

The Pharming Guide (part 2) : [Whitepapers](#) : [Home](#)

The Pharming Guide (part 2)

Goto Part (1): How DNS Works

Section 3: Pharming Attack Vectors

3.1. Understanding the Attack Vectors

It should now be clear that there are a lot of background processes being executed each time a customer wishes to connect to a named host or online service. Each process relies upon an extraordinary number of systems and routines to function correctly, and different results are possible depending upon the physical location of the customer and the timing of their resolution request. Consequently, there exist a number of vectors through which a Pharmer could conduct their attack.

While there are indeed many ways in which the existing name resolution processes (and DNS in particular) can be attacked, not all are useful within a pharming attack. Since the ultimate goal of a Pharmer is to obtain personal information about a customer (ranging from website authentication details, through to complete identity theft), some attack vectors are more successful than others.

This section focuses upon the different attack vectors available to the Pharmer to conduct their attack and, where possible, the consequences of the attack.

3.1.1. Grouping Attack Targets

Before examining each individual attack vector available to the Pharmer, it is important to understand which groups of hosts or services are likely to come under attack. Each grouping typically has its own unique attack vectors and likelihoods of success. The following figure divides the host resolution process into five target groups:




Figure 11: Host resolution services – attack target groupings

Further explanation of groupings:

- **Local Network** – The local network grouping includes the customer’s host, the physical LAN, any proxy servers and egress firewalls. In addition, if the customer is located within a business environment, local DNS services may also be included.
- **ISP DNS Services** – This group includes all the DNS servers used by the customer, located on the Internet, used for DNS resolution. It includes ISP DNS servers that cache lookup results as well as and resolving services.
- **Global DNS Services** - This group includes all the globally managed services used as part of the resolving process to identify the authoritative name servers for a domain. It includes all Root and TLD servers.
- **Corporate Domain** – This group includes all the services typically owned by a corporate entity to do carryout the IP address resolution of named hosts. As such it includes the authoritative name services for their domain, and any other final delegation processes.
- **Related Resolution Services** – This group includes services not directly related to the DNS lookup process, but which have a substantial effect on the resolution of hosts or online resources. As such, it includes the domain registration process and online search engines or portals.

3.1.2. Attack Classification

Using the groupings discussed above makes for an easier classification of vectors likely to be used during a Pharming attack. The following table categorises attack vectors used by Pharmers and their association with the five different target groupings.




Figure12: Key to attack targets

Attack Vector	Target Group				
	A	B	C	D	E
Human Factors The insider edge	Yes	Yes	Pos.	Yes	
Local Host and Local Network Attacks Modification of lookup processes	Yes				

Traffic observation and modification Man-in-the-middle attacks	Yes Yes	Pos.			
Domain Registration Attacks Domain hijacking Similar domain name registrations Botnet name server registration					Yes Yes Yes
Domain Configuration Attacks DNS wildcards Poorly managed DNS servers		Yes		Yes Yes	
DNS Spoofing DNS cache poisoning DNS ID spoofing with sniffing DNS ID spoofing without sniffing The birthday attack	Yes Yes Pos. Pos.	Yes Yes Yes			
The “New DNS” Attacks Page rank escalation					Yes

3.2. Human Factors

The automated process of resolving a named host to a particular IP address is wholly dependant upon the careful management and configuration of the hosts by technical administrators. These administrators must often manually edit individual DNS server configuration files in order to tune and optimise the services under their care, as well as manage the important host data the services contain.

Because manual intervention is a necessary part of running the DNS system, there are ample opportunities for human factors to affect the security and integrity of the information it contains. The complexity of the DNS system and the interplay between various globally load-balanced servers, coupled with regional optimizations and content delivery options, means that any flaws in the configuration of a single DNS server or name server can often be difficult to identify. Therefore, should a DNS system administrator choose to, he could most likely make malicious alterations to the information about a particular organisations host which may not be spotted for some time – if ever.

For instance, in January 2005 the DNS address for the domain “panix.com” was changed, and the ownership of the New York State ISP was changed from New York to Australia. As part of the attack, requests to reach the panix.com web services were directed to the United Kingdom while email was redirected to Canada.

3.2.1. The Insider Edge

In the past most of the unauthorised modifications made to DNS host entries have been of nuisance value, and rarely for financial gain. However, in the future, the potential profit to be gained by an attacker undertaking a pharming scam is expected to increase the probability of DNS system administrators making temporary changes for cash. With organised crime now taking a keen interest in online identity theft opportunities, the insider threat has never been higher.

Depending upon the location of the conspiring system administrator and the DNS services they have access to, the following risks are present:

- Internal network DNS servers – Within these networks the local system administrator would typically have no difficulty in modifying any static host records, adding custom entries, or modifying cached information without detection. These modifications would only affect the customers operating within the internal network.
- ISP DNS servers – The opportunity for a rogue system administrator to make short-term modifications to DNS entries and conduct pharming attacks is highest within this network level. Since ISPs typically operate several DNS servers simultaneously to cater for different network subscriber demands, modifications to a single DNS server are unlikely to be detected. Any customers relying upon the corrupted DNS server would receive incorrect IP address information and be directed to a different host. Given the high volume of traffic using an ISP’s DNS services, even a single malicious entry is likely to yield substantial numbers succumbing to the pharming attack.

