ADD / XOR / ROL

Monday, July 21, 2008
On Dan's request for "no speculation please"

I know that Dan asked the public researchers to "not speculate publicly" about the vulnerability, in order to buy people time. This is a commendable goal. I respect Dans viewpoint, but I disagree that this buys anyone time (more on this below). I am fully in agreement with the entire way he handled the vulnerability (e.g. getting the vendors on board, getting the patches made and released, and I understand his decision not to disclose extra information) except the proposed "discussion blackout".

In a strange way, if nobody speculates publicly, we are pulling wool over the eyes of the general public, and ourselves. Consider the following:

Let's assume that the DNS problem is sufficiently complicated that an average person that has _some_ background in security, but little idea of protocols or DNS, would take N days to figure out what is problem is.
So clearly, the assumption behind the "discussion blackout" is that no evil person will figure it out before the end of the N days.

Let's say instead of having an average person with _some_ background in security, we have a particularly bright evil person. Perhaps someone whose income depends on phishing, and who is at the same time bright enough to build a reasonably complicated rootkit. This person is smart, and has a clear financial incentive to figure this out. I'd argue that it would take him N/4 days.

By asking the community not to publicly speculate, we make sure that we have no idea what N actually is. We are not buying anybody time, we are buying people a warm and fuzzy feeling.

It is imaginable that N is something like 4 days. We don't know, because there's no public speculation.

So in that case, we are giving people 29 days of "Thank us for buying you time.", when in fact we have bought them a false perception of having time. The actual time they have is N/4th, and we're just making sure they think that N/4th > 30. Which it might not be. It might be ... 1.

It all reminds me of a strange joke I was told last week. It's a russian joke that makes fun of the former east german government, so it might not be funny to everyone. I apologize up front: I am both german and a mathematician, so by definition the following can't be funny.

"Lenin travels with the train through Russia, and the train grinds to a halt. Engine failure. Lenin sends all workers in the factory that might be responsible to a labor camp.

Stalin travels with the train through Russia a few years later, and the train grinds to a halt. Engine failure. Stalin has all workers in the factory that might be responsible shot.

Honecker (the former head of State of the GDR) travels with the train through Russia. The train grinds to a halt. Engine failure. Honecker has a brilliant idea: "The people that are responsible should be forced to rock the train, so we can sit inside and feel like it is still running." "

It feels like we're all trying to rock the train.

If there was public speculation, we'd at least get a lower boundary on the "real" N, not the N we wish for.

So I will speculate.

The last weeks I was in the middle of preparing for an exam, so I really didn't have time to spend on the DNS flaw. I couldn't help myself though and spent a few minutes every other evening or so reading a DNS-for-dummies-text. I have done pretty much no protocol work in my life, so I have little hope for having gotten close to the truth.

As such, anyone with a clue will probably laugh at my naive ideas. Here's my speculation:

Mallory wants to poison DNS lookups on server ns.polya.com for the domain www.gmx.net. The nameserver for gmx.net is ns.gmx.net. Mallory's IP is 244.244.244.244.

Mallory begins to send bogus requests for www.ulam00001.com, www.ulam00002.com ... to ns.polya.com. ns.polya.com doesn't have these requests cached, so it asks a root server "where can I find the .com NS?" It then receives a referral to the .com NS. It asks the nameserver for .com where to find the nameserver for ulam00001.com, ulam00002.com etc.

Mallory spoofs referrals claiming to come from the .com nameserver to ns.polya.com. In these referrals, it says that the nameserver responsible for ulamYYYYY.com is a server called ns.gmx.net and that this server is located at 244.244.244.244. Also, the time to live of this referral is ... long ...

Now eventually, Mallory will get one such referral spoofed right, e.g. the TXID etc. will be guessed properly.

ns.polya.com will then cache that ns.gmx.net can be found at ... 244.244.244.244. Yay.

The above is almost certainly wrong. Can someone with more insight into DNS tell me why it won't work ?
posted by halvar.flake at 1:17 AM

35 Comments:
greader said...
I don't think this would work because you are racing the .com server's NXDOMAIN (name not found) reply.

There are also 13 root servers and you would have to spoof replies from all of them since you can't know which one the target is using.

/olle

3:05 AM
one.miguel said...
That's how I thought it worked as well, with the non random source ports and the TXIDs being easy to guess.

BTW, that joke was not funny at all. :-)

5:06 AM
Steve said...
I think Klein's DNS work is going to end up being the basis for whatever Kaminsky found. The presentation that Klein gave at RSA in April was very impressive in its efficient guessing of TXIDs and poisoning multiple vendor DNS caches.

