

DNS Cache Poisoning (DCP)

Introduction

The old problem of DNS cache poisoning has again reared its ugly head. While some would argue that the domain name system protocol is inherently vulnerable to this style of attack due to the weakness of 16-bit transaction IDs, we cannot ignore the immediate threat while waiting for something better to come along. There are new attacks, which make DNS cache poisoning trivial to execute against a large number of nameservers running today. The purpose of this article is to shed light on these new attacks and recommend ways to defend against them.

A Brief History of Cache Poisoning

In 1993, Christoph Schuba released a paper entitled "Addressing Weaknesses in the Domain Name System Protocol". In it, he outlined several vulnerabilities, including the technique of DNS cache poisoning. In the earliest incarnation, it was possible to provide extra information in a DNS reply packet that would be cached by the daemon. This allowed an attacker to inject false information into the DNS cache for a network, allowing them to perform man-in-the-middle attacks or other mayhem.

In 1997, CERT released advisory CA-1997-22, describing a vulnerability in BIND, the Berkeley Internet Name Domain software which is used by nearly all of the nameservers on the Internet. This time a very basic principle was finally realized: BIND did not randomize its transaction IDs - they were purely sequential. Aside from layer 3 and 4 protocol checking (source and destination IP addresses and ports must match), the transaction ID is the sole form of authentication for a DNS reply. Because an attacker could easily predict the next transaction ID after making their own request, a cache poisoning attack could be carried out using a spoofed query followed by a spoofed answer. To solve this, all new versions of BIND were updated to use randomized transaction IDs.

In 2002, Vagner Sacramento released an advisory showing another problem with BIND's implementation of the DNS protocol. He found that BIND would send multiple simultaneous recursive queries for the same IP address. Because of this a mathematical phenomenon comes into play known as the "Birthday Paradox". This causes the probability of a successful attack to rise to near 100% with only a few hundred packets instead of the tens of thousands previously believed to be needed.

While researching Sacramento's findings, the CERT team also realized there might be another attack possible, based on the work of Michal Zalewski in the area of TCP sequence numbers and phase space analysis of the pseudo-random number generators used by different operating systems to generate them. Zalewski found that using a certain type of analysis it was often trivial to guess the next sequence number in certain implementations. The CERT team felt that this might also apply to the random number generators in BIND. This article will attempt to show that their assumption was correct.

Attack #1 - The Birthday Attack

To perform this attack, one needs to send a sufficient number of queries to a vulnerable nameserver, while sending an equal number of phony replies at the same time. Because the flaw in the BIND software generates multiple queries for the same domain name at the same time, one encounters statistically improved odds of hitting the exact transaction ID. This is the classic "Birthday Attack", which is derived from the "Birthday Paradox", described below:

A birthday attack is a name used to refer to a class of brute-force attacks. It gets its name from the surprising result that the probability that two or more people in a group of 23 share the same birthday is greater than 1/2; such a result is called a birthday paradox.

If some function, when supplied with a random input, returns one of k equally-likely values, then by repeatedly evaluating the function for different inputs, we expect to obtain the same output after about $1.2k^{1/2}$. For the above birthday paradox, replace k with 365. (unknown author, <http://www.x5.net/faqs/crypto/q95.html>)

We can apply the same methodology to pseudo-random number sequences, such as the one that generates transaction IDs in BIND.

With conventional spoofing, we would send n spoofed replies for one query. Our probability of success is $n / 65535$. With the BIND birthday attack, we send n number of spoofed replies for n queries. For this, we can predict the probability of success using the formula below where t is the total number of possible values in the master set, and n is the number of values in the spoofing subset.

The power of this method of spoofing versus conventional DNS spoofing is shown in illustration 2. The birthday attack nears 100% success around 700 packets. At this point the conventional spoofing attack would only have a success probability of 700 divided by 65535 (1.07%) The steepness of the curve is such that one needs only 300 packets to achieve a 50% success ratio. This is well within the realm of a trivial attack to anyone with a broadband Internet connection.

