VERISIGN DISTRIBUTED DENIAL OF SERVICE TRENDS REPORT

ISSUE 2 – 2ND QUARTER 2014

VERISIGN®

VERISIGN

CONTENTS

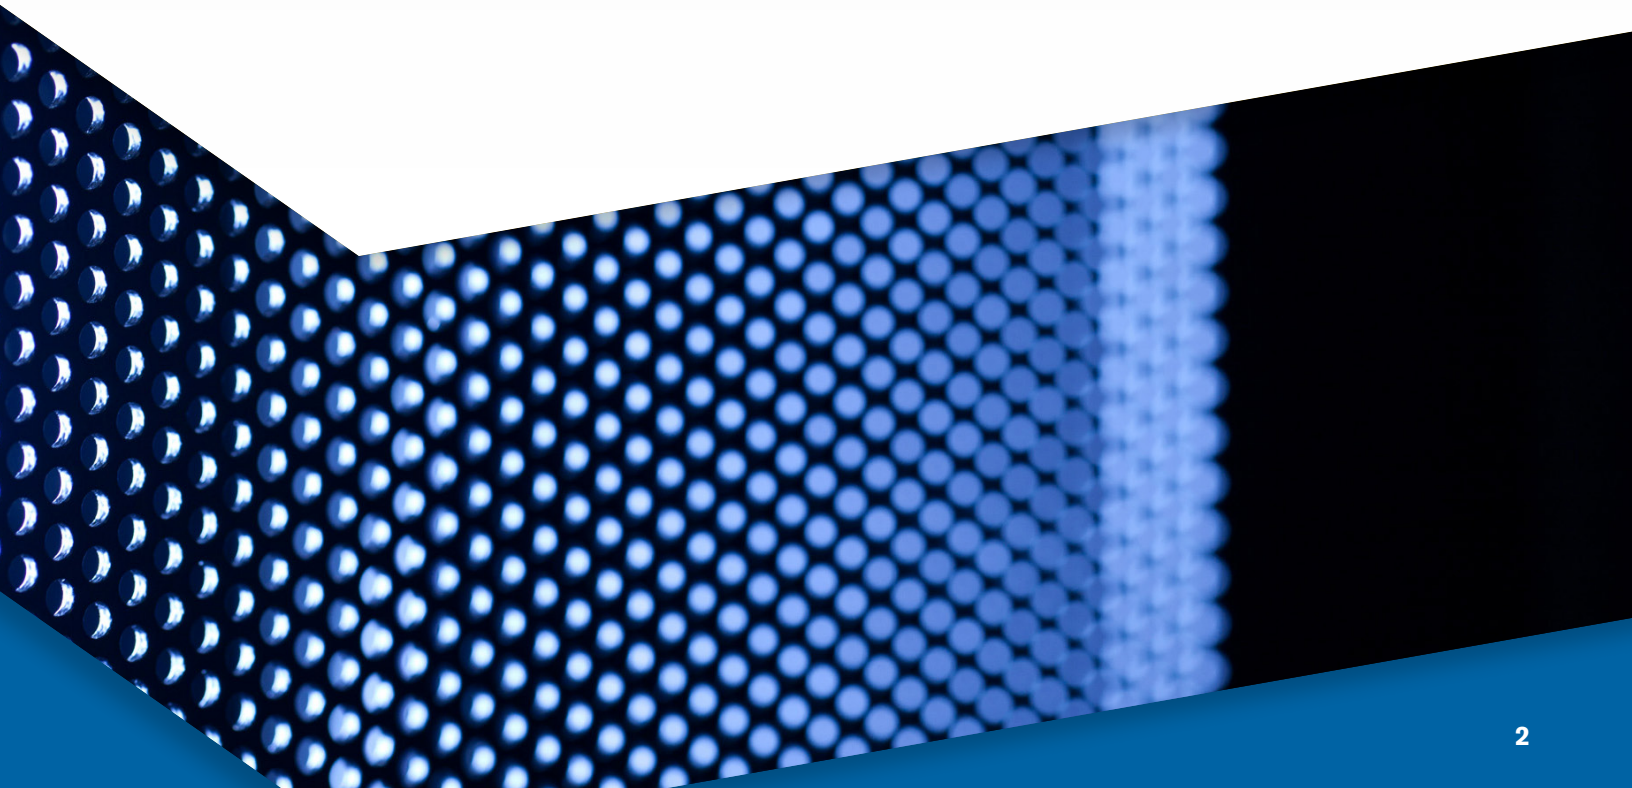

VERISIGN

65
PERCENT
of attacks were greater than 1 Gbps

216
PERCENT
INCREASE
in average attack size since Q1 2014

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. It represents a unique view into the attack trends unfolding online for the previous quarter, including attack statistics, behavioral trends and future outlook.

