

An in-depth perspective on software vulnerabilities and exploits, malware, potentially unwanted software, and malicious websites

Microsoft Security Intelligence Report

Volume 15

January through June, 2013

Cloud security:
Conflict and cooperation

Microsoft Security Intelligence Report

This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

This document is provided "as-is." Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it.

Copyright © 2013 Microsoft Corporation. All rights reserved.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Authors

Dennis Batchelder <i>Microsoft Malware Protection Center (MMPC)</i>	Aneesh Kulkarni <i>Windows Services Safety Platform</i>	Anthony Penta <i>Windows Services Safety Platform</i>
Joe Blackbird <i>MMPC</i>	Marc Lauricella <i>Microsoft Trustworthy Computing</i>	Tim Rains <i>Microsoft Trustworthy Computing</i>
David Felstead <i>Bing</i>	Russ McRee <i>Online Services Security & Compliance</i>	Vidya Sekhar <i>MMPC</i>
Paul Henry <i>Wadeware LLC</i>	Chad Mills <i>Windows Services Safety Platform</i>	Holly Stewart <i>MMPC</i>
Ben Hope <i>MMPC</i>	Nam Ng <i>Microsoft Trustworthy Computing</i>	Matt Thomlinson <i>Microsoft Trustworthy Computing</i>
Jeff Jones <i>Microsoft Trustworthy Computing</i>	Daryl Pecelj <i>Microsoft IT Information Security and Risk Management</i>	Todd Thompson <i>Microsoft IT Information Security and Risk Management</i>
		Terry Zink <i>Microsoft Exchange Online Protection</i>

Contributors

Danielle Alyias <i>Microsoft Trustworthy Computing</i>	Satomi Hayakawa <i>CSS Japan Security Response Team</i>	Bill Pfeifer <i>MMPC</i>
Joe Faulhaber <i>MMPC</i>	Aaron Hulett <i>MMPC</i>	Cynthia Sandvick <i>Microsoft Trustworthy Computing</i>
Methuselah Cebrian Ferrer <i>MMPC</i>	Jimmy Kuo <i>MMPC</i>	Richard Saunders <i>Microsoft Trustworthy Computing</i>
Peter Ferrie <i>MMPC</i>	Hilda Larina Ragragio <i>MMPC</i>	Jasmine Sesso <i>MMPC</i>
Tanmay Ganacharya <i>MMPC</i>	Jenn LeMond <i>Microsoft IT Information Security and Risk Management</i>	Frank Simorjay <i>Microsoft Trustworthy Computing</i>
Kathryn Gillespie <i>Microsoft IT Information Security and Risk Management</i>	Ken Malcolmson <i>Microsoft Trustworthy Computing</i>	Francis Tan Seng <i>MMPC</i>
Enrique Gonzalez <i>MMPC</i>	Marianne Mallen <i>MMPC</i>	Henk van Roest <i>CSS Security EMEA</i>
Jonathan Green <i>MMPC</i>	Scott Molenkamp <i>MMPC</i>	Steve Wacker <i>Wadeware LLC</i>
Angela Gunn <i>Microsoft Trustworthy Computing</i>	Daric Morton <i>Microsoft Services</i>	Shawn Wang <i>MMPC</i>
Joe Gura <i>Microsoft Trustworthy Computing</i>	Yurika Muraki <i>CSS Japan Security Response Team</i>	Bob White <i>Microsoft IT Information Security and Risk Management</i>
Chris Hale <i>Microsoft Trustworthy Computing</i>	Takumi Onodera <i>Microsoft Premier Field Engineering, Japan</i>	Iaan Wiltshire <i>MMPC</i>
		Dan Wolff <i>MMPC</i>

Table of contents

About this report	iv
Trustworthy Computing: Security engineering at Microsoft	v
Cloud security: Conflict and cooperation	1
Domain Name System (DNS) attacks	3
Distributed Denial of Service (DDoS) attacks.....	9
Guidance: Preventing and mitigating DNS and DDoS attacks	11

About this report

The *Microsoft Security Intelligence Report (SIR)* focuses on software vulnerabilities, software vulnerability exploits, and malicious and potentially unwanted software. Past reports and related resources are available for download at www.microsoft.com/sir. We hope that readers find the data, insights, and guidance provided in this report useful in helping them protect their organizations, software, and users.