The full links get cut off in blogger, but the OSVDB entry has them. Check out both trusteer.com links:

http://osvdb.org/show/osvdb/41143

5:07 AM
Rudd-O said...
This IS PRECISELY the flaw. I have a modicum of protocol knowledge and, once I saw how little randomness TXID includes, this was immediately obvious to me. It should have been immediately obvious to anyone with half a brain too. Even the advisory told everyone (in a half-assed, "we're not telling you but we are" kinda way) what the problem was.

Why security-savvy people either failed to recognize it or tried to keep public discussion under wraps, fails me. Ego, misguided intentions, whatever the reasons, they all are monumentally stupid attempts to prevent the unpreventable, that end up harming us all.

1:37 PM
nate said...
Halvar, you're awesome. I'm pretty sure you guessed it.

-Nate

3:40 PM
Nate McFeters said...
Mmmm... that's a spicy meatball. Halvar, I think that if this is not Dan's stuff, then you have a flaw of your own :O.

BTW, will you be at black hat vegas this year?

-Nate

3:51 PM
natron said...
Your theory was very, very close, but missed one thing: domain servers will check that the additional resource record answer matches the requested one via a process called "bailiwick" checking (according to matasano). A quick google doesn't turn up any refs on that phrase, however, except matasano's now removed post.

In your example, "it says that the nameserver responsible for ulamYYYYY.com is a server called ns.gmx.net". Because of bailiwick checking, ulamYYYYY responses can't poison ns.gmx.net. If it could, it would be trivial to poison any other name server. UlamYYYYY's DNS server has no business telling another DNS server where ns1.google.com is, and being able to submit that reply would be a problem.

That's the explanation, anyway; I'm not sure how it works when you legitimately have a DNS server that's not on your domain. (Like most hosted domains do, e.g. www.customdomain.com's name server is ns1.godaddy.com or similar.)

Anyone know what fields bailiwick checking actually checks?

In any event, you could get around it by just using subdomains and taking advantage of wildcards.

More details here: http://blog.invisibledenizen.org/2008/07/kaminskys-dns-issue-accidentally-leaked.html

4:20 PM
Steve said...
Awesome, Halvar.

4:31 PM
Rudd-O said...
The problem is not just the guessing but that you can include malicious records which lets you poison SOA of a second-level domain. THAT is the bug.

4:39 PM
caf said...
Halvar,

That shouldn't work, because when the resolver later makes an actual query for www.gmx.net, the "gmx.net NS ns.gmx.net" result will also be accompanied by a glue record (A record for ns.gmx.net), which should replace the previous cached result for ns.gmx.net. No idea if that's what the resolvers actually do in practice, though.

By the way, I liked the joke.

5:54 PM
nullbomb said...
heh.. :)
I think the speculation is simply one way to social-engineer the bug details from those who actually more know about it ;-)

www.matasano.com/log/1103/reliable-dns-forgery-in-2008-kaminskys-discovery/ ->
http://blogs.buanzo.com.ar/2008/07/matasano-kaminsky-dns-forgery.html

6:05 PM
nate said...
Halvar,

You're right but you have to push it up a level. Instead of spoofing some random TLD under .com (foo1234.com), spoof a subdomain of gmx.net -- foo1234.gmx.net.

-Nate

6:40 PM

AySz88 said...
You've apparently hit the nail on the head. But now that we know it's really that simple, I think the beans story is appropriate to explain why it's a bad idea to have posted this publicly where every university student (such as me) can find it.

(Ironically, the "don't speculate" appeal is also an instance of "beans".)

8:08 PM
Kago said...
Ok you may be right, and if you are this is the single biggest DICK move in the history of security research IMHO. So Dan's talk at BH is now moot. Well done! Not only have you taken away two weeks of possible patching time for critical gov infrastructure etc, you have stolen a fellow security dude's thunder. I used to think you were cool. Now I just think you are a dick.

8:43 PM
murda said...
That joke was funny. I am from Bangladesh so can see similarities in that.
This is a joke about Bangladesh:

An Bangladeshi dies and goes to hell. There he finds that there is a different hell for each country. He goes to the German hell and asks, "What do they do here?" He is told, "First they put you in an electric chair for an hour. Then they lay you on a bed of nails for another hour. Then the German devil comes in and whips you for the rest of the day." The man does not like the sound of that at all, so he moves on. He checks out the USA hell as well as the Russian hell and many more. He discovers that they are all more or less the same as the German hell. Then he comes to the Bangladeshi hell and finds that there is a long line of people waiting to get in. Amazed, he asks, "What do they do here?" He is told, "First they put you in an electric chair for an hour. Then they lay you on a bed of nails for another hour. Then the Bangladeshi devil comes in and whips you for the rest of the day." But that is exactly the same as all the other hells - why are there so many people waiting to get in?" Because maintenance is so bad that the electric chair does not work, someone has stolen all the nails from the bed, and the devil is a former Govt servant, so he comes in, signs the register and then goes to the canteen..."