- Corporate name servers – Modifications to resolution information within the authoritative name servers for an organisations domain will yield maximum impact for the Pharmer, but also the highest probability of being detected. Malicious changes to entries within the authoritative name server(s) will result in all network traffic being directed to an alternative IP address. However, unless there is adequate security logging at the name server, short-period (e.g. a few hours each day) alterations are unlikely to be detected by the parent organisation.
- Global DNS Servers – It is possible that system operators could modify values and have a widespread affect on the DNS resolution process. However changes at this level are likely to be noticed fairly rapidly.

3.3. Local Host and Local Network Attacks

Over the last few years there has been a substantial increase in attacks that focus upon gaining control of desktop systems – be that a customers home PC or a corporate workstation. The delivery mechanism chosen for the attack is often related to the type of control over the computer that the attacker wishes to achieve. Since the ultimate goal of the Pharmer is to steal customer identity credentials, there are a number of unique attack vectors focusing upon the local host or LAN that go beyond those previously discussed in “The Phishing Guide”.

3.3.1. Modification of Lookup Processes

Each computer used directly by the Customer must use a predetermined routine to resolve a host name to an Internet routable IP address. As discussed in the earlier section “Local Lookup” (2.2.7), there are a number of local methods that may serve as alternatives to using standard DNS servers. If the attacker is able to gain the ability to modify local lookup preferences, or the files they rely upon, it is possible to conduct a pharming attack.

HOSTS file modification

Since the most common desktop operating systems are based upon the Microsoft Windows platform, many of the successful pharming attacks focus upon modifying the operating systems HOSTS file. With access to this file, the Pharmer is able to add extra entries that will divert the customer’s traffic to a different IP address. For instance, by exploiting a vulnerability in the customers web browser software (typically the customer would have browsed an “infected” webpage that contained specific exploit code waiting to trap any person who visited the page and had a vulnerable version of the browser software), the Pharmer may be able to overwrite the existing HOSTS file with a specially crafted one such as the following:

```
# Pharming Hosts file
200.1.1.10    www.mybank.com
200.1.1.10    mail.mybank.com
200.1.1.10    www.hotmail.com
200.1.1.10    passport.microsoft.com
200.1.1.10    login.passport.net
200.1.1.10    webmail.yahoo.com
200.1.1.10    www.hushmail.com
200.1.1.10    mail.google.com
200.1.1.10    www.google.com
```

In this example, the Pharmer knows that the HOSTS file is used in preference to querying external DNS services and has pointed all the above host names to an IP address he controls. Typically the Pharmer would be running a web server with a number of virtual sites that look like each real site. The Pharmer can choose to just capture the customers login credentials as they enter them into the fake sites and generate an error afterwards, or may choose to transparently proxy the requests to the real servers and capture further confidential details as the customer “uses” the real site.

DNS Network Settings

If the Pharmer has control of the local network hosts (e.g. is a corporate network administrator, runs an Internet café, or can pay an insider to do it for them) it is a simple process to modify the network settings of the computer to point all DNS queries to a DNS server that the Pharmer controls.



Figure 13: Modification of local host DNS server preferences.

3.3.2. Traffic Observation and Modification

With access to the local network, the Pharmer can typically directly observe and modify the destinations of network traffic. While the implications of network sniffing or rogue proxies are obvious, the ability to take control of computer configurations through initialisation services such as DHCP and WPAD are less so.

Rogue DHCP servers

For many networked environments, DHCP is typically used to automatically assign IP addresses and routing information for each computer host as it starts up. These computers must also renew or update the DHCP information from time to time (usually configured centrally by a network administrator). It is possible for an attacker to install a rogue DHCP on a particular network segment and have the local computers use information from it in preference to a centralised server located on a different network segment (mostly due to the speed of response).

By controlling the DHCP settings of a computer, an attacker can state which DNS servers must be used by the customer's computer. If the attacker also has control of a remote DNS server (or has installed his own DNS server somewhere else) he can also provide incorrect host resolution information and direct the customer to hosts of his choice.

In addition, DHCP as implemented in Microsoft operating systems also allows for the definition of a WPAD location.

Rogue WPAD services

The Web Proxy Automatic Discovery (WPAD) service allows web browsers (and other related software) to use a variety of methods to automatically locate suitable proxy services for their traffic. The WPAD service relies upon a number of well-known network protocols to identify and register proxies with computers configured to use the service. Since many popular software products are configured by default to use the service, a rogue WPAD server or suitably constructed entries in the local DNS server can be used by a Pharmer to redirect network traffic to a proxy server of their choice – thereby carrying out a man-in-the-middle attack.

3.3.3. Man-in-the-middle Attacks

Man-in-the-middle attacks often form a key component of a sophisticated Phishing or Pharming attack. With access to a customer's local network, this attack delivery platform is much simpler and often harder to detect. For more discussion about man-in-the-middle attacks, readers are referred to section 2.3.1 of "The Phishing Guide".

