

Internet Footprinting, by incognit

Notes taken from the first chapter of *Hacking Exposed: Network Security Secrets and Solutions*, by Stuart McClure, et al.

Internet footprinting is comprised of four steps

I. Determine the Scope of Your Activities

II. Network Enumeration

III. DNS Interrogation

IV. Network Reconnaissance

V. Summary

I. Determine the Scope of Your Activities

Footprint entire organization, or limit to certain locations?

Items of interest

- Locations
- Related companies and entities
- Merger or acquisition news
- Phone numbers
- Contact names and email addresses
- Privacy or security policies indicating the types of security mechanisms in place
- Links to other web servers related to the organization

Open Source Search

Peruse the target's web page; review HTML source code for comments.

Perform open searches for information relating to the target organization (News articles, press releases, etc); this may give insight to organization's current state and their security posture.

Any recent security incidents? (Use web search engines)

Search USENET for postings related to @targetdomain.com; esp. regarding new HW/SW they are setting up or asking questions about. Postings can give insight into the technical prowess of an organization's staff.

Use AltaVista or Hotbot to search for all sites that have links back to the target organization's domain. This list may include rogue web sites put up by employees that aren't sanctioned or secure.

EDGAR Search

EDGAR database search (for publicly traded companies): <http://www.sec.gov>. Organizations typically have problems managing their Internet connections, especially when they are actively acquiring or merging with other entities. Focus on newly acquired entities. Two of the best SEC publications to review are the 10-Q and 10-K.

- 10Q: Snapshot of what the organization did over last quarter; including purchase or disposition of other entities.
- 10K: Yearly update of what the organization has done

Peruse the documents by searching for "subsidiary" or "subsequent events." This may provide you with information on a newly acquired entity. Often, organizations will scramble to connect their newly acquired entities to their corporate networks with little regard for security. This may allow you to leap-frog into the parent company. Take advantage of the chaos that results from combining two networks.

Countermeasure: Public Database Security

Remove any unnecessary information from your web pages that may aid an attacker in gaining access to your network.

II. Network Enumeration

Identify domain names and associated networks related to a particular organization.

Query types of interest

- *Organizational*: Displays all information related to a particular organization
- *Domain*: Displays all information related to a particular domain
- *Network*: Displays all information related to a particular network or single IP address
- *Point of Contact (POC)*: Displays all information related to a specific person, typically the administrative contact.

Organizational Query

Perform a "whois" on target organization: \$ whois "Target Org". Are the results real networks? Or are they reserved for future use? The InterNIC will truncate it to 50 names.

You can still download entire .edu domains from <ftp://rs.internic.net/domain>. After downloading, manipulate it using shell commands or Perl.

To bypass the 50 name limit, go to <http://www.websitez.com>. This site has most domains indexed, and will provide all records associated with a particular domain.

Whois searching techniques and data sources:

Mechanism	Resources	Platform
Web interface	http://www.networksolutions.com/	Any platform with a web client
Web interface	http://www.arin.net/	Any platform with a web client
Whois client	Unix; Fwhois <ccappuc@santafe.edu>	UNIX
WS Ping Pong	http://www.ipswitch.com/	Win 9x/NT
Sam Spade	http://www.blighty.com/products/spade	Win 9x/NT
Sam Spade Web Interface	http://www.samspade.org/	Any platform with a web client
Netscan tools	http://www.nwpsw.com/	Win 9x/NT
Xwhois toolkit	http://www.goatnet.ml.org/software.html	UNIX with X and GTK+ GUI

Government, military, and international company sources of Whois database:

Whois Server	Address
European IP Address Allocations	http://whois.ripe.net/
Asia Pacific IP Address Allocations	http://whois.apnic.net/
U.S. Military	http://whois.nic.mil/
U.S. Government	http://whois.nic.gov/

Domain query

Perform “whois” on most likely domain: \$ whois mostlikelydomain.com

Domain query information

- The registrant
- The domain name
- The administrative contact
- When the record was created and updated
- The domain name system servers (DNSes)

Analyze the information for clues that will provide you more information.

Enticements – excess information or information leakage that may entice an attacker into mounting a more focused attack.