10:34 PM
Rory McCune said...
One thing in all this, having heard the detailed explanation from matasano on the vulnerability, is wouldn't it be possible to mitigate this by changing the behaviour of the authoritative name server..?

If I'm understandning things correctly as the authoritative name server for a domain you'd see a whole load of requests for invalid subdomains to your domain (eg, AAAA.MYDOMAIN.COM AAAB.MYDOMAIN.COM) and usually you just respond with NXDOMAIN.

Would it be possible to change that behaviour to respond as the attacker would do with the RR for your valid hosts, so causing the caching DNS server to cache them on the first attempt and preventing the attacker from getting the incorrect entries in first..?

1:13 AM
halvar.flake said...
For those of you that think my post took anything away from Dan's talk:

Imagine there's a world-renowned export on particle physics coming to town, and you want to go see what his theories are. He will give a 1-hour long lecture on his newest discoveries. On your way to the lecture, some random dude on the street corner comes up to you and goes:

"Hey, I think I know what he'll say, it's (..30 seconds of vague mumbling follows..)".

Do you then decide that watching the physics expert is no longer worth it ?

Seriously, if you think that my vague mumblings take anything away from Dan's talk, you're insulting Dan. He's one of the leading experts on DNS, and he'll give a talk about much more than the 8 lines of potential bullshit that I wrote.

1:46 AM
caf said...
natron,

Typically, in that case, when you query the .com servers for www.customdomain.com, they return "customdomain.com NS ns1.godaddy.com" with no corresponding glue record - so your resolver then has to go look up ns1.godaddy.com's A record for itself (obviously, you eventually end up having to rely on glue records somewhere along the line).

You might think "oh, you can just ignore all glue records that are out-of-domain (eg. if looking up x.foo.com, accept a glue record for ns.foo.com but not one for ns.bar.com)" but that fails if foo.com's nameservers are under bar.com and bar.com's name servers are under foo.com (ie. a circular reference). Now that's a pretty stupid setup, so you might well be happy to break it...

2:10 AM
Egill H said...
Here is why it works:

Malory wants to poison the server ns.polya.com

Malory sends NS requests for ulam00001.com, ulam00002.com ... to ns.polya.com.

Malory then sends a forged answer, saying that the NS for www.ulam00002.com is ns.google.com *AND* puts a glue record saying that ns.google.com is 66.6.6.6

Because the glue records corresponds with the answer record, (same domain) the targetted nameserver will cache or replace it's curent record of ns.google.com to be 66.6.6.6

2:58 AM
Marco Massenzio said...
Halvar - it's actually worse than that... if a "dude" can figure it out after reading a "DNS-for-dummies" book (and, from what I've seen, looks as if you got it *almost* right) it totally does my head in how on earth it's possible 'experts' have totally missed it.

That's not being asleep at the wheel: it's more like having taken a dose of sleeping drug before setting off...

BTW - that's the Italian Hell :-)

3:20 AM
Steve said...
Aysz88,

Who gives a damn about stealing someone's thunder? Some of us have networks to secure against smart people that don't post their speculations where the rest of us can benefit from them. Halvar helped us all out by giving us an opportunity to make informed decisions about our networks. Dan can still bask in his glory, and the rest of us can get back to work.

3:41 AM
3247 said...
No, it does not work. If ns.polya.com gets the correct reply for ns.gmx.net first, it's in ns.polya.com's cache.

The odds that the attacker manages to guess the correct TXID the first time is 2^-16. The odds that he also guesses the correct root server is 1/13. The odds that he is also able to outrun the server without having the forge reach ns.polya.com to early are not very good, either.

4:01 AM
mokum von Amsterdam said...
LOL@halvar [not at the joke :P]

Nice chain of thought on the DNS issue.

8:17 AM
RichieB said...
Having listened to many of Dan's talks at Black Hat, he will probably ramble for an hour about unrelated DNS/network voodoo and then unfold the problem which won't be much different from what's been explained above. No disrespect to

Dan, he discovered it, he wrote a prove of concept, he get all the credit. He just should not have tried to keep it quiet for so long. Thank you Halvar.

9:26 AM
Nima said...
Nice blog! If you like we can exchange links on our blogs! My blog talks about information security software tools and resources which is being updated daily, you can also subscribe to see the updates on your Google page:

Information Security Software Tools
http://cryptoexperts.blogspot.com

10:55 AM
natron said...
@Rory, It's actually irrelevant to the attack that the subdomains are irrelevant. You are only kicking off the requests so that you have a (much) greater chance of guessing the TXID. You only have to get one right.