3.4. Domain Registration Attacks

Domain registration attacks abuse the way in which a domain may be registered with a registrar. The most common vectors for attack are:

- Domain Hijacking
- Similar domain name registrations
- Botnet name server registrations

3.4.1. Domain Hijacking

In order for an organisation to make use of a domain name, they must first register it. This process is done through various domain registration authorities; typically by paying a small fee to the registration authority to maintain ownership of the domain for a set number of years (typically 1-3 years). Domain Hijacking is the process by which a domain with a lapsed registration is purchased by another person and is then used for some other purpose.

Registration information is managed by various Internet registrars, and can be queried using several tools (such as whois) and other related online services. For instance, querying the registrar for information about the TECHNICALINFO.NET domain, we retrieve the following information:

```
Domain Name: TECHNICALINFO.NET
Registrar: MELBOURNE IT, LTD. D/B/A INTERNET NAMES WORLDWIDE
Whois Server: whois.melbourneit.com
Referral URL: http://www.melbourneit.com
Name Server: NS0.PHASE8.NET
Name Server: NS1.PHASE8.NET
Name Server: NS2.PHASE8.NET
Status: REGISTRAR-LOCK
Updated Date: 14-dec-2004
Creation Date: 20-jun-2002
Expiration Date: 20-jun-2006
```

It is important to note that, in the above example, the ownership of TECHNICALINFO.NET will expire on the 20th June 2006. If the current owner does not renew his registration by that date, anyone could purchase the domain and take ownership.

By 'hijacking' an existing domain, as opposed to registering a new domain, the new owner can take advantage of any existing links to it – thereby guaranteeing a number of "backlinks" and associated traffic. Domain hijacking an increasingly popular mechanism for advertisers and other organisations that generate revenue from customers connecting to their sites.

In a Pharming attack, the attacker would seek to take ownership of the domain as soon as the current owner neglects to re-register the domain. With ownership, the Pharmer would construct a new website (and other related Internet services such as email) to replicate the earlier version and fool any customers who connect to the site.

Alternatively, the Pharmer may choose to utilise a heavily trafficked site (not associated with an organisation they plan to target) to provide links (hidden or otherwise) to an alternative site under their control and increase its search engine rankings. This attack vector is explained fully in a later section.

3.4.2. Similar Domain Names

Perhaps one of the simplest attack vectors of all, the Pharmer registers multiple spelling and key mashing (e.g. "fat fingers" hitting neighbouring keys simultaneously) permutations of the target host name hoping that a customer will mistype it. If a customer does mistype the host name, instead of getting a "no such host" message or connecting to a host that is clearly not their desired destination, the attacker has created a fake version of the site to fool the customer and steal their credentials.

The registration of similar domain names has been abused for many years, but the nature of the attack was often somewhat different. Many adult entertainment or advertising sites make use of alternative domain name registrations to capture the attention of a potential customer or generate site traffic. A past example includes www.whitehouse.com, which was once a porn site and obviously not affiliated with www.whitehouse.gov – a US government information site.

More recently, attackers have made use of key mashing permutations of popular websites to direct customers to malicious websites. For example, in April 2005 an attacker registered google.com and msmn.com for the purpose of secretly infecting the computers of users of Google and MSN who has mistyped their host names with Spyware, Adware, and other malicious software.

It is expected that this particular attack vector will become increasingly popular as a mechanism of infecting customer computers with malicious software and stealing personal authentication information.

3.4.3. Botnet Name Server Registration

The domain registration process allows the registrant to list the authoritative servers responsible for managing the IP address lookups of the hosts within that domain. It is normally recommended that at least two name servers be listed (a primary and a backup) and that they be located on different network segments to help cope with network resiliency issues.

Botnets have been used in phishing attacks in the past – mainly to host multiple copies of the faked website at several IP addresses. As each host or IP address is closed down by the owner ISP, the Phisher just modifies his DNS to point to alternative location. If the ISP has control of the DNS entry (i.e. the Phisher is using an ISP's DNS server as their authoritative domain server), it is a simple process for the ISP to remove the entry (or point to the real website) and effectively close down access to all the fake websites in one go.

While it is recommended that at least two name servers be listed, the registrant can choose to list many more. This ability to list multiple entries can be abused during a pharming attack and, when combined with an established Botnet, can be very difficult to shut down once identified.

By ensuring that the multiple name servers registered by the Pharmer are spread across several ISP's, no single ISP can shut down the DNS service. Instead, the organisation being targeted by the attack must attempt to deal with the registrar of the malicious domain to close it down. Unfortunately many registrars do not have formal procedures for dealing with this kind of request.

3.5. Domain Configuration Attacks

3.5.1. DNS Wildcards

DNS Wildcards are special entries within a DNS configuration file for handling “catchall” name resolutions. For instance, an organisation may have two internet hosts – www.mybank.com for web traffic and mail.mybank.com for email – but wish to make use of more intuitive host names for their customers at a later date. Instead of continuously updating their DNS configuration, they may add a pair of wildcard entries to direct network traffic to these two hosts that were destined for other host names.