For the period starting April 1, 2014 and ending June 30, 2014, Verisign observed the following key trends:

- 65% of attacks were greater than 1 Gbps.
- Verisign has seen a marked increase in volumetric DDoS activity in the second quarter of 2014, with attacks reaching peak of approximately 300 Gbps/24 Mpps for UDP floods and about 35 Gbps/91 Mpps for TCP; this constitutes a 291% increase in average attack size over the same quarter last year (Q2 2013: 3.17 Gbps) and a 216% increase over the previous quarter (Q1 2014: 3.92 Gbps)
- Continuing the trend for the first quarter, the largest number of attacks (43%) targeted Media and Entertainment customers, but was nearly mirrored by attacks directed against IT Services / Cloud / SaaS verticals (41%).
- Verisign observed increasing complexity in second-quarter DDoS attacks, including attacks that quickly and unpredictably changed vectors over the course of the mitigation. Verisign saw sophisticated TCP and UDP floods that targeted specific custom application ports and continuously switched vectors.
- The primary attack vector continues to be UDP based NTP reflective attacks generating significant volumetric attack scale against online businesses.
- The attacks Verisign observed in the second quarter were most often very short in duration but high in intensity. In the second quarter, Verisign mitigated multiple attacks in the 200+ Gbps range within a 24-hour period.

VERISIGN

VERISIGN-OBSERVED DDOS TRENDS FROM THE SECOND QUARTER OF 2014

Attack Stats

In Q2 2014, Verisign saw heightened DDoS attack activity against online businesses and applications, including one of the largest DDoS attacks recorded this year against a Verisign customer.

Mitigations by Attack Size

The average attack size increased 216% to 12.42 Gbps (Figure 1) from the previous quarter, which represents a 291% increase over the same quarter last year.

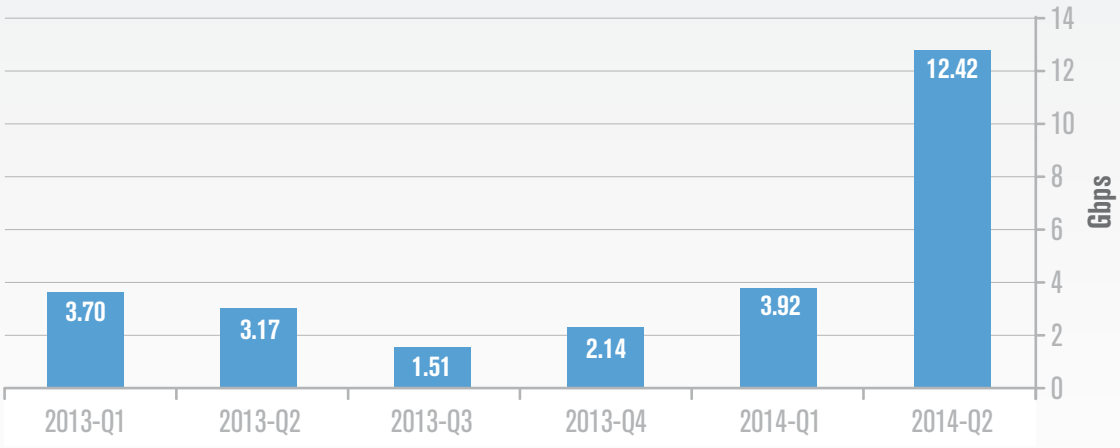

Figure 1: Average Quarterly Peak Attack Size

Volumetric Attacks Increased:

UDP
300 GBPS/
24 MPPS

TCP
35 GBPS/
91 MPPS

Additionally, Verisign saw an increase in volumetric DDoS scale, with attacks reaching 300 Gbps and 24 Mpps at peak for UDP floods and about 35 Gbps and 91 Mpps for TCP. The percentage of attacks in the 5-10 Gbps range increased to 33% and the number of attacks over 10 Gbps increased to 16% compared with 17% and 9% respectively in Q1 2014 (Figure 2).