This shows that even a perfect random-number generator is vulnerable to attack when it generates multiple numbers for the same transaction. This is something that should be taken into account by any software designer who is working with random numbers. It has long been used in brute-force attacks on one-way hash systems, as described by Bruce Schneier in his book Applied Cryptography.

The BIND birthday attack would follow the sequence shown in illustration 3. The attacker merely needs to send a few hundred queries to the ISP nameserver asking for the IP address of the domain name to be hijacked. At the same time, he will send the same number of replies formulated to look as if they were sent from the authoritative nameserver. In each packet he will assign a random transaction ID. In order to be successful, one of his spoofed packets transaction IDs, source and destination IP addresses and ports must match a legitimate recursive query packet from the victim nameserver.

Finding the correct IP addresses is easy; we know our target, and we know the addresses of the legitimate nameservers for the domain to be hijacked. Finding the port is slightly harder. We know that the destination port of the recursive query is UDP port 53, but the source port is a moving target. Fortunately for our attacker, BIND will more often than not reuse the same source port for queries on behalf of the same client. So, if the attacker is working from an authoritative nameserver, he can first issue a request for a DNS lookup of a hostname on his server. When the recursive query packet arrives, he can look at the source port. Chances are this will be the same source port used when the victim sends the queries for the domain to be hijacked. Look at the tcpdump output of four subsequent queries for different domain names:

Code:

```
10:54:12.423228 192.168.1.2.33748 > 66.218.71.63.53:
21345 [1au] A? http://www.yahoo.com. (42) (DF)
10:54:21.313293 192.168.1.2.33748 > 216.239.38.10.53:
53735 [1au] A? http://www.google.com. (43) (DF)
10:54:27.182852 192.168.1.2.33748 > 149.174.213.7.53:
19315 [1au] A? http://www.netscape.com. (45) (DF)
10:54:43.252461 192.168.1.2.33748 > 66.35.250.11.53:
43129 [1au] A? http://www.linux.com. (42) (DF)
```

All four queries used source port 33748 while querying four different nameservers. If BIND used randomized source ports, it could improve its chances of warding off spoofing attacks.

At this point, all the attacker needs to do is win the race between the first successful collision of his spoofed transactions and the legitimate answer from the true authoritative nameserver. This race is already slanted in favor of the attacker; however, he could utilize other methods to gain an even bigger edge, such as flooding the authoritative nameserver with bogus packets in order to slow down its response time.

After a collision between a legitimate recursive query and a falsified reply packet, the targeted nameserver at the ISP will cache the spoofed record for the time indicated in the TTL section of the reply. At this point the attacker is finished, but the effect lives on for the time the ISP holds the phony record in its nameserver cache. The victim user at the ISP is exposed to the attack any time it makes a query for the domain name in question.

Attack #2: Phase Space Analysis Spoofing

The CERT vulnerability note at <http://www.kb.cert.org/vuls/id/457875> made the following observation:

Additionally, Michal Zalewski's paper "Strange Attractors and TCP/IP Sequence Number Analysis" [ZALEWSKI] describes a method for analyzing the predictability of transaction IDs which we believe could be extended to analyze Transaction ID / UDP port pairs as well.

Once an attacker has managed to poison a DNS cache, there are a number of ways she can subvert protocols that rely on DNS. Some of the potential methods are listed below.

