

Modern Trends in Network Fingerprinting

SecTor [11.21.07]



Jay Graver
Ryan Poppa

// Fingerprinting Topics

- Why, What, Who & How?
- Tools in action
- Why Tools Break
- Tools EOL
- New Approaches
- New Tool

// Why Fingerprint?

□ WhiteHat

- needs accurate identification of hosts in a PenTest report

□ BlackHat

- reconnaissance

□ SysAdmins

- track down and identify new services or hosts when they appear on their network

// What is a Fingerprint?

- Looking at something *common* ...




```
192.168.2.187:8004 192.168.2.187 [152]
 48 54 54 50 2f 31 2e 31  20 32 30 30 20 4f 4b 0d  HTTP/1.1 200 OK.
 0a 43 6f 6e 6e 65 63 74  69 6f 6e 3a 20 63 6c 6f  .Connection: clo
 73 65 0d 0a 41 6c 6c 6f  77 3a 20 4f 50 54 49 4f  se..Allow: OPTIO
 4e 53 2c 20 47 45 54 2c  20 48 45 41 44 2c 20 50  NS, GET, HEAD, P
 4f 53 54 0d 0a 43 6f 6e  74 65 6e 74 2d 4c 65 6e  OST..Content-Len
 67 74 68 3a 20 30 0d 0a  44 61 74 65 3a 20 46 72  gth: 0..Date: Fr
 69 2c 20 30 32 20 4e 6f  76 20 32 30 30 37 20 32  i, 02 Nov 2007 2
 32 3a 32 35 3a 31 38 20  47 4d 54 0d 0a 53 65 72  2:25:18 GMT..Ser
 76 65 72 3a 20 6c 69 67  68 74 74 70 64 2f 31 2e  ver: lighttpd/1.
 34 2e 31 35 0d 0a 0d 0a                                     4.15....
```

// What is a Fingerprint?

□ ... and finding something *unique*



```
192.168.2.187:8004 192.168.2.187 [152]
 48 54 54 50 2f 31 2e 31 20 32 30 30 20 4f 4b 0d HTTP/1.1 200 OK.
 0a 43 6f 6e 6e 65 63 74 69 6f 6e 3a 20 63 6c 6f .Connection: clo
 73 65 0d 0a 41 6c 6c 6f 77 3a 20 4f 50 54 49 4f se..Allow: OPTIO
 4e 53 2c 20 47 45 54 2c 20 48 45 41 44 2c 20 50 NS, GET, HEAD, P
 4f 53 54 0d 0a 43 6f 6e 74 65 6e 74 2d 4c 65 6e OST..Content-Len
 67 74 68 3a 20 30 0d 0a 44 61 74 65 3a 20 46 72 gth: 0..Date: Fr
 69 2c 20 30 32 20 4e 6f 76 20 32 30 30 37 20 32 i, 02 Nov 2007 2
 32 3a 32 35 3a 31 38 20 47 4d 54 0d 0a 53 65 72 2:25:18 GMT..Ser
 76 65 72 3a 20 6c 69 67 68 74 74 70 64 2f 31 2e ver: lighttpd/1.
 34 2e 31 35 0d 0a 0d 0a 4.15....
```