For example, the authoritative DNS server for the domain mybank.com may be configured to have the following entries:

www.mybank.com	IN	A		150.10.1.21
mail.mybank.com	IN	A		150.10.1.20
mybank.com	IN	A		150.10.1.21
mybank.com	IN	MX	10	mail.mybank.com
*.mybank.com	IN	MX	10	mail.mybank.com
*.lon.mybank.com	IN	A		150.10.1.21

Here we see that the * wildcard entries are designed to do the following:

★ Direct all emails destined for [something].mybank.com to the mail server mail.mybank.com. For example, this would handle emails addressed to security@newyork.mybank.com and gunter@foo.bar.london.mybank.com.

★ Direct all connections to hosts with names ending with lon.mybank.com to the IP address 150.10.1.21. For example, this would direct customers requesting http://customergateway.lon.mybank.com to the same host IP address as www.mybank.com.

In the past Phishers have spoofed email source addresses of organisations that did not have DNS wildcard entries for their mail servers (i.e. MX records) so that, should a recipient of one of their fake emails attempt to reply to it, the response would never be received/intercepted by the organisation and thus have a lower likelihood of discovering that their organisation was an unwitting participant in a phishing attack.

Pharmers (and many spammers) make use of DNS wildcard entries to obfuscate the true destination host of their attack. For example:

★ If the Pharmer owns the top level domain (e.g. “pharmer.com”) he may use a host name http://www.mybank.com.Login.html.134534534.pharmer.com/

★ The Pharmer may abuse common link sites to redirect the victim to a server of their choice. For instance, in March 2005 an attacker abused the DNS wildcard configuration of a third-party redirection service (Kickme.to) to target Barclays banking customers with URL's such as:

- ▲ http://barclays.co.uk|YJ3EMOHOqljQ8J5oW2ZKyTaRMQOahSWaxTrFTEQK9I9VVQj6jDtyq10d2
- ▲ http://barclays.co.uk|34fdbcb4rvidnp9phxbahhvbs6l56a2uyx%2edivxmovies%2ea%74/41pvaw3/

★ Spammers often use DNS wildcard entries to embed unique tracking information within the host name to verify real email accounts and bypass anti-spam filtering software.

3.5.2. Poorly Managed DNS Servers

While it is true that new vulnerabilities within core DNS software are being discovered continuously, there is also a parallel stream of patches and security fixes being issued by the various vendors to correct them. If a DNS server is being managed correctly, these patches and updates will be installed shortly after being made available – thereby limiting the window of opportunity for an attacker seeking to exploit the new vulnerability.

Poorly managed DNS servers tend to be several patch cycles behind and tend to suffer from poor configurations and lax authentication controls. This means that an attacker can often easily gain control of the DNS server. Depending upon the role of the DNS server (e.g. ISP-level DNS caching, DNS resolving, corporate name server, etc.), the attacker can use the compromised DNS as part of a successful Pharming attack as if they were an insider (see section 3.2.1).

3.6. DNS Spoofing

A DNS spoofing attack can be defined as the successful insertion of incorrect resolution information by a host that has no authority to provide that information. It may be conducted using a number of techniques ranging from social engineering through to exploitation of vulnerabilities within the DNS server software itself. Using these techniques, an attacker may insert IP address information that will redirect a customer from a legitimate website or mail server to one under the attacker's control – thereby capturing customer information through common man-in-the-middle mechanisms.

According to the most recent "Domain Health Survey" (Feb 2003), a third of all DNS servers on the Internet are vulnerable to spoofing.

The Attack

Operating normally, a customer can expect to query their DNS server to discover the IP address of the named host they wish to connect to. The following diagram reflects this process.




Figure 14: The normal DNS resolution process

1. The customer queries the DNS server – "What is the IP address of www.mybank.com?"
2. The DNS responds to the customer query with "The IP address of www.mybank.com is 150.10.1.21"
3. The Customer then connects to the host at 150.10.1.21 – expecting it to be www.mybank.com.

However, with a successful DNS spoofing attack, the process has been altered. The following diagram reflects this process.




Figure 15: The DNS resolution process having fallen victim to a DNS spoofing attack

1. The attacker targets the DNS service used by the customer and adds/alters the entry for www.mybank.com – changing the stored IP address from 150.10.1.21 to the attackers fake site IP address (200.1.1.10).
2. The customer queries the DNS server – “What is the IP address of www.mybank.com?”
3. The DNS responds to the customer query with “The IP address of www.mybank.com is 200.1.1.10” – not the real IP address.
4. The Customer then connects to the host at 200.1.1.10 – expecting it to be www.mybank.com, but in fact reaching the attackers fake site.

Use of Botnets

Botnets may be used within many of DNS spoofing attacks as a force multiplier in the following ways:

- As a denial of service agent against an authoritative name server to cause it to respond slowly (or not at all) to resolver requests for valid host IP address information.
- As a delivery mechanism for fake UDP responses from the spoofed IP address of the authoritative name server.