Redirecting Web Traffic

An attack of this nature might range from a simple annoyance to a financial nightmare for a great number of people. The goal here is to set up a website that looks enough like the original so as to not raise any suspicion. Then the domain is hijacked via cache poisoning for as many ISPs/companies as possible, causing their traffic to hit the phony site instead. Some of the sub-attacks here are:

- * Redirect a popular search engine to a pop-up ad site

- * Redirect a bank website to gain access to account passwords
- * Redirect news site to inject false stories and manipulate stocks

Man-in-the-Middle

In this scenario, an attacker tries to intercept secure communication between two parties. For example, Xavier wants to make an online purchase at Yuri's website. The attacker poisons the cache at Xavier's ISP, pointing traffic to Yuri's site to Zamfir's system. Zamfir accepts the incoming SSL connection, decrypts it, reads all the traffic, and makes the same request via SSL to Yuri's site. Replies from Yuri are read by Zamfir then sent back to Xavier over the same encrypted session. Zamfir now has Xavier's credit card number and all other details needed to make illicit use of it.

Recommended Defenses Against DNS Cache Poisoning

Users of BIND

When attempting to protect yourself against a DNS spoofing attack, you have to consider the different aspects of the attack and where you fit in. For instance, do you want to prevent against your users being returned bogus data? Or are you trying to prevent your domain name from being hijacked? In any DNS spoofing attack there are two victims - the hijacked domain owner, who is losing traffic to his site, and the end user who gets redirected to a phony IP address.

Domain Owner

For a domain owner, there is little you can do to protect against someone spoofing your domain name to a vulnerable nameserver. If you are running a webserver, consider using SSL to authenticate yourself to browsers. Eventually DNSSEC will allow all domain servers to have cryptographically signed records, but it is not widely implemented at this time. Even detecting such an attack would be difficult, since the hijack would be largely independent of your servers. Be on the lookout for short denial-of-service attacks; they may indicate someone trying to slow your server down temporarily or crash it in order to complete a spoofing attempt.

ISP Nameserver Admin

You can upgrade BIND to the latest version in the 9.x series, which is not vulnerable to this attack. Alternatively you may try using djbdns, an alternative to BIND written by D. J. Bernstein, author of the MTA program qmail. The djbdns software comes with a security guarantee, basically offering a monetary reward to anyone who publicly discloses legitimate buffer-overflow vulnerabilities in djbdns. Although the guarantee doesn't cover cache-poisoning attacks, I will show later in this article that djbdns offers much greater protection against such attacks when deployed properly.

Disable recursive queries from the outside world, using split-split DNS if possible (see illustration 7). Split-split DNS means you have 2 nameservers; one to serve your public domain information to the outside world, and one to do recursive queries for your users. The public server should not allow recursive queries, and the recursive (caching) server should be protected from the Internet by a firewall.

This is sometimes confused with split DNS, where the internal server forwards requests to the outside server which makes recursive queries on its behalf. This arrangement offers no protection against cache poisoning, and should be avoided.

If you cannot use split-split DNS, you should at least try to restrict who can do recursive queries from your

nameserver. Using the "allow-recursion" option doesn't give you very much protection - remember, the attacker is already spoofing the source IP address, so it is just a matter of them knowing which addresses are allowed to do recursive queries. If possible, use the "listen-on" option to bind the nameserver daemon to an interface that is protected from the outside world.

End User

Urge your ISP/company to upgrade BIND. If they are not open to this, you can always run your own recursive resolver and bypass the ISP's nameservers. Practice safe computing; keep antivirus software updated and signatures current. Always confirm SSL certificates when making secure online transactions. If you suspect a site is being spoofed, you can make use of the ARIN whois records to determine whether an IP address actually belongs to the organization that owns the domain name.

Vendor

This vulnerability could be remedied by changing the behavior of BIND 8 and 4 to only send one request for any number of queries for the same name. BIND 9 already does this, so it is more likely the vendor will just urge people to upgrade instead. However, BIND 9 also demonstrates the same behavior of reusing source ports. This should be corrected, as it would make any new attacks on the PRNG harder to execute.

Proof-of-Concept Code: spooftest.pl

I have written the following program to demonstrate the feasibility of the BIND Birthday attack. It is a passive demonstration only; it cannot be used to execute an actual attack. It will send n number of queries to a recursive server, while generating a ?spoofing set? which is just an array of psuedo-random numbers. If any of the transaction IDs of the recursive queries coming back from the targeted nameserver match a number in the spoofing set, an attack would have succeeded.