Review the information in detail:

- Ascertain if the domain belongs to the entity that we are trying to footprint
- The administrative contacts (maybe the name of the person responsible for the Internet connection or firewall); Attacker may pose as the administrative contact, using social engineering on unsuspecting users in the organization. An attacker will send spoofed email messages posing as the administrative contact to a gullible user. Many user will change their password to whatever you like, as long as it looks like the request is being sent from a trusted technical support person.
- Voice and fax numbers: Enormous help when doing dial-in penetration review (war dialers).
- Record creation and modification dates indicate how accurate the information is (all info may be out-of-date).
- Authoritative DNS servers: Primary, secondary, tertiary, etc (used for our DNS interrogation). Try to use the network range listed as a starting point for our network query of the ARIN database.

Try to discover other domains that a given DNS server is authoritative:

1. Execute a whois query on your domain with a “\$ whois targetdomain.com”
2. Locate the first DNS server
3. Execute a whois query on that DNS server using “\$ whois x.x.x.x”
4. Locate the HST record for the DNS server
5. Execute a whois query with the server directive using \$ whois “server NS9999-HST”

Network query

Begin a network query to determine if a real network is associated with our target domain (use ARIN database b/c InterNIC database only contains domain-related information).

NOTE: 10.10.10.x = dummy IP

\$ whois 10.10.10.0@whois.arin.net (the @ option allows you to specify an alternative database); it maybe possible to use the ‘-a’ option to specify the ARIN database (i.e. BSD servers); Web-based query mechanism: <http://www.arin.net/whois/arinwhois.html>

Results should indicate the main backbone provider, and the type of network assigned to organization (and range).

BOOK: *TCP/IP Illustrated Volume 1* by Richard Stevens

If all information matches up, we could conclude that it is a valid network owned by the organization.

POC query

The administrative contact may be the administrative contact for multiple organizations, so perform a POC query. You may uncover a domain that you were unaware of.

```
$ whois XYZ30
```

Also search for @domain.com to obtain a listing of all mail addresses for a given domain.

```
$ whois "@domain.com"@whois.internic.net
```

For more info on whois queries, see RFC 954 – NICKNAME/WHOIS, or for a complete help reference, do:

```
$ whois ?
```

To obtain the entire membership list of a group or organization, or a list of all authorized users of a host, precede the name of the host or organization by an asterisk, i.e. *SRI-NIC (this may take a while if there are a lot of members). You may use an exclamation point and asterisk, or a period and asterisk together.

Countermeasure: Public Database Security

Ensure that the information listed in the database is accurate; update contacts as necessary (if someone leaves the organization, and they are still listed as a contact, they can change the information).

Use toll-free numbers (or a number not in the organizations phone exchange) for listed phone numbers to prevent war dialing.

Consider using a fictitious contact for the administrative contact (if an employee receives an email from this fictitious person, it could tip you off to the attacker).

Do not use the default authentication method (the FROM field via email) to authenticate the domain registrant's identity on InterNIC. Anyone can trivially forge an email address and change the information associated with your domain. They could redirect all traffic to an alternate domain.

III. DNS Interrogation

After identifying all associated domains, you can begin to query the DNS.

Zone transfers

One of the most serious misconfigurations a system administrator can make is allowing untrusted Internet users to perform a DNS zone transfer.

zone transfer – allows a secondary master server to update its zone database from the primary master. This provides for redundancy when running DNS, should the primary name server become unavailable.

Generally, a DNS zone transfer only needs to be performed by secondary master DNS servers. Many DNS servers are misconfigured to provide a copy to anyone who asks.

If the DNS server does not use a public/private DNS mechanism to segregate the external DNS information (public) from the internal (private), then the internal hostnames and IP addresses could be

disclosed to an attacker. This would be like giving the attacker a blueprint of the organization's internal network.

Zone transfer methods:

nslookup:

```
$ nslookup
```

```
> server 10.10.10.2
```

```
> set type=any
```

```
> ls -d domain.net. > /tmp/zone_out
```

To get all the DNS records we want, change the server to the target's primary DNS server.

Set the record type to *any*. This will allow you to pull any DNS records available (man nslookup) for a complete list.

Use the "ls" option to list all the associated records for the domain. The `-dai` switch is used to list all records for the domain. We append a "." to the end to signify a fully qualified domain name—however, you can leave this off most times.