3.6.1. DNS Cache Poisoning

One attack vector for DNS spoofing is through cache poisoning. In this attack the attacker abuses caching vulnerabilities within the DNS server to add multiple resolution entries for hosts not originally asked for and is not authorised to provide. While most new DNS service implementations are not vulnerable to cache poisoning, there are still a large number of vulnerable DNS servers that are.

The process in which a DNS server may have its cache poisoned can be explained in the following diagram and walkthrough.




Figure 16: The DNS cache poisoning process

1. The attacker queries the DNS server for the IP address for a host that is managed by a name server owned by the attacker – “What is the IP address of `www.attackerowned.com`?”
2. The DNS Caching server does not have a cached entry for `www.attackerowned.com` and must resolve the IP address by querying the authoritative name server for the `attackerowned.com` domain. This authoritative name server belongs to the attacker.
3. The attackers name server informs the DNS caching server that the IP address of `www.attackerowned.com` is `200.1.1.10`. In addition, the attackers name server also includes additional (faked) resolution records such as:
 - a. `www.mybank.com` is `200.1.1.11`
 - b. `mail.mybank.com` is `200.1.1.11`
 - c. `secure.mybank.com` is `200.1.1.11`
4. The DNS caching server responds to the attacker’s original query with – “The IP address of `www.attackerowned.com` is `200.1.1.10`.” This result, along with the extra resolution records, is cached by the DNS server for a period equivalent to the TTL supplied by the attackers name server.
5. At a later date, an ordinary customer who also uses this DNS caching server queries it for the IP address of `www.mybank.com` – “What is the IP address of `www.mybank.com`?”
6. The corrupted DNS caching server responds to the customer query by supplied the previously cached (and fake) answer – “The IP address of `www.mybank.com` is `200.1.1.11`” – instead of the real `150.10.1.21` address.

For instance, In July 1997 Eugene Kashpureff of AlterNIC used a program to “poison” the caches of major name servers around the world. This caused traffic originally destined for `www.internic.net`’s address to go to the IP address of the AlterNIC web server. No attempt was made to disguise the attack, and customers who tried to reach `www.internic.net` were confronted with the AlterNIC website.

3.6.2. DNS ID Spoofing with Sniffing

DNS lookup queries by Customer hosts rely upon the UDP protocol (an important Internet protocol that, unlike TCP, does not use any form of handshaking) to request and obtain resolution information. Each UDP-based DNS query originated from the customers computer will also be assigned a unique identifier (ID for short) to help manage multiple lookup responses. Any queried DNS sever is supposed to include the same query ID as the request. If the ID supplied is not the same as that of the originating request, the customers computer should ignore the response.

DNS ID spoofing is an attack vector which focuses upon providing incorrect or malicious DNS resolution information to customer requests after having observed the ID of their request. To be successful, the attacker must observe the customers request (most often achieved by sniffing the network traffic) and be capable of constructing a spoofed response faster than the DNS server can supply the legitimate answer.




Figure 17: The DNS ID spoofing process

In the figure above, the process of DNS ID spoofing is as follows:

1. The customer sends a UDP-based request to the DNS server – “What is the IP address of www.mybank.com?” – with an ID of 117.
2. The attacker has positioned a laptop on the network to sniff all network traffic and respond to DNS queries.
3. Having identified a DNS request for www.mybank.com with an ID of 117, the attacker's laptop automatically responds with “The IP address of www.mybank.com is 200.10.1.11” using the same ID.
4. A few fractions of a second later, the real DNS server sends its response “The IP address of www.mybank.com is 150.10.1.21”, but is ignored by the customer's computer since it has already received a response from the attacker's machine.

3.6.3. DNS ID Spoofing without Sniffing

While DNS ID spoofing is a relatively simple process if the attacker is able to monitor a customer's network traffic for DNS lookup requests, the attack is restricted to environments where the attacker can gain physical network access to a network segment shared with the customer. To overcome this limitation, an attacker may use a combination of techniques to achieve ID spoofing without relying upon network sniffing.

A limitation of the UDP-based DNS lookup process is that the ID is coded to 2 bytes, meaning that there are only 65535 possible values. Therefore, for an attacker to succeed in carrying out the previous attack without sniffing, he would have to either guess the correct ID or rapidly produce 65535 spoofed responses before the DNS server could respond.

An additional problem is to know exactly when to launch the attack. This problem can be overcome by using a process similar to the one discussed previously on DNS cache poisoning – the attacker actually launches the initial lookup request and the DNS caching server (the victim in this attack) must query the authoritative name server for the IP address information.

Early versions of a popular DNS software implementation suffered from a security flaw that resulted in all DNS transaction ID's being non-random – instead they were sequential. This of course made the “guessing” of a request's ID a simple process. Following a CERT advisory (CA-1997-22), organisations using this software were advised to upgrade to a new version that implemented random transaction ID's.