Code:

```
#!/usr/bin/perl

#####
# spooftest.pl #
# By William Palmer =me (do not post this script on anyother site without
my consent #
# Tests to see if a BIND server is vulnerable to the Birthday#
# Attack spoofing technique described by Vagner Sacramento in#
# This script MUST be run from an authoritative nameserver's #
# IP address and must be running in place of the real #
# nameserver daemon. It sends a number of queries and sees if#
# it can guess a correct transaction ID from the target IP #
# address. If it guesses correctly, the remote host is #
# considered vulnerable to the BIND Birthday Attack. It does#
# not send any replies, so it is not capable of carrying out #
# an actual attack (sorry script kiddies) #
#####
```

```

use IO::Socket;
use strict;
$| = 1;
my $usage = "Usage: $0 [ip address] [number of packets]\n";

my $target = $ARGV[0];
my $m = $ARGV[1];

die "$usage" unless $ARGV[0] && $ARGV[1];
die "$usage" if $m !~ /^\\d+$/ || $m == 0;

my $domain = "mydomain.com";
# domain name with NS RR pointing to us
my $r = 0xdead;
# initial transaction ID for our queries
my @spoofingset;
# list of our guesses
my $total = 65536;
# total possible packets
my $maxlen = 1500;
my $collisions = 0;
my $datagram;

printf "Probability of success using $m packets: %.2f%%\n",
      100 - (((1 - (1 / $total)) ** ($m * ($m - 1)) / 2)
* 100);

for (0..($m - 1)) {
    $spoofingset[$_] = sprintf("%x", int(rand
($total - 1)));
}

#print "Spoofing set: ", join(" ", @spoofingset), "\n";

my $server = IO::Socket::INET->new(LocalPort => 53,
                                Proto      => "udp")
    or die "Couldn't be a udp server on port 53 : $@\n";

my $client = IO::Socket::INET->new(PeerAddr => $target,
                                PeerPort => 53,
                                Proto      => "udp")
    or die "Couldn't be a udp client on port 53 : $@\n";

my ($second, $stop) = split(/\\.\/, $domain);

```

```

my $findlabel = chr(length($second)) . $second;

my $request = "\x01\x00\x00\x01\x00\x00\x00\x00\x00\x00\x03" .
              "www$findlabel\x03$top\x00\x00\x01\x00\x01";

for (0..($m - 1)) {
    ## send query with incrementing transaction ID
    $client->send(pack("H*", sprintf("%x", $r++)) . $request);
}
print "Sent $m DNS initial queries to target...\n";

$SIG{'ALRM'} = sub { die "timeout"};
my $count = 0;
eval {
    alarm(30);
    while ($count ( $m) {
        $server->recv($datagram, $maxlen);
        my $tid = sprintf("%x", hex(unpack("H*", substr
            ($datagram,0,2))));
        printf "Received recursive query with transaction
            ID: $tid\r";
        for (@spoofingset) {
            if ($tid eq $_) {
                print "\nMatched TID $tid in spoofing set.
                    Success.\n";
                $collisions++;
            }
        }
        $count++;
    }
    alarm(0);
};

if ($?) { if ($? !~ /timeout/) { alarm(0); die $!; } }

print "\nReceived $count recursive queries for $m
initial queries\n";

if ($count == 0) {
    print "Target not listening or does not answer recursive
queries\n";
}
if ($count == 1) {
    print "Target does not appear to be vulnerable\n";
}

```

```
}  
if ($collisions > 0) {  
    print "Spoofing attack would have been successful with " ,  
        "these parameters.\n";  
} else {  
    print "Spoofing attack unsuccessful in this run.\n";  
}
```

Source: <http://www.go4expert.com/articles/dns-cache-poisoning-dcp-t11840/>