Reporting period

This volume of the *Microsoft Security Intelligence Report* focuses on the first and second quarters of 2013, with trend data for the last several quarters presented on a quarterly basis. Because vulnerability disclosures can be highly inconsistent from quarter to quarter and often occur disproportionately at certain times of the year, statistics about vulnerability disclosures are presented on a half-yearly basis.

Throughout the report, half-yearly and quarterly time periods are referenced using the *nHyy* or *nQyy* formats, in which *yy* indicates the calendar year and *n* indicates the half or quarter. For example, 1H13 represents the first half of 2013 (January 1 through June 30), and 4Q12 represents the fourth quarter of 2012 (October 1 through December 31). To avoid confusion, please note the reporting period or periods being referenced when considering the statistics in this report.

Conventions

This report uses the Microsoft Malware Protection Center (MMPC) naming standard for families and variants of malware and potentially unwanted software. For information about this standard, see **"Error! Reference source not found."** on page **Error! Bookmark not defined.** In this report, any threat or group of threats that share a common unique base name is considered a family for the sake of presentation. This consideration includes threats that may not otherwise be considered families according to common industry practices, such as adware programs and generic detections. For the purposes of this report, a "threat" is defined as a malware or potentially unwanted software family or variant that is detected by the Microsoft Malware Protection Engine.

Trustworthy Computing: Security engineering at Microsoft

Amid the increasing complexity of today's computing threat landscape and the growing sophistication of criminal attacks, enterprise organizations and governments are more focused than ever on protecting their computing environments so that they and their constituents are safer online. With more than a billion systems using its products and services worldwide, Microsoft collaborates with partners, industry, and governments to help create a safer, more trusted Internet.

The Microsoft Trustworthy Computing organization focuses on creating and delivering secure, private, and reliable computing experiences based on sound business practices. Most of the intelligence provided in this report comes from Trustworthy Computing security centers—the Microsoft Malware Protection Center (MMPC), Microsoft Security Response Center (MSRC), and Microsoft Security Engineering Center (MSEC)—which deliver in-depth threat intelligence, threat response, and security science. Additional information comes from product groups across Microsoft and from Microsoft IT, the group that manages global IT services for Microsoft. The report is designed to give Microsoft customers, partners, and the software industry a well-rounded understanding of the threat landscape so that they will be in a better position to protect themselves and their assets from criminal activity.

Cloud security: Conflict and cooperation

As one of the largest and fastest growing operators of cloud services in the world, Microsoft makes cloud security a top priority. Incidents are handled by multiple teams throughout the company, and many business groups have their own incident response teams with specific focus areas and authority. Despite this decentralized structure, all Microsoft cloud incident response teams face certain intrinsic challenges. For example, the infrastructure required to serve hundreds of millions of customer accounts on every continent generates an astronomical amount of data in the form of logs, alerts, and other telemetry. Over the course of one recent month, the domain controller logs for servers that manage primary Microsoft production environment domains generated 57.1 billion Windows security events. Add in network data (including NetFlow telemetry), firewall events, and intrusion prevention system (IPS) events, and event counts easily reach the trillions. And that's primarily from non-virtual systems!

Even at this scale, the Microsoft cloud infrastructure faces many of the same security challenges and attack patterns that affect much smaller computing environments. The scale may be vastly different, but many of the challenges that Microsoft cloud services administrators and security response teams face are similar or identical in nature to issues faced by every IT administrator reading this report. For example, administrators who manage monthly security updates from Microsoft might find it interesting to consider that the Microsoft cloud team deploys the same set of updates to a server base numbering in the hundreds of thousands. Automation plays an invaluable role, but system administration in massive, distributed cloud infrastructures is still a significant undertaking.

Similarly, some of the high-profile attack vectors that have been deeply problematic for system administrators around the world in recent times have not gone unnoticed by Microsoft cloud security teams. This section of the *Microsoft Security Intelligence Report* examines two of these attack vectors from the perspective of Microsoft cloud services and incident response teams.