Figure 18: The DNS ID spoofing attack not relying on sniffing

In the figure above, the process of DNS ID spoofing without relying on network sniffing is as follows:

1. The attacker makes a request to the DNS server – “What is the IP address of www.mybank.com?”
2. The DNS server must query the mybank.com authoritative name server (150.10.1.2) for the IP address of www.mybank.com and has assigned a DNS ID of 45889 to this request.
3. While the DNS server is attempting to resolve the IP address of www.mybank.com and awaiting a response from the authoritative name server, the attacker has launched a flood of UDP responses from a host under their control using the spoofed IP address of the mybank.com name server (150.10.1.2). Each spoofed response has a different DNS ID and states that the IP address of www.mybank.com is 200.10.1.11.
Note: The UDP packets must be sent to a specific port as well. While this in theory must also be guessed, in practice many DNS servers use the same source port for their queries. Therefore the attacker would, in advance of this spoofing attack, have constructed their own authoritative name server and observed the DNS source port after having queried the DNS server for the IP address of a host listed on the attackers name server.
4. If the attacker has succeeded (i.e. managed to guess the DNS response with the correct ID using the spoofed name server IP address to the DNS server before the real mybank.com name server did) he will receive a response from the DNS caching server stating “The IP address of www.mybank.com is 200.10.1.11”. This IP address information is then cached for the stipulated TTL and will be supplied by the DNS server in response to any later customer requests.
5. The mybank.com authoritative name server responds with the correct IP address for www.mybank.com (150.10.1.21), but is ignored by the requesting DNS server if it has already received a “valid” spoofed response from the attackers system.

3.6.4. The Birthday Attack

Closely related to the previous attack vector, a “Birthday Attack” exploits a weakness discovered in 2002 relating to the fact that the most popular DNS implementation (BIND) would send multiple simultaneous recursive queries for the same IP address (now fixed in the latest versions of the software). This repetitive behaviour means that a “Birthday Paradox” could be used to mathematically increase the speed and probability of a successful attack by reducing the number of spoofed guesses of the DNS transaction ID from tens of thousands down to a few hundred.




Figure 19: The DNS Birthday Attack

In the figure above, the birthday attack is carried out as follows:

1. The attacker launches repeated requests to the DNS caching server asking "What is the IP address of www.mybank.com" as fast as possible.
2. Simultaneously, the attacker also sends repeated spoofed responses using different DNS transaction ID's stating that "The IP address of www.mybank.com is 200.10.1.11".
3. For each request from the attacker in (1), the DNS server tries to resolve the IP address for www.mybank.com by querying the authoritative mybank.com name server – typically using a different DNS transaction ID for each request. Based upon the mathematical properties of the Birthday Paradox, there is a higher probability that the attacker can "guess" a correct DNS transaction ID (thereby "answering" the DNS servers query) faster than the real name server can respond.

To further increase the odds of the attacker supplying a correct DNS transaction ID with the spoofed message, the attacker could target the authoritative name server with other requests or denial of service techniques to slow down its response to the DNS caching server.

Why does the attack work?

The Birthday Attack is named after a mathematical result that establishes that the probability that two or more people in a group of 23 share the same birthday is greater than 50% - the so called "Birthday Paradox". This mathematical principle can be applied to pseudo-random number generation; which in this case is the process for generating DNS transaction ID's.

During a conventional DNS ID spoofing attack (section 3.6.3) the attacker would send n spoofed responses for a single query – resulting in a probability of success of $n/65535$. During a DNS Birthday Attack the attacker sends n spoofed replies for n queries for which the probability of success (P) becomes:

$$P = 1 - \left(1 - \frac{1}{t}\right)^{\frac{n(n-1)}{2}}$$

Plotting this equation, where t represents the maximum number range of DNS transaction ID's (65535), we observe the following results for the first 1000 values of n .




Figure 20: The DNS Birthday Attack probability of success graph

From this quick analysis, we can see that the probability of success reaches 50% when n is approximately 300 and 99% as n approaches 800. Using a conventional DNS ID spoofing attack (section 3.6.3), 300 spoofed responses would have only yielded less than 0.5% success.

3.7. The “New DNS” Attacks

The growing popularity of the “new DNS” coupled with the way customers rely upon these services to locate frequently accessed Internet resources represents an ideal target population – some would say “low hanging fruit” - for Pharmers. With some forward thinking and relatively little effort, a Pharmer can corrupt the names associations these services offer at regional or global levels and target specific customer bases.

The marketing opportunities that many popular search engines provide, means that Pharmers can purchase “sponsored links” or similar services which will place their hyperlinked resources (i.e. link to their faked website) at the top of a customers search page response. Closely related to the exploitation of banner advertising by Phishers (see section 2.2.2 – Web-based Delivery – of “The Phishing Guide”), Pharmers may exploit the validation processes of some search engine providers that are not as rigorous as others, and can provide links to their fake sites using keywords or phrases normally associated with the targeted organisation.

3.7.1. Page Rank Escalation

With prior planning, a Pharmer can seek to increase their search engine page ranking by abusing the way that provider actually calculates the ranking. By taking advantage of the page ranking system the Pharmers goal is to get their fraudulent link (and extracted page text) to appear in the place of where a customer would normally expect to find the real link, preferably first on the list.