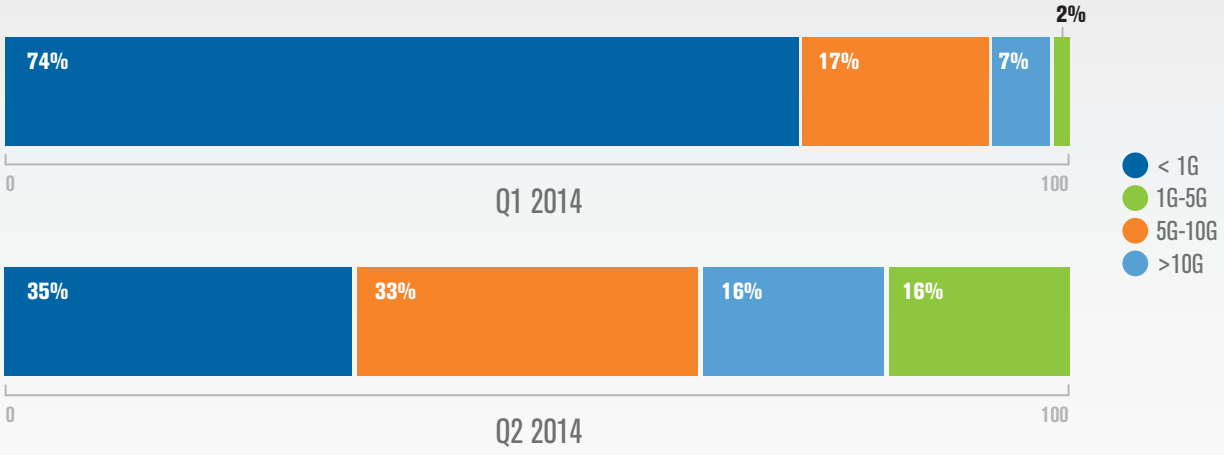

Figure 2: Q1-Q2 2014 Peak Attack Sizes

Two out of three attacks peaked at over 1 Gbps. Considering that most small and medium-sized businesses likely have less than 1 Gbps of upstream bandwidth, attacks like these could be devastating for organizations without adequate protection.

Mitigations by Vertical

In Q2 2014, attackers continue to target verticals beyond Financial Services which were the most-attacked vertical in 2013. Media and Entertainment was the most frequently attacked vertical in Q2 2014, followed by IT Services / Cloud / SaaS (Figure 3).

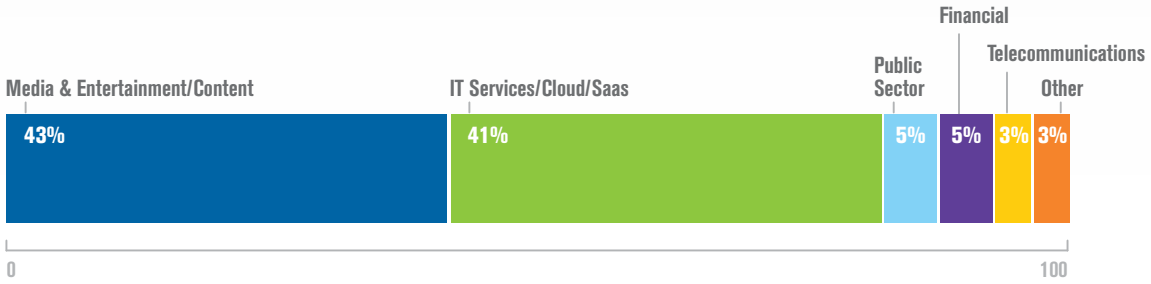

Figure 3: Q2 Mitigations by Vertical

VERISIGN

87
PERCENT
Year-Over-Year
Increase in Peak
Attack Size

The percentage of mitigations in the Media/Entertainment and IT Services/Cloud/SaaS verticals increased 16% compared with Q1 2014. The Media and Entertainment vertical experienced the largest attack, peaking at 300 Gbps (See feature), while the largest Verisign customer attack size in Q2 2013 peaked at just over 160 Gbps and targeted a Financial Services customer, representing a 87% increase in peak attack size year over year.

BEHAVIORAL TRENDS

UDP-Based NTP Reflection Attacks Continue

In Q2 2014, the primary attack vector for UDP floods continues to be large Network Time Protocol (NTP) floods. Verisign observed the emergence of large NTP reflection attacks in December 2013, and these have continued steadily through the second quarter of 2014, despite recent public reports to the contrary.