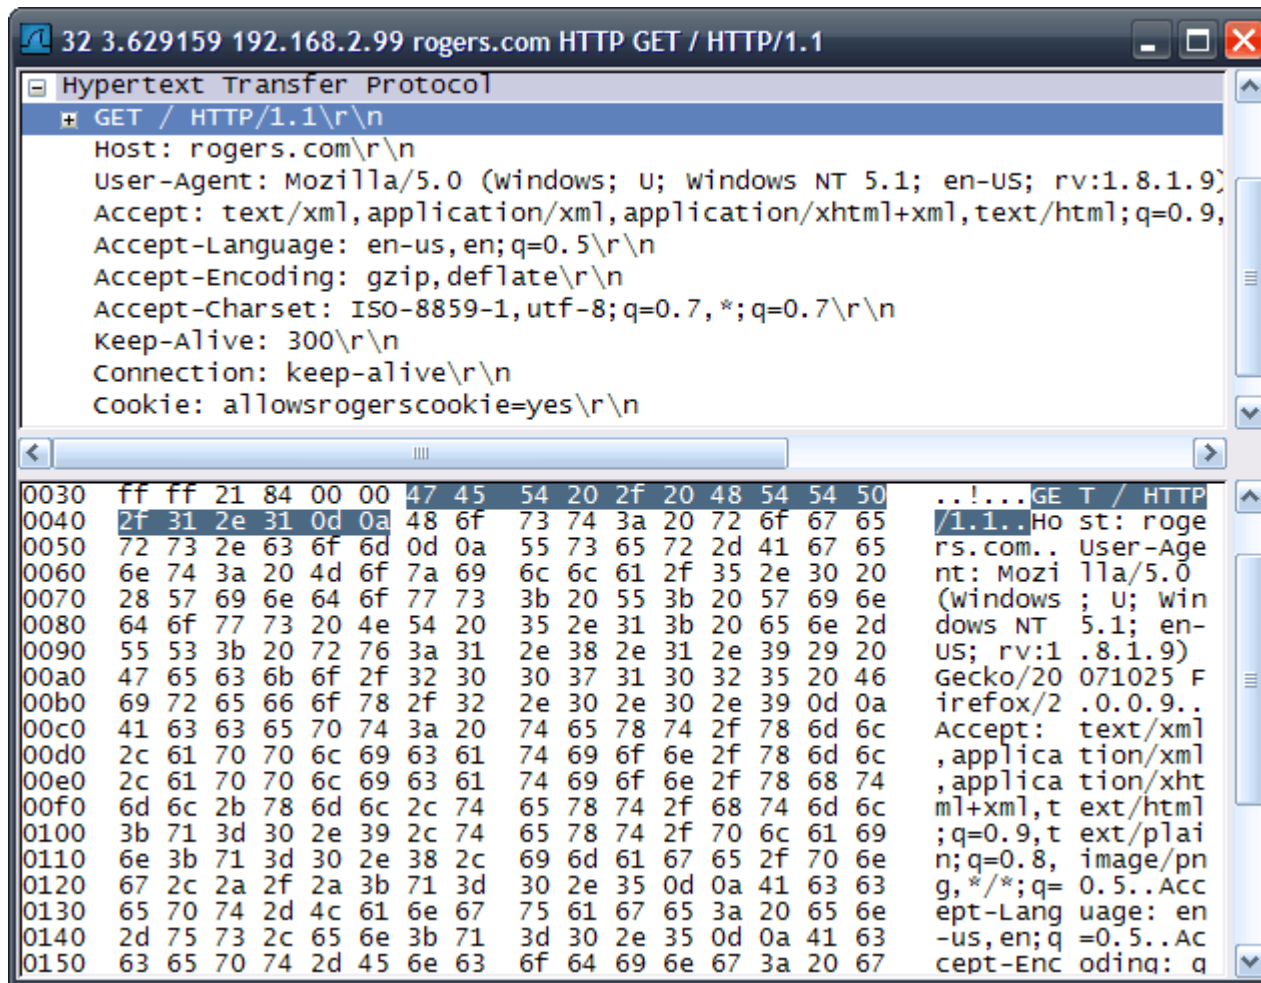


// Who Fingerprints?

- Primarily Network Security Researchers
- Names like:
 - Fyodor
 - Ofir Arkin
 - Michael Zalewski
 - THC (The Hacker's Choice)
 - Jeremiah Grossman
 - Saumil Shah

// How to Fingerprint?