For Example, using the Google search engine, customers who type in “Citibank” would normally expect to see several links to Citibank. However, by exploiting the page weightings explained in section 2.3.2, it may be possible for the Pharmer to cause the following to appear in response to the query:



Figure 21: Search engine result influencing

In the example figure above, we note that:

- The Pharmer has paid for the “sponsored link” to point to their host – www.citibank.hackeddomain.com. This purchase could have been made through the use of stolen credit card details, or other fraudulent financial transactions.
- The first non-sponsored link of the returned search contains the expected “Welcome to Citibank” message and supporting text. However, it is only by taking a closer look at the URL in green that the customer would notice anything untoward. In this case, the link points to the fake site hosted at www.latam.citibank.hackeddomain.com.
- The fact that the Pharmer has managed to get his fake site listed first is important for customers that make use of search engines that offer an “I’m feeling lucky” search option. With such an option available, the customer would have been automatically directed to this top-listed faked result.

The successful manipulation of search page rankings provides a very good delivery platform for the Pharmer because of the way it can be targeted to a specific customer audience or region (i.e. most popular search engines offer regional responses to search queries) and the difficulty for the victim organisation to identify or shut down the attack.

Section 4: Defence Mechanisms

Pharming attacks tend to be harder to defend against than traditional Phishing attacks due to the distributed nature of the attack focus and the use of resources not under the control of the victim organisation. In addition, the manipulation of the DNS resolution process occurs at such a fundamental level that there are very few methods available to reliably detect any malicious changes.

4.1. Classic Phishing Defences

Many of the defences used to thwart phishing attacks can be used to help prevent or limit the scope of future Pharming attacks. While readers are referred to the detailed coverage of these defence tactics explained in “The Phishing Guide”, a brief summary of these key defences is as follows:

Client-side

- Desktop protection technologies
- Utilisation of appropriate, less sophisticated, communication settings
- User application-level monitoring solutions
- Locking-down browser capabilities
- Digital signing and validation of email
- General security awareness

Server-side

- Improving customer awareness
- Providing validation information for official communications
- Ensuring that the Internet web application is securely developed and doesn't include easily exploitable attack vectors
- Using strong token-based authentication systems
- Keeping naming systems simple and understandable

Enterprise

- Automatic validation of sending email server addresses,
- Digital signing of email services,
- Monitoring of corporate domains and notification of "similar" registrations,
- Perimeter or gateway protection agents,
- Third-party managed services.

4.2. Additional Pharming-specific Defences

While Phishing attacks typically use email as the attack delivery platform, Pharming attacks do not require any email obfuscation attacks to succeed – therefore Phishing defences that rely upon email security play a lesser role. The defences that will be most successful in preventing Pharming attacks focus upon the following areas:

- Change management, monitoring and alerting
- Third-party host resolution verification
- DNS server patching, updating and configuration
- Search engine control

4.2.1. Change Management, Monitoring and Alerting

The potential for an administrator or other authoritative employee to maliciously modify DNS resolution information without detection is great. As financial incentives increase, organisations and ISP's will need to ensure that adequate change control, monitoring and alerting mechanisms are in place and enforced.

It is recommended that:

- Wherever editing is possible, access to DNS configuration files and caching data is limited to approved employees only.
- A change management process is used to log and monitor all changes to DNS configuration information.
- Auditing of DNS record changes is instigated by a team external to any DNS administrative personnel; with automatic alerting of changes conducted in real time.
- Regular audits and comparative analysis of secondary DNS and caching servers should be conducted.

4.2.2. Third-party Host Resolution Verification Services

Toolbars

Many third-party developed plug-in toolbars originally designed to detect Phishing attacks are deceived by Pharming attacks. Typically, these Phishing toolbars show the IP address and reverse lookup information for the host that the browser has connected to, so that customer can clearly see if he has reached a fake site. Some managed toolbars (normally available through a subscription service) also compare the host name or URL of the current site to an updatable list (or real-time querying) of known phishing sites.

Some toolbars now offer limited anti-pharming protection by maintaining a stored list of previously validated "good" IP addresses associated with a particular web address or host name. Should the customer connect to an IP address not previously associated with the host name, a warning is raised. However, problems can occur with organisations that change the IP addresses of their online services, or have large numbers of IP addresses associated with a particular host name.

In addition, some toolbars provide IP address allocation information such as clearly stating the geographic region associated with a particular netblock. This is useful for identifying possible fake Pharming sites that have been setup in Poland pretending to be for an Australian bank for instance.

Server Certificates

To help prevent pharming attacks, an additional layer can be added to the authentication process, such as getting the server to prove it is what it says it is. This can be achieved through the use of server certificates.

Most web browsers have the ability to read and validate server identification certificates. The process would require the server host (or organisation) obtain a certificate from a trusted certificate authority, such as Verisign, and present it to the customer's browser upon connection for validation.