Domain Name System (DNS) attacks

Attacks on the global Domain Name System (DNS) are some of the most serious and potentially damaging attacks affecting the Internet today. A group of malicious hackers calling itself the "Syrian Electronic Army" made headlines in mid-2013 when it successfully compromised a registrar that manages DNS

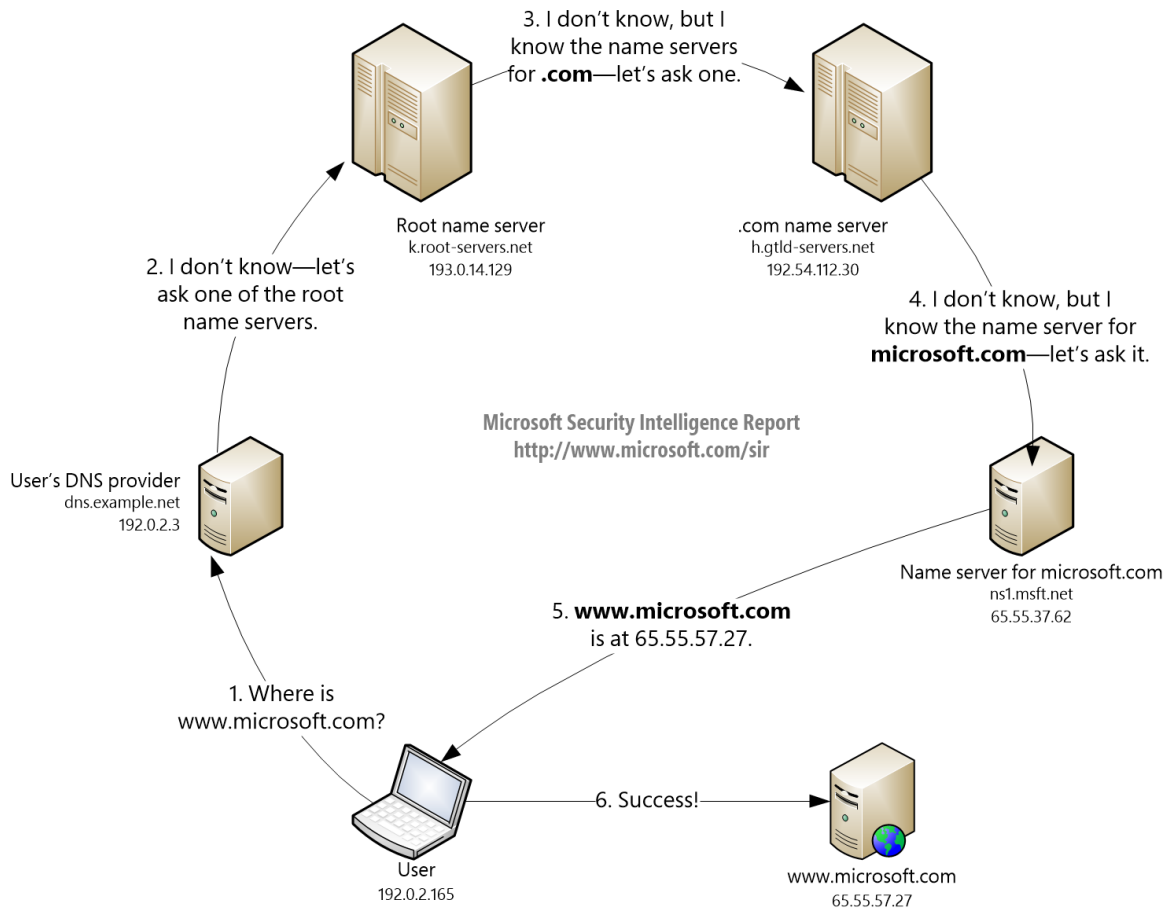
records for *The New York Times* and Twitter.¹ Over the last few years, Microsoft has experienced similar attacks, some of which were politically motivated, against registries managing its DNS records in specific markets. This malicious manipulation of DNS records has an adverse impact not only on Microsoft but on the global online community as well, including Microsoft industry peers, partners, and customers.

When a computer user requests a domain-based URL from a web browser, the computer usually must query at least one DNS name server to resolve the alphanumeric domain string into an IP address that can be used to locate and retrieve the desired web page. In a typical case, visiting a URL such as *www.microsoft.com* might require querying at least four different name servers:²

¹ Timothy B. Lee, "The New York Times Web site was taken down by DNS hijacking. Here's what that means," *The Washington Post*, August 27, 2013, www.washingtonpost.com/blogs/the-switch/wp/2013/08/27/the-new-york-times-web-site-was-taken-down-by-dns-hijacking-heres-what-that-means/.