Many organizations do not use or trust external systems for their NTP, so in this case the solution can be as easy as restricting or rate limiting NTP ports inbound/outbound to only the authenticated/known hosts. The real danger of an NTP attack is volume as a result of feasible amplification vectors it provides. If an organization is the target of a large NTP reflection attack, in a size range above or close to their upstream capacity, a cloud-based mitigation is the most ideal solution. In 2014, Verisign has defended against NTP-based DDoS attacks in the 50-70 Gbps range.

Application-Layer Targeting Slows

While volumetric Layer-3 and Layer-4 attacks were definitely more common, approximately 10% of the attacks against our customers in Q2 targeted the application layer, down from 30% in Q1 2014. Layer-7 attacks target specific Web applications, protocol headers and application parameters and are typically smaller in sheer volume, but require complex mitigation techniques including packet inspection, real-time signature generation, resource management techniques and diverse client capability-handling techniques. Verisign's proprietary Athena mitigation platform and integrated intelligence from Verisign iDefense have been effective in mitigating these complex attacks.

Future Outlook

DDoS attackers are deploying new techniques and leveraging new network and server vulnerabilities every day in pursuit of their goals. Amplification attacks using techniques such as DNS reflection, NTP, SNMP, etc. are becoming more prevalent targeting businesses. Verisign has seen an increase in repeated volumetric attacks in the 200-300 Gbps range – orders of magnitude larger than most enterprises or even many ISPs can survive without extensive collateral damage or service degradation, and the support of a capable cloud-based DDoS mitigation services provider is necessary for preserving services availability and business continuity.

While some organizations have chosen to over-provision bandwidth and other resources, the volumes seen in modern attacks make this needlessly expensive at best and an arms race the defender will often lose at worst. DDoS mitigation providers use specialized systems to find commonalities in malicious packets and scrub them, allowing legitimate traffic to pass.

Verisign has additionally seen attackers continue to adjust their tactics over the course of an attack in a determined effort to disable their victims. It's also clear that attackers are increasingly aware of cloud-based DDoS services and may be structuring attacks in an attempt to circumvent provider infrastructure, as indicated in the case of attacks crafted to stream successfully through GRE tunneling network transport discussed in the following feature case study.

VERISIGN

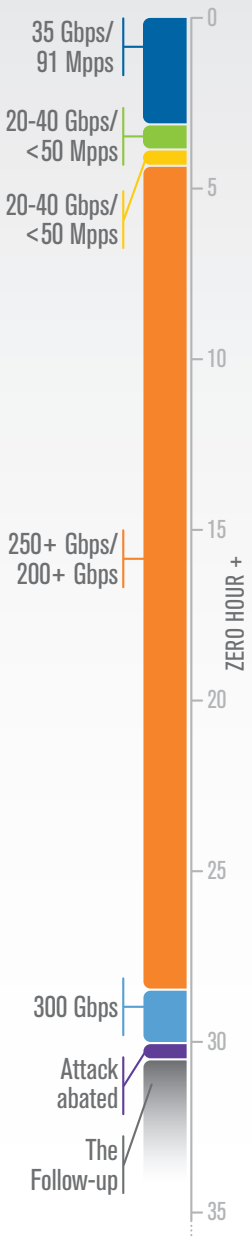

Figure 4: Timeline Showing Attack Progression

Feature:
VERISIGN THWARTS MASSIVE 300 GBPS MULTI-VECTOR DDoS ATTACK AGAINST GLOBAL MEDIA CUSTOMER

Introduction

The largest DDoS event Verisign observed and mitigated in Q2 2014 weighed in at 300 Gbps, and was directed against a Media and Entertainment Services customer.

Verisign believes this attack was notable for a number of reasons: the size (aggregate bandwidth and packets per second) was exceptionally large and multiple vectors were employed, using numerous TCP and UDP transport layer attack vectors. Verisign asserts that such adaptable and well orchestrated DDoS attacks point to a rapidly evolving threat landscape, where asymmetries afforded to adversaries require rapid response and layered security defenses that can accommodate an ever-expanding array of attack vectors.

Attack Timeline

Following is a high-level timeline of the attack, including what the Verisign DDoS Protection Services Team saw and did to mitigate the attack.