4.2.3. DNS Server Patching, Updating and Configuration

As with any Internet-based host, it is vital that all accessible services be configured in a secure manner and that all current security updates or patches be applied. Failure to do so is likely to result in an exploitation of any security weaknesses, resulting in a loss of data integrity.

Given the number of possible attacks that can be achieved by an attacker whom manages to compromise an organisation's DNS servers, these hosts are frequently targeted by attackers. Therefore it is vital that security patches and updates be applied as quickly as possible – typically organisations should aim to apply fixes within hours of release.

Similarly, it is important that organisations use up to date versions of the service wherever possible. As we have already discussed in section 3.6, each new version of the DNS software usually contains substantial changes to protect against the latest attack vectors (e.g. randomising DNS ID's, randomising port numbers, etc.)

Configuration

DNS servers typically offer organisations and their administrators an extensive number of configuration options. Therefore great care must be made during the installation and configuration process if the service is to be deployed securely.

Many common configuration mistakes are exploited for spoofing attacks. To help defend against spoofing attacks, the following advice should be heeded:

- Name servers are typically vulnerable to spoofing attacks from the Internet when configured to accept recursive queries. As explained in section 3.6.1, an attacker can query the name server for a host belonging to a domain managed by a DNS server under their control. To help prevent this:
 - If possible, turn off recursion.
 - Restrict the addresses to which the name server will respond to queries from.
 - Restrict the addresses to which the name server will respond to recursive queries from.
- When configured to disable recursion, the name server is operating in a passive mode – not sending queries on behalf of other name servers or resolvers. This has the following effects:
 - Since a non-recursive name server doesn't send out queries, it doesn't cache any data, and is difficult to spoof.
 - It is not advisable to disable recursion on a name server if it is relied upon by other servers for resolution or as a forwarder
- If it is not possible to turn off recursion, organisations should aim to restrict the types of queries that the name server accept.
 - Restrict queries based upon the addresses they came from.
 - Restrict queries based upon the zones that are allowed to ask about.
- "Glue Fetching" is the term used to define the process in which a name server attempts to retrieve A records to accompany any NS record requests. This process has the potential for receiving spoofed responses. By turning off glue fetching, the following occurs:
 - The name server will not generate these queries.
 - The name server will not build up a cache.

4.2.4. Search Engine Control

Internet search engines are undergoing constant development. Many of the methods used by attackers to increase their page ranking statistics are known of by the search engine developers, and a constant cycle of detection and refinement can be observed by both parties. For instance, Google modified its search algorithm to "reset" the page rank statistics of web sites that had recently changed ownership – this was to reduce the impact of instant "backlinks" and the weighting they attach to a ranking.

Traditionally the emphasis on increasing a pages ranking has been for revenue or lead generation – most closely associated with advertising. However, the increasing pace at which customers are relying upon search engines to access key services (such as online banking) means that a Pharmer who can get his fake site ranked at the top is likely to acquire a high number of victims.

Organisations should ensure that they regularly review keyword associations with their online services. Ideally automated processes should be developed to constantly monitor all the popular search engines for key search words or phrases customers are likely to use to locate their key services. It is also important that region-specific search engines also be monitored.

Section 5: Summations

Attacks focused upon host name resolution processes are likely to be of increasing importance to attackers seeking financial gain or to conduct identity theft operations. The lack of understanding customers (and many organisations) have of the background processes necessary to resolve IP address information to named hosts or services means that attacks that manipulate these DNS services are likely to go unnoticed.

Building upon the success of Phishing attacks, this new class of Pharming attacks enables the attacker to reach a wider customer audience with very little effort and a lower probability of detection.

To combat the new Pharming threat, it is vital that organisations understand how global DNS resolution services function and how they may be manipulated by an attacker. Armed with this knowledge, organisations can develop better monitoring and alerting processes to detect Pharming attacks early on – attacks that would most likely have escaped detection.

5.1. Resources

[“The Phishing Guide”](#), Gunter Ollmann, 2004

“DNS and Bind”, O'Reilly, 2001

“Addressing Weaknesses in the Domain Name System Protocol”, Christoph Schuba, Purdue University, August 1993

[“Security Best Practice: Host Naming and URL Conventions”](#), Gunter Ollmann, 2005

“DNS Cache Poisoning – The Next Generation”, Joe Stewart, 2003

“The Anatomy of a Large-Scale Hypertextual Web Search Engine”, Sergey Brin and Lawrence Page, 2000

Information Links

Current status and physical location of Root Servers - <http://netmon.grnet.gr/stathost/rootns/>

Top level generic domain name listing - <http://www.iana.org/gtld/gtld.htm>

Top level country-code domain name listings - <http://www.iana.org/cctld/cctld-whois.htm>

DNS Spoofing Techniques - <http://www.securesphere.net/download/papers/dnsspoof.htm>

Birthday Paradox - http://en.wikipedia.org/wiki/Birthday_paradox

Cache Poisoning - <http://www.lurhq.com/cachepoisoning.html>

SANS Warning - <http://isc.sans.org/diary.php?date=2005-03-04>

Copyright 2001-2007 © Gunter Ollmann