² In practice, techniques such as DNS caching and hosts file lookups usually eliminate one or more of these steps for most queries.

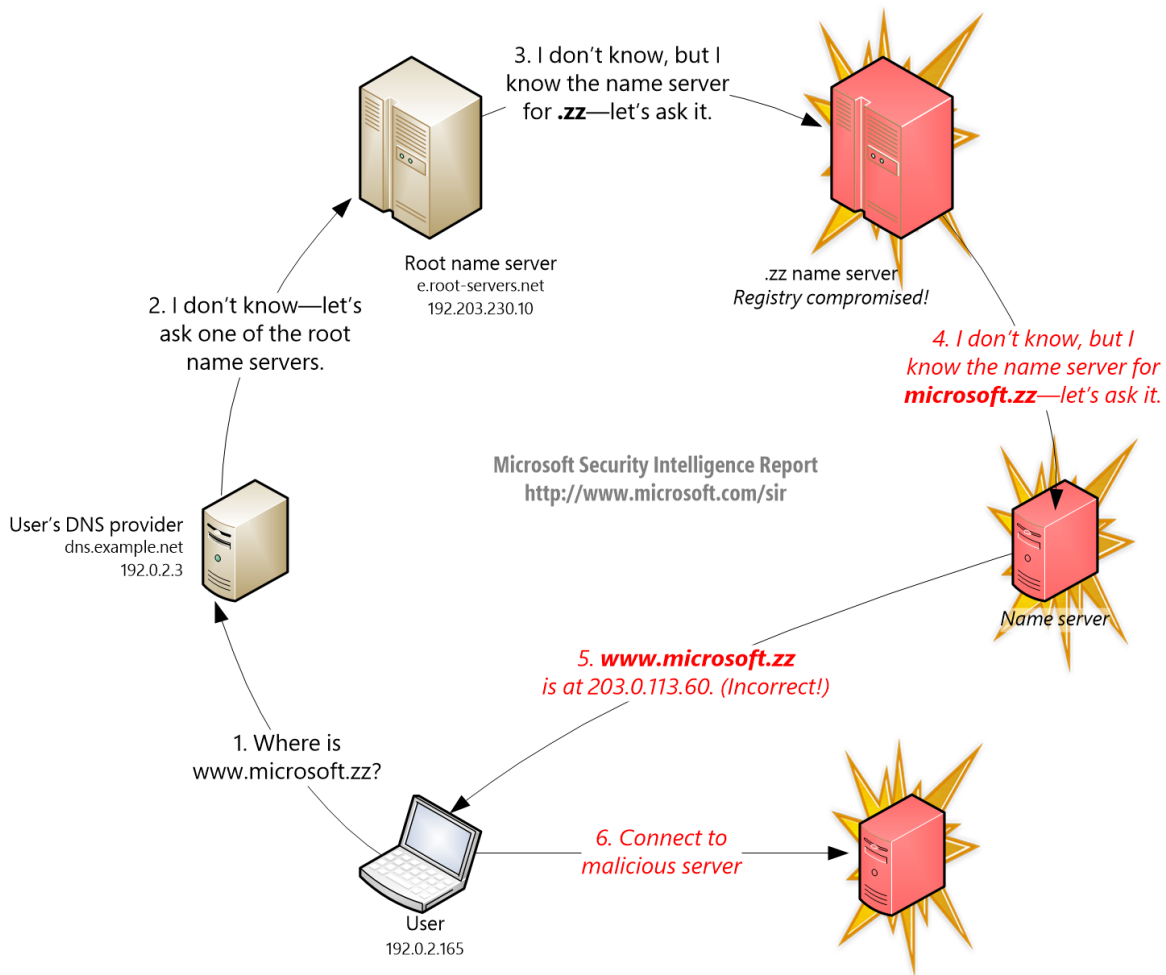
Figure 1. A simplified diagram of the DNS address resolution process



1. The computer queries the recursive DNS server for the network connection being used. A recursive DNS server handles DNS queries for its clients by locating and querying other DNS servers (called authoritative name servers), which are designated to provide authoritative address lookups for specific individual domains.
2. If the recursive name server doesn't have the answer, it queries one of the 13 root name servers (which correspond to hundreds of physical servers located around the world).
3. The root name server maintains a record of the authoritative name servers for the `.com` top-level domain (TLD) and queries one of them.
4. The `.com` name server maintains a record of the authoritative name server for the `microsoft.com` domain, and queries that server.
5. The `microsoft.com` name server maintains a record of the IP address for the `www` subdomain, and returns the IP address.