- **"Zero Hour" (Z+0 hrs.)**: Verisign begins redirecting victim IP space through its global network of mitigation centers (See Sidebar) and subsequently begins receiving and filtering SYN flood DDoS attack traffic. Victim begins receiving clean traffic and their network is stabilized.
- **"The Second Wave" (Z+3:30 hrs.)**: For the first 3.5 hours, Verisign mitigates periodic SYN & TCP floods with invalid flag combinations. Attack sizes average 20-40 Gbps, and spike as high as 91 Mpps.
- **"The Change-up" (Z+4:00 hrs)**: Perhaps seeing his efforts thwarted, the attacker morphs attack type to a UDP flood using large packets. Verisign adapts mitigation techniques to mitigate the new attack vectors. The attacker constantly mixed up the attack patterns and switched between high packets per second TCP and high bits per second UDP packets.
- **"The Peak" (Z+4:10 hrs.)**: The UDP flood attack reaches 250+ Gbps. Verisign continues to mitigate the attack using its diverse mitigation platforms (including Athena), in conjunction with its global network and capacity. In parallel, Verisign DDoS experts and iDefense intelligence analysts collaborate on understanding attacker tactics, possible attack signatures and bots to refine the ongoing mitigation techniques.
- **"The Siege" (Z+4:11-28:00 hrs.)**: The attacker proves persistent and re-launches attack in the form of 200+ Gbps UDP flood and TCP floods in 5-15-minute waves and bursts – more than 30 in this 24-hour period. The attack is well distributed in nature, originates from several regions of the world and is sourced from multiple compromised computers, or "bots."

- **"The Last Swing" (Z+30:00 hrs.)**: The attacker's UDP flood spikes to 300 Gbps. Again, Verisign's distributed network absorbs the attack while maintaining network availability.
- **"The End" (Z+30:15 hrs.)**: After more than 30 hours, the attacker finally gives up and the attack subsides.
- **"The Follow-Up" (Z+30:15- TBD)** Verisign continues to monitor the customer network for several days to ensure the attack is indeed over, and then returns customer traffic to its normal flow and operation. Verisign performs a post-attack review with the customer and identifies any opportunities for improvement on either side.

The Attack

The initial attacks were TCP SYN and NULL Floods consisting of packets with invalid flag combinations (in some cases no flags defined) that grew to 35 Gbps in size and 91 Mpps in volume. Once this attack vector was mitigated, the attack type morphed into a UDP flood that grew to a peak of approximately 300 Gbps and 24 Mpps.

The attack came in multiple waves that were very short in duration but high in intensity. In one 24- hour period Verisign recorded more than 30 attacks in the 200+ Gbps range (See Figure 5).

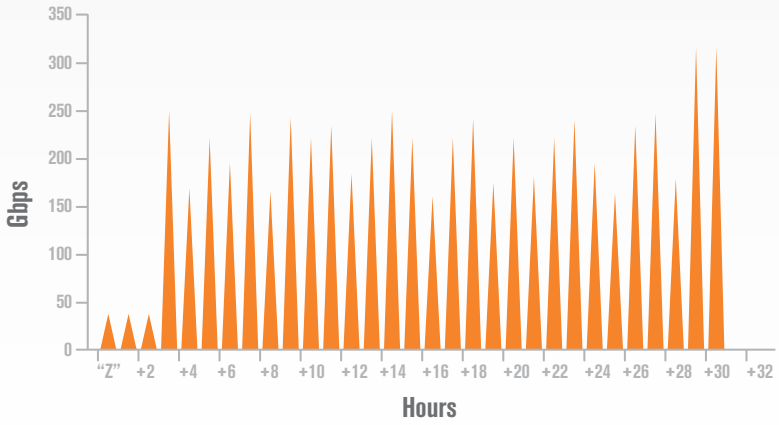

Figure 5: Attack Peaks over Time

THE IMPORTANCE OF A STRONG NETWORK IN COMBATING DDoS ATTACKS

Verisign's DDoS Mitigation Service is strategically engineered to handle both nominal traffic loads in addition to large traffic spikes that occur under DDoS attack conditions. The flexible MPLS configuration of Verisign's network allows Verisign's DDoS engineers to selectively route traffic flows in response to global DDoS attack dynamics ensuring that no single component is overwhelmed.

For more on how Verisign's network makes our DDoS mitigation service faster and easier, click **here** to download our white paper.

One interesting aspect of the UDP portion of the attack was that it targeted a custom, application-specific port. The attackers crafted large malformed UDP payloads to generate the attack size. Another notable behavior around the UDP attack was that the packets were sized (intentionally or not) to fit inside GRE tunnels. This may indicate that the attacker was aware that the victim was being protected by a cloud-based mitigation service using GRE tunneling. By crafting packets that fit within GRE tunneling, the attack traffic would continue to the victim's network had filtering proved ineffective.