Redirect output to the file `/tmp/zone_out` so we can manipulate the output later.

Review the output.

The "A" record denotes the IP address of the system name located to the right.

The "HINFO" record identifies the platform or OS running (see RFC-952). These are not needed, but provide a wealth of information to an attacker. An attacker could use `grep`, `sed`, `awk`, or `Perl` to find machines running whatever OS the attacker knows best.

```
grep -i irix zone_out |wc -l (maybe 1 instead of "1")
```

Search for test systems. These normally don't have many security features enabled, often have easily guessed passwords, and administrators tend not to notice or care who logs in to them. Use `grep` to search for "test."

```
grep -i test /tmp/zone_out |wc -l (maybe 1 instead of "1")
```

Use this data to zero in on systems with known vulnerabilities.

This method only queries one name server at a time. You have to perform the same tasks for all name servers that are authoritative for the target domain.

If there are subdomains, you would have to perform the same type query for each subdomain.

If you receive a message stating that you can't list the domain or that the query was refused, this means that the server was configured to disallow zone transfers from unauthorized users. If there are multiple DNS servers, you may be able to find one that allows zone transfers.

Tools to speed the process:

- host
- Sam Spade

- axfr
- dig

host:

```
host -l domain.net
```

```
host -l -v -t any domain.net
```

```
host -l domain.net |cut -f 4 -d " " > /tmp/ip_out  
(to feed the IP address to a shell script)
```

axfr:

one of the best tools for performing zone transfers

get it at: <ftp://ftp.trinux.org/pub/trinux/tools/netmap/axfr-0.5.2.tar.gz> (by Gaius)

Will recursively transfer zone information, and create a compressed database of zone and host files for each domain queried. Also, you can pass top-level domains like *com* and *edu* to get all the domains associated with com and edu, respectively. However, this is not recommended.

```
$ axfr domain.net
```

To query the axfr database for information you just obtained:

```
$ axfrcat domain.net
```

Determine Mail Exchange (MX) Records

Determining where mail is handled is a great starting place to locate the target organization's firewall network. Often in a commercial environment, mail is handled on the same system as the firewall, or at least on the same network.

Use host to help harvest even more information.

```
$ host domain.net
```

Without any parameters, host will resolve *A* records first, then *MX* records. The information should cross-reference with the whois ARIN search.

Countermeasure: DNS Security

Reduce the amount of information available via the Internet.

Restrict zone transfers to only authorized servers. In modern versions of BIND, the *xfernets* directive in the *named.boot* file can be used to enforce the restrictions. To restrict zone transfers in Microsoft's DNS, you can use the *Notify* option (see <http://support.microsoft.com/support/kb/articles/q193/8/37.asp> for more information). For other name servers, consult the documentation to find out how to restrict zone transfers.

On the network side, you could reconfigure the firewall or packet-filtering router to deny all unauthorized inbound connections to TCP port 53. Since name lookup requests are UDP and zone transfer

requests are TCP, this will thwart a zone transfer attempt.

Also, you can set your access control device or intrusion detection system (IDS) to log this information as a potential hostile activity.

Restricting zone transfers will increase the time necessary for attackers to probe for IP addresses and hostnames. However, attackers could still perform lookups manually against all IP addresses for a given net block. Therefore, configure external name servers to provide information only about systems directly connected to the Internet.

External name servers should never be configured to divulge internal network information.

Don't use HINFO records.

IV. Network Reconnaissance

Now that we have determined potential networks, we can attempt to determine their network topology as well as potential access paths into the network.

Tracerouting

tracert (tracert – WinNT):

A hop counter; comes with most Unices; or get it at <ftp://ftp.ee.lbl.gov/traceroute.tar.Z>

Use it to determine the exact path that our packets are taking. It may allow you to discover the network topology employed by the target network, in addition to identifying access control devices (application-based or packet-packet-filtering routers) that may be filtering our traffic.

```
$ traceroute domain.net
traceroute to domain.net (10.10.10.1), 30 hops max, ...
 1 gate2 (192.168.10.1) 5.391 ms 5.107 ms 5.559 ms
 2 rtr1.bigisp.net (10.10.12.13) 33.374 ms 33.443 ms ...
 3 rtr2.bigisp.net (10.10.12.14) 35.100 ms 34.427 ms ...
 4 hssitrt.bigisp.net (10.11.31.14) 43.030 ms 43.941 ...
 5 gate.domain.net (10.10.10.1) 43.803 MS 44.041 MS ...
```

Notice the path of the packets leaving the router (gate2) and traveling three (2-4) to the final destination. The packets go through various hops without being blocked.