If attackers successfully compromise one of the name servers or registries in this chain, they can redirect DNS queries to a malicious name server. For example, a compromise of the authoritative name server for microsoft.com could result in requests for *www.microsoft.com* being redirected to an IP address of the attacker's choosing, which may serve malware or contain a maliciously altered version of the Microsoft website. The potential for greater damage increases as one travels up the DNS hierarchy; a hypothetical compromise of one of the root name servers could conceivably put every domain on the Internet in jeopardy.

Figure 2. A compromised registry can result in malicious responses being issued to DNS queries

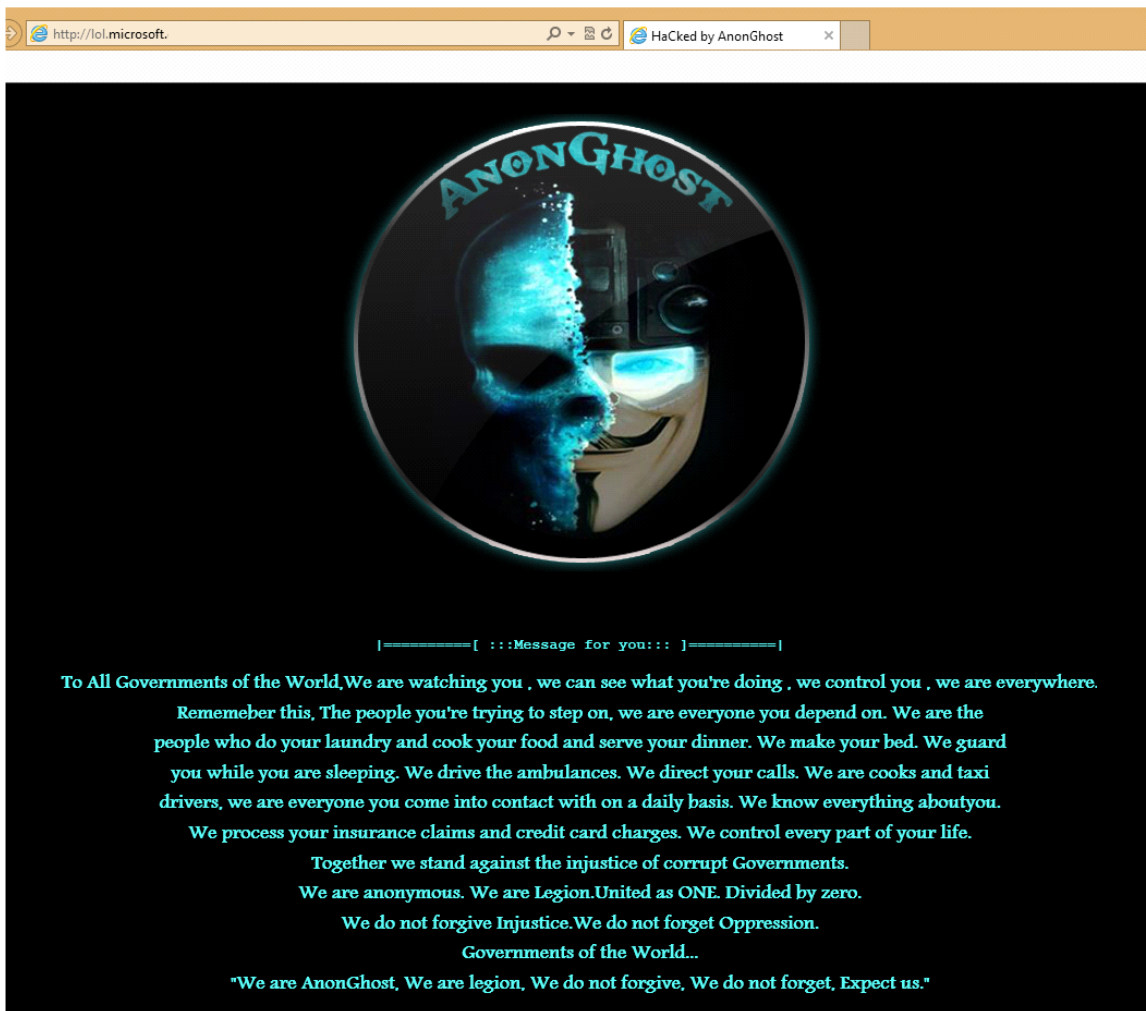


The exploitation of vulnerabilities that are specific to country-code top-level domain (ccTLD) registries has become increasingly common, especially in relatively small markets. A ccTLD is a top-level domain that is generally used or reserved for a country or region, such as *.ca* for Canada. There are currently

more than 300 ccTLD name registries responsible for servicing hundreds of millions of domain names worldwide. Domains registered under ccTLDs are typically websites or other resources that cater to the associated country or region for those who wish a web presence in their country of origin, or for companies that seek to grow their presence and market share in such countries. For example, Microsoft maintains registered domains under a number of different ccTLDs for its regional subsidiaries, such as microsoft.ca for Microsoft Canada and microsoft.co.jp for Microsoft Japan. Domains that are registered under ccTLDs help create positive Internet experiences for users in different communities by providing locally targeted resources at familiar and predictable domain names. Unfortunately, the name servers run by some ccTLD registrars are vulnerable to attack, which can negatively affect individuals, nonprofits, and government organizations as well small companies and large corporations such as Microsoft. Between May 2012 and July 2013, 17 ccTLDs that manage DNS records for Microsoft (and many other organizations) in specific countries and regions were compromised, often through a combination of Structured Query Language (SQL) injection exploits and social engineering.

When computer users attempt to reach a website whose DNS record has been hijacked, they are typically redirected to a server controlled by an attacker. This server may contain web browser exploit kits or malware, or may display malicious or inappropriate content. For example, in May 2013 a group of malicious hackers calling itself "AnonGhost" redirected queries for a Microsoft regional website to a server it controlled, as shown in Figure 3.

