense analysts discovered a variant of the DBOT backdoor which runs on Unix-like systems and is primarily used for DDoS attacks.

This malware is controlled through an Internet Relay Chat (IRC) command-and-control (C&C) channel, and will set its process name to look like common system processes (such as syslogd or crond). It is used not only to perform DDoS attacks, but includes full reverse-shell access and mail-sending capabilities (e.g., for spam) on compromised systems.

Attackers use the IRC protocol to take control of the DBOT backdoor, and then use the established IRC C&C server's channel to issue commands that will instruct the compromised server to execute tasks.

Verisign iDefense analysts observed the following DDoS commands:

Command	Description
udp1	UDP flood against a user: Defined IP address and port with a payload consisting of the string "Tr0x" repeated a random number of times
udp2	Simultaneous flood against a user: Defined IP address on incrementing ports for UDP, TCP, ICMP and IGMP with a payload consisting of repeated "A" character
udp3	UDP flood against a user: Defined IP address and random port with a payload of all zeroes; this function has some flaws that may cause it not to work reliably.
tcp	TCP SYN flood against a user: Defined IP address and port consisting of 1,000 simultaneous connections that are closed immediately after the connection is established
http	Connect to TCP port 80 on a user: Defined IP address and repeatedly perform HTTP "GET /" requests

Verisign iDefense has established that no IP address spoofing currently occurs during the execution of any of the aforementioned built-in DDoS attack commands; as such, most observed attacker IPs will be legitimate, increasing mitigation speed. However, the DBOT malware allows, via its reverse shell function, for arbitrary command execution on a compromised system, giving an attacker the unlimited ability to manually modify attack patterns or install additional DDoS tools as needed. Samples of the malware analyzed by iDefense have an MD5 hash of 579190b74b86f591097b9b6773c1176b.

VERISIGN

"SHELLSHOCK" Leveraged to Deploy Linux DDoS Malware

Verisign iDefense researchers analyzed ELF malware, which was observed to be delivered via the "Shellshock" Bash vulnerability.

"Shellshock" is the common name for a series of critical vulnerabilities (CVE-2014-6271 and later, CVE-2014-6277, CVE-2014-6278, CVE-2014-7169, CVE-2014-7186, and CVE-2014-7187) in the Bash shell application, which is used pervasively in a wide array of operating systems (including OSX, RedHat, Debian and many embedded devices).

The vulnerability is caused by a flaw in the command and argument parser of GNU Bash versions 1.14 through 4.3. The flaw results in incorrect processing of commands placed after function definitions in the added environment variable. The issue allows malicious actors to execute arbitrary and malicious binary code via a crafted environment that enables network-based exploitation.

The malware that leverages this vulnerability communicates with specific hard-coded C&C servers. Connections through these C&C servers result in back-and-forth communication, the receipt of commands and links to additional malicious contents or payloads in the form of raw Pastebin links.

The malware existed in the wild prior to being deployed in the latest campaign leveraging Shellshock. Verisign iDefense has records of the exact strings distributed by the malware as posted on Pastebin as early as Aug. 20, 2014.

The malware checks the infected system for the commonly used set of usernames and weak passwords; (i.e., root, admin, user, login, etc.) to launch DDoS attacks and carry out massive scans for vulnerable systems on the Internet. During the course of analysis, Verisign iDefense researchers discovered several samples that leveraged the recently found vulnerability. Malware samples analyzed by iDefense have an MD5 hash of 5B345869F7785F980E8FF7EBC001E0C7.

VERISIGN

NOTES