@Egil, that won't work because of the bailiwick checking discussed throughout this thread.

@caf, Thanks for the response! I actually found out a lot more about determining bailiwick from a couple of different sources. This linuxjournal article does a decent job of explaining it:

http://www.linuxjournal.com/article/9905

1:01 PM
vendelhaz said...
My guess is that the problem is related to two (dependent) querys. Both of them should be spoofed. Tihis is the difference from the trivial hacks. But this is just a random guess. Boldi

2:33 PM
MooseBoys said...
It seems to me this would only work for redirecting entries that were unpopular enough to have not been in the cache already. I'm not familiar on the eviction methods, but if the lookup were that unpopular in the first place, wouldn't it likely be bumped out of the cache before anyone actually tried to reference it?

4:50 PM
Paul said...
It looks simple to me. Just send a response that's spoofed from all 13 root servers with about 5000 different TXIDs. you can do this 14 times in a row and one of them is bound to be right. Even better, if you start this before the original request you can cause a denial of service to ensure that the actual root server can't get a response in before yours is accepted. If you have a bot net doing this it's even better.

6:37 PM
the system said...
"The odds that the attacker manages to guess the correct TXID the first time is 2^-16. The odds that he also guesses the correct root server is 1/13. The odds that he is also able to outrun the server without having the forge reach ns.polya.com to early are not very good, either."

Over time, given enough opportunities, the probability of the exploit succeeding tends towards 1.

9:14 PM
Allen Baranov, CISSP said...
I'm sorry to rain on your parade but pretty much everyone has guessed that the weakness is to do with the TXID.

Get rid of the TXID complexity and voila - random ports are your only protection against DNS cache poisoning.

Increasing the complexity of TXID creation may help but is not so easy to sort out quickly.

The question is "how" did Dan do it? Mayber I am giving him too much credit but I doubt he will stand up at a hacker conference and tell them to flood the DNS server with thousands of guesses. Your IPS should block this and the real request will arrive long before you hit the correct TXID.

Flood-guessing the correct TXID is the equivalent of a magician going through a pack of cards and saying "is this your chosen card? no? This one? no? this one? no? this one? no...etc". No magic. I think Dan has found a way to guess the card.

Keeping with the analogy - this blog post is like saying "Dan will guess which card you chose". The magic will be the "how".

12:07 AM
caf said...
allen baranov,

Not exactly. Here's a better analogy for what Halvar's suggesting:

The magician, through the power of suggestion, makes you want to choose a card. You do so, and send it off to a third party (who you cannot see). Then the magician starts shouting guesses: "FIVE OF HEARTS... ACE OF DIAMONDS..." etc, as quickly as he can. Eventually, the third party yells out the correct card, and the magician has failed.

Most of the time.

But the magician can do this over, and over, and over, again; until through sheer chance, he does manage to guess it right, before the real answer comes back. Wow! you think - this guy is for real!

Back in the land of DNS, at this point the attacker has managed to feed you a bogus response for some random subdomain you don't care about anyway. However, you have a certain amount of misplaced trust in this result, which the attacker can transfer to other bogus results, that you do care about.

The magician has all your money.

2:15 AM
murda said...
Sorry halvar-have to stick up for you here and say that I think you're right to discuss the whole issue.
"DNS For Dummies"? More like DNS For a Smart Guy.
Still have to wait and see what the rest of the stuff is that Dan will talk about-like you said, he is a very intelligent dude.
I'm also impressed that you let that post through from the guy being abusive.
This discussion of what can be 'discussed' reminds me of something I read in a Carl Jung book(IIRC); He said that when he was young, he was told it was forbidden to blaspheme against that holy spirit. What was blasphemy against the Holy Spirit? He didn't know, because it was forbidden to talk about that.
So...logically...

4:31 PM
Changlinn said...
@caf thats how I took it as well.

A couple things on this though, it was discovered and an article was written back in 2001 and earlier iirc.

The reason though not to make this widely known is if it is widely known someone will write a tool to exploit it, then every script kiddie and his troll will get that tool and bring done thousands of networks.

There is TIME and a way to do full disclosure doing it without giving the affected (in this case everyone) time to attempt a fix is just silly.

http://morganstorey.com

4:52 AM
MrConceited said...
You've stumbled upon the gist of an attack that has been known for many years. In fact, a more sophisticated version that uses the birthday paradox has been known since 2002.

The new attack is an improvement on that.

7:05 AM