Figure 3. The appearance of a website defacement resulting from a compromised DNS record



To the computer user it appears as though the website itself has been compromised, even though the owner of the targeted website usually has no control over the ccTLD and is not responsible for the incident. Users typically can't differentiate between a problem with the ccTLD or the organization that runs the website they wish to browse, and even advanced users may have considerable trouble distinguishing between a website problem and a DNS problem. This type of DNS hijacking diminishes public confidence in the victimized organizations and adversely affects their reputations.

Although security best practices, reviews, training, and awareness can help prevent these types of attacks, the frequency and impact of such attacks have prompted Microsoft to offer help to registries. Microsoft now offers the ccTLD Registry Security Assessment Service, which helps registry operators find and fix

vulnerabilities at no charge before they are exploited.³ Microsoft believes that close collaboration in this effort between industry peers, partners, and industry groups such as ICANN can help increase awareness for ccTLDs and reduce the unfortunate impact of DNS records manipulation.


Distributed Denial of Service (DDoS) attacks

Another common attack vector that has been used to attempt to adversely affect cloud and online services at Microsoft is Distributed Denial of Service (DDoS), including attacks that result from *DNS amplification* (a technique that involves using publicly accessible open DNS servers to flood the target system with DNS traffic). DNS amplification made headlines in March 2013, when attackers used the technique to attack the Spamhaus spam prevention service with as much as 300 gigabits per second (Gbps) of traffic.⁴

On a daily basis, Microsoft's DDoS protective measures apply mitigations to prevent impact from DoS and DDoS attacks to ensure uptime and availability for services and customers. Common types of attack include SYN floods, DNS amplification, malformed packets (TCP and UDP), and application layer abuses specific to HTTP and DNS. One common attack technique used by a number of freely available DDoS toolkits involves using fragmented IP packets with a fixed payload, as described below.

A DDoS attack in progress quickly shows up on monitoring telemetry as a significant elevation of both packets-per-second and bits-per-second traffic, as seen in Figure 4. The 30Mbps attack shown here is nominal, but if left unchecked could impact the availability of the service.