Request Data from Hosts



```
32 3.629159 192.168.2.99 rogers.com HTTP GET / HTTP/1.1
Hypertext Transfer Protocol
GET / HTTP/1.1\r\n
Host: rogers.com\r\n
User-Agent: Mozilla/5.0 (windows; U; windows NT 5.1; en-US; rv:1.8.1.9)
Accept: text/xml,application/xml,application/xhtml+xml,text/html;q=0.9,
Accept-Language: en-us,en;q=0.5\r\n
Accept-Encoding: gzip,deflate\r\n
Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7\r\n
Keep-Alive: 300\r\n
Connection: keep-alive\r\n
Cookie: allowsrogerscookie=yes\r\n
..!...GE T / HTTP
/1.1..Ho st: roge
rs.com.. User-Age
nt: Mozi lla/5.0
(windows ; U; win
dows NT 5.1; en-
US; rv:1 .8.1.9)
Gecko/20 071025 F
irefox/2 .0.0.9..
Accept: text/xml
,application/xml
,application/xht
ml+xml,t ext/html
;q=0.9,t ext/plai
n;q=0.8, image/pn
g,*/*;q= 0.5..Acc
ept-Lang uage: en
-us,en;q =0.5..Ac
cept-Enc odinq: q
```

// How to Fingerprint?

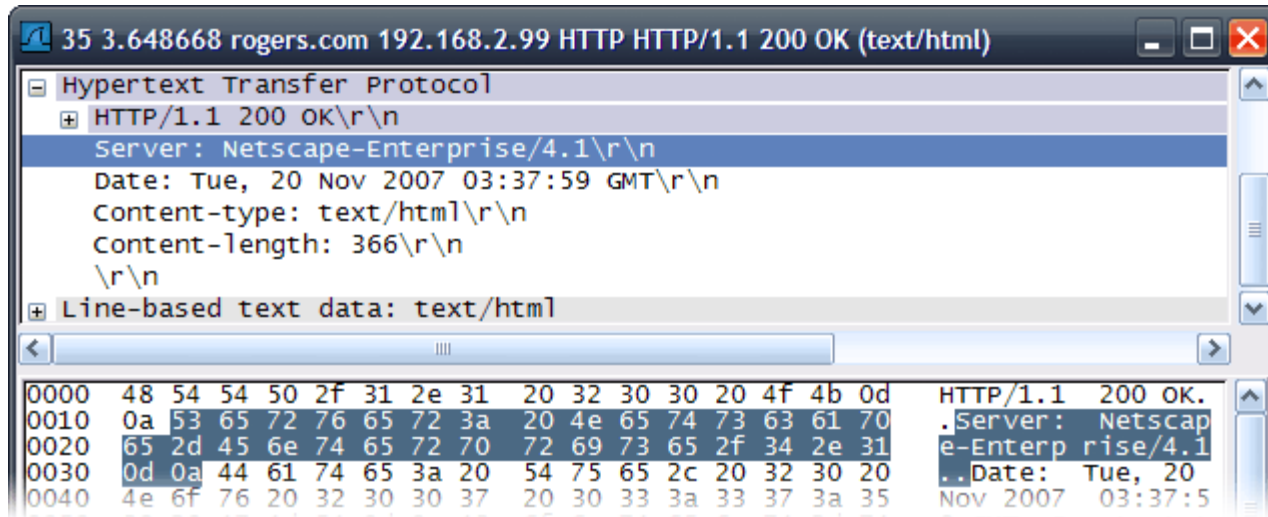
❑ Collect data from Host 1

```
35 3.648668 rogers.com 192.168.2.99 HTTP HTTP/1.1 200 OK (text/html)
Hypertext Transfer Protocol
  HTTP/1.1 200 OK\r\n
  Server: Netscape-Enterprise/4.1\r\n
  Date: Tue, 20 Nov 2007 03:37:59 GMT\r\n
  Content-type: text/html\r\n
  Content-length: 366\r\n
  \r\n
  Line-based text data: text/html
0000 48 54 54 50 2f 31 2e 31 20 32 30 30 20 4f 4b 0d HTTP/1.1 200 OK.
0010 0a 53 65 72 76 65 72 3a 20 4e 65 74 73 63 61 70 .Server: Netscap
0020 65 2d 45 6e 74 65 72 70 72 69 73 65 2f 34 2e 31 e-Enterp rise/4.1
0030 0d 0a 44 61 74 65 3a 20 54 75 65 2c 20 32 30 20 ..Date: Tue, 20
0040 4e 6f 76 20 32 30 30 37 20 30 33 3a 33 37 3a 35 Nov 2007 03:37:5
0050 39 20 47 4d 54 0d 0a 43 6f 6e 74 65 6e 74 2d 74 9 GMT..C ontent-t
0060 79 70 65 3a 20 74 65 78 74 2f 68 74 6d 6c 0d 0a ype: tex t/html..
0070 43 6f 6e 74 65 6e 74 2d 6c 65 6e 67 74 68 3a 20 Content- length:
0080 33 36 36 0d 0a 0d 0a 3c 68 74 6d 6c 3e 0d 0a 3c 366....< html>..<
0090 68 65 61 64 3e 0d 0a 3c 74 69 74 6c 65 3e 52 6f head>..< title>Ro
00a0 67 65 72 73 2e 63 6f 6d 3c 2f 74 69 74 6c 65 3e gers.com </title>
00b0 0d 0a 3c 73 63 72 69 70 74 20 74 79 70 65 3d 22 ..<scrip t type="
00c0 74 65 78 74 2f 6a 61 76 61 73 63 72 69 70 74 22 text/jav ascript"
00d0 20 73 72 63 3d 22 2f 6a 73 2f 72 6f 67 65 72 73 src="/j s/rogers
00e0 2e 6a 73 22 3e 3c 2f 73 63 72 69 70 74 3e 0d 0a .js"></s cript>..
00f0 20 20 20 20 3c 73 63 72 69 70 74 3e 20 20 20 20 <scr ipt>
0100 20 20 20 20 20 20 20 20 0d 0a 20 20 20 20 20 ..
```