From our earlier work, we know that the MX record for domain.net points to gate.domain.net. Thus, we can assume this is a live host, and that the hop before it (4) is the firewall, or it could be a simple packet-filtering device router—we are not sure yet.

Generally, once you hit a live system on a network, the system before it is usually a device performing routing functions (for example, a router or firewall).

This is a very simplistic example; in a complex environment, there may be multiple routing paths, that is, routing devices with multiple interfaces (for example, Cisco 7500 series router). Moreover, each

interface may have different Access Control Lists (ACLs) applied.

In many cases, some interfaces will pass your traceroute requests, while others will deny it because of the ACL applied. Thus, it is important to map your entire network using traceroute. After you traceroute to multiple systems on the network, you can begin to create a network diagram that depicts the architecture of the Internet gateway, and location of devices that are providing access control functionality (an *access path diagram*).

Most flavors of traceroute in UNIX default to sending User Datagram Protocol (UDP) packets, with the option of using Internet Control Messaging Protocol (ICMP) packets with the `-I` switch. In Windows NT, however, the default behavior is to use ICMP *echo request* packets. Thus, your mileage may vary using each tool if the site blocks UDP versus ICMP and vice versa.

The `-g` option allows users to specify loose source routing. Thus, if you believe the target gateway will accept source-routed packets (a cardinal sin), you might try to enable this option with the appropriate hop pointers (run `man traceroute` in UNIX for more information).

There are several other switches that may allow you to bypass access control devices during our probe. The `-p n` option allows you to specify a starting UDP port number (`n`) that will be incremented by 1 when the probe is launched. Thus, we will not be able to use a fixed port number without some modifications to traceroute. There is a patch by Michael Schiffman that adds the `-S` switch to stop port incrementation for traceroute version 1.4a5. This allows you to force every packet we send to have a fixed port number, in the hopes the access control device will pass this traffic. A good starting port number would be UDP port 53 (DNS queries). Since many sites allow inbound DNS queries, there is a high probability that the access control device will allow our probes through.

Additionally, if you send a probe to a system that has UDP port 53 listening, you will not receive a normal ICMP unreachable message back. Thus, you will not see a host displayed when the packet reaches its ultimate destination.

VisualRoute, <http://www.visualroute.com> is a graphical tool to perform tracerouting. It provides a graphical depiction of each network hop and integrates this with whois queries. It is appealing to the eye; however, it does not scale well for large-scale network reconnaissance.

There are additional techniques that will allow you to determine specific ACLs that are in place for a given access control device. *Firewall port scanning* is one such technique.

Countermeasure: Thwarting Network Reconnaissance

Many commercial IDS programs will detect this type of network reconnaissance. Alternatively, you can use free programs like `tdetect` from Vadim Lolontsov (<ftp://ftp.deva.net/pub/sources/networking/ids/tdect-0.2.tar.gz>). It is a simple program that will detect and log via *syslog* any UDP and ICMP traceroute packets with a TTL field equal to 1.

If someone traceroutes you, Humble from Rhino9 developed a program called RotoRouter (<ftp://coast.cs.purdue.edu/pub/tools/unix/trinux/netmon/rr-1.0.tgz>). This utility is used to log incoming traceroute requests and generate fake responses.

Depending on your site's security paradigm, you may be able to configure your border routers to limit ICMP and UDP traffic to specific systems, thus minimizing your exposure.

V. Summary

The above discussion was limited to common tools and techniques; new tools are released daily. You may have the daunting task of footprinting tens or hundreds of domains. Use a combination of shell and EXPECT scripts or Perl programs to automate as many tasks as possible. There are many attackers well schooled in performing network reconnaissance activities without ever being discovered, and they are suitably equipped. Thus, it is important to remember to minimize the amount and types of information leaked by your Internet presence and to implement vigilant monitoring.