Figure 4. Flow monitoring telemetry during a DDoS attack



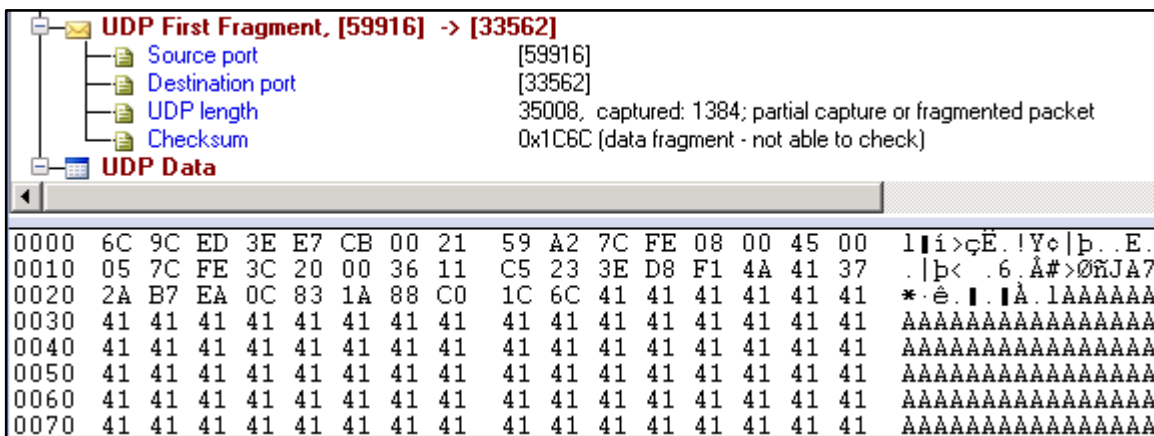
³ For more information, see the entry "[Microsoft Offers Security Assessment Service for Country-Code Top-Level Domain Registries \(ccTLD\)](#)" (February 26, 2013) on the Microsoft Security Blog at blogs.technet.com/security.

⁴ Michael McNally, "What is a DNS Amplification Attack?", *ISC Knowledge Base*, April 1, 2013, <https://deephought.isc.org/article/AA-00897/0/What-is-a-DNS-Amplification-Attack.html>.

A typical attack involving IP fragments might consist of a padded payload consisting of a single ASCII letter, such as A (0x41 in hexadecimal), repeated many times, and transmitted using multiple communications protocols, including User Datagram Protocol (UDP), Transmission Control Protocol (TCP), Internet Control Message Protocol (ICMP), KRYPTOLAN, Versatile Message Transaction Protocol (VMTP), Internet Protocol version 6 (IPv6), Extensible Name Service (XNS), and others. Packets often include full 1,518-byte payloads, and the UDP fragments are directed to multiple destination ports.

Figure 5 represents a UDP fragment that was captured during an attack.

Figure 5. A UDP fragment from a DDoS attack



During one 60-second window, Microsoft detected more than 8,985 unique IP addresses sending fragmented traffic during the attack. As the service was forced to drop incoming packets during the attacks, it is believed that the actual volume of the attack may have been considerably greater than what Microsoft was able to analyze.

An investigation of a host known to have participated in a recent attack, acquired via appropriate legal means by the Microsoft Digital Crimes Unit (DCU), revealed a common attack tool (currently detected as [Backdoor:Perl/IRCbot.E](#)) that was used for UDP flooding.

Figure 6. Perl code from a UDP flooding trojan

```
#####
if ($funcarg =~ /^udp2\s+(.*)\s+(\d+)\s+(\d+)/) {
    sendraw($IRC_cur_socket, "PRIVMSG $print1 :!4,1 [UDP-2 DDOS]! '9,1Attacking '12'.$1." '9,1with '12'.$2." '9,1Kb Packets for '12'.$3."
    '9,1seconds.' '1');
    my ($dtime, $pacotes) = udpflooder("$1", "$2", "$3");
    $dtime = 1 if $dtime == 0;
    my $bytes;
    $bytes(igmp) = $2 * $pacotes(igmp);
    $bytes(icmp) = $2 * $pacotes(icmp);
    $bytes(o) = $2 * $pacotes(o);
    $bytes(udp) = $2 * $pacotes(udp);
    $bytes(tcp) = $2 * $pacotes(tcp);
    sendraw($IRC_cur_socket, "PRIVMSG $print1 :!4,1 [UDP-2 DDOS]! '9,1Results '12'.int(($bytes(icmp)+$bytes(igmp)+$bytes(udp) + $bytes(o)
    /1024).' '9,1Kb in '12'.$dtime." '9,1seconds to '12'.$1.'" '9,1.' '1');
}
}
```

Tools such as this IRCbot provide even the most unsophisticated attackers a platform from which to launch potentially damaging attacks on cloud services. Although the defensive measures and tactics employed by Microsoft help mitigate such attacks, it can nonetheless be burdensome and resource intensive to do so.