The botnets used in the attack seem to have been created by exploiting servers vulnerable to the "supermicro ipmi" exploit (See FAQ below). Although the actual number of compromised systems has not been identified, Verisign iDefense Security Intelligence Services analysts believe that up to 100,000 servers could be vulnerable to this exploit.

Over the course of this particular attack, Verisign was able to dynamically update the blacklists used during the mitigation to help neutralize the event and push the mitigation out to the edge of Verisign's globally connected network backbone. Verisign's custom build proprietary mitigation platform, Athena, can support the use of blacklists in excess of 1 million IP addresses.

Tactics, Techniques, and Procedures (TTPs)

Verisign observed two types of DDoS attacks on a content delivery network (CDN) for the customer site. TCP SYN floods against an application-specific port reached peaks around 90 Mpps. UDP floods (also against the same port) instead used volumetric techniques and reached a peak of 300 Gbps. Though the lack of anti-spoofing measures at the network layer and either new connection setup handshakes for stateful protocols, or connectionless protocols such as UDP are effective spoof-based or reflective attack vectors, one address appeared far more frequently than any other, and is assigned to a European ISP. Past social media network communications from the alleged attacker indicate that he may have used compromised FreeBSD servers or routers from other application-related service providers.

FAQ: THE "SUPERMICRO IPMI" VULNERABILITY

Q: What is it?
A: Supermicro motherboards store administrator passwords in plain text, which is available to any attacker who can connect to port 49152.

Q: How Does it Work?
A: The "/nv" directory contains the BMC password file "IPMIdevicedesc.xml," which is accessible from any UPnP-enabled Supermicro motherboard running IPMI on a public interface.

Additionally, all the contents of the "/nv/" directory, including the server.pemfile, the wsman administrator password, and the netconfig files, are accessible via the browser.

Q: What is the Threat?
A: Exploiting this vulnerability can allow an attacker to remotely access password files and steal or expose the victim's passwords.

Wikholm, Zachary. "CARISIRT: Yet Another BMC Vulnerability (And some added extras)." CARI.net; http://blog.cari.net/carisirt-yet-another-bmc-vulnerability-and-some-added-extras/. Accessed 7/28/14

VERISIGN

Adversary Motivations

The motivations for the attacks appear to stem from recent controversies around server monetization and potentially other disputes around moderation on independent forums. Attacks like this could constitute online activism, or "hacktivism," because they seem to have been motivated by ideology. DDoS attacks have become one of the two main weapons of choice (along with SQL injection) when it comes to cyber protest and hacktivism.

The Vulnerability

On June 19, 2014, Zachary Wikholm, a researcher from hosting provider CARI, released an advisory disclosing details about a vulnerability within unpatched Universal Plug-and-Play (UPnP)-enabled Super Micro Computer Inc. (or Supermicro) microprocessor boards.

The vulnerability is due to the availability of an unencrypted password file, "PSBlock," for download via port 49152, without any access restrictions, which allows any attacker to simply connect to port 49152 and gain access to the unencrypted password file. Wikholm reportedly disclosed this flaw to Supermicro. In its response, Supermicro informed the researcher that it had fixed this flaw and issued new firmware to patch this vulnerability in 2013. An attacker can connect to any one of them and gain access to the file containing the passwords. This access will then enable the attacker to compromise the management interface of the victim server's motherboard.

Although Supermicro had released a new firmware that patches this vulnerability, in the real world, users may be unwilling to shut down their servers to install the new firmware, as it could constitute a loss of revenue.

Unfortunately, this scenario plays out all too often, and to the detriment of all but the attackers. The lifecycle of a vulnerability does not immediately come to an end once the vendor releases a patch. It takes time for users to install patches, and as long as there remain unpatched applications in the wild, the exploitation lifecycle of the relevant vulnerability will continue.

Source: ibid.

VERISIGN

CONCLUSION

Verisign believes this attack is notable not only for its size, but also for its sophistication, and believes this type of multi-vector, adaptive and well-orchestrated attack is indicative of the continued evolution of the DDoS threat. Increasingly prepared attackers using packets crafted to the size of GRE tunnels, and targeting specific ports may point toward more resilient, unpredictable and expedient attacks in the future, which will necessitate that protection providers increase their expertise, intelligence and technology innovation to keep up with, if not stay ahead of, the DDoS threat.

VERISIGN

NOTES