Guidance: Preventing and mitigating DNS and DDoS attacks

For owners of websites in vulnerable ccTLDs, preventing DNS attacks at the TLD level can be very difficult or impossible. Website owners should urge their ccTLD registrars to visit www.microsoft.com/cctldregsec and take advantage of the Microsoft ccTLD Registry Security Assessment Service to find and mitigate any vulnerabilities that may leave domains open to attack.

Because attackers also target individual domains for DNS hijacking directly, website owners should act to ensure that their designated authoritative name servers cannot be changed without their approval. Many domain name registrars offer domain locking services that can help prevent DNS records from being changed without the domain owner's approval. Website owners should take advantage of any locking services offered by their registrars, and should urge registrars to offer such services if they do not. Site owners should also take general precautions to secure their domain names against unauthorized changes, such as carefully protecting the usernames and passwords they use to access their domain registry accounts, and only using SSL connections to review their accounts or make changes.

Because DDoS attacks are so difficult to mitigate, it's important that DNS administrators everywhere be willing to cooperate with each other to prevent attacks from happening in the first place. The United States Computer Emergency Readiness Team (US-CERT) has provided some suggestions to help

administrators stop attackers from taking advantage of their DNS servers to launch attacks.⁵

- Most DNS amplification attacks take advantage of open DNS name servers, which resolve DNS queries submitted to them by any computer on the Internet. System administrators should configure their DNS servers to ignore queries they receive from hosts outside their domain. A number of tools are available for helping administrators detect misconfigured DNS servers within their networks, including:
 - The Open Resolver Project (openresolverproject.org) maintains a list of open DNS resolvers and provides an interface for searching an IP range for open resolvers.
 - The Measurement Factory (dns.measurement-factory.com) also maintains a list of open resolvers and offers a free tool to test a single server to determine if it allows open recursion.
 - DNSInspect (dnsinspect.com) is another free tool for testing DNS resolvers, and it can also test an entire DNS zone for other possible configuration and security issues.
- Administrators of DNS resolvers can take a number of steps to prevent their resources from being used in attacks, including:
 - *Source IP verification.* Even well-configured DNS resolvers can be exploited by attackers who use source IP address spoofing to issue DNS queries. The Internet Engineering Task Force has released two Best Current Practice documents (tools.ietf.org/html/bcp38, tools.ietf.org/html/bcp84) that can help system administrators perform network ingress filtering, which rejects packets that appear to originate from addresses that cannot be reached via the paths the packets actually take.
 - *Disabling recursion on authoritative name servers.* An authoritative name server is one that provides public name resolution for a specified domain (such as *microsoft.com*) and optionally one or more subdomains (such as *www.microsoft.com*). Because authoritative name servers must be publicly accessible, they should be configured to reject recursive queries from clients. For help disabling recursion in Windows Server, see

⁵ See <https://www.us-cert.gov/ncas/alerts/TA13-088A> for the full alert from US-CERT.

[“Disable Recursion on the DNS Server”](#) at Microsoft Technet (technet.microsoft.com).

- *Limiting recursion to authorized clients.* DNS servers that are deployed within an organization or Internet service provider (ISP) should be configured to perform recursive queries on behalf of authorized clients only, preferably restricted to clients within the organization’s network.

Although attacks on popular cloud services tend to make the most headlines, DDoS attacks can—and do—happen to anyone. In fact, well-run cloud services tend to be much better prepared to deal with DDoS attacks than most enterprise IT infrastructures, because successfully overwhelming a large cloud service requires a level of coordination that few prospective attackers are likely to achieve. Organizations that have struggled with DDoS attacks on their websites or other vital parts of their network infrastructures should consider moving some resources to the cloud to take advantage of the security and operations benefits that cloud services provide.



One Microsoft Way
Redmond, WA 98052-6399
microsoft.com/security