Frame (420 bytes) Reassembled TCP (501 bytes)

// How to Fingerprint?

□ Transform this



```
35 3.648668 rogers.com 192.168.2.99 HTTP HTTP/1.1 200 OK (text/html)
Hypertext Transfer Protocol
+ HTTP/1.1 200 OK\r\n
Server: Netscape-Enterprise/4.1\r\n
Date: Tue, 20 Nov 2007 03:37:59 GMT\r\n
Content-type: text/html\r\n
Content-length: 366\r\n
\r\n
+ Line-based text data: text/html
0000 48 54 54 50 2f 31 2e 31 20 32 30 30 20 4f 4b 0d HTTP/1.1 200 OK.
0010 0a 53 65 72 76 65 72 3a 20 4e 65 74 73 63 61 70 .Server: Netscap
0020 65 2d 45 6e 74 65 72 70 72 69 73 65 2f 34 2e 31 e-Enterp rise/4.1
0030 0d 0a 44 61 74 65 3a 20 54 75 65 2c 20 32 30 20 ..Date: Tue, 20
0040 4e 6f 76 20 32 30 30 37 20 30 33 3a 33 37 3a 35 Nov 2007 03:37:5
```

□ Into this

Netscape 0:0:2:0:1:1:0:0:1.1:301:2:0:0:0:0:0:0:
Date,Content-length,Content-type,Location,Connection

// What has been Fingerprinted?

- ❑ IP Stack Operating Systems
- ❑ SMTP Mail Servers
- ❑ FTP File Servers
- ❑ NTP Time Servers
- ❑ HTTP Web Servers
- ❑ DNS Name Servers
- ❑ Client Apps Web Browsers

// Past Tools

- ❑ httpprint
- ❑ Amap
- ❑ queso
- ❑ Xprobe
- ❑ Nessus
- ❑ nmap
- ❑ HMAP
- ❑ smtpscan

// httpprint v301



- ❑ Released by Saumil Shah's Net-Square



- ❑ Freeware (not open source)

- ❑ Uses 23 Sendcases

- ❑ Database: 111 Servers

- ❑ Updated: December 2005

A light purple, multi-pointed starburst graphic with a black outline, containing text.

Voting Algorithm
Confidence Metrics


// Amap v5.2



- ❑ van Hauser and DJ RevMoon of THC (The Hacker's Choice)
- ❑ Open source (GPLish)
- ❑ Uses 30 Sendcases (only 1 HTTP)
- ❑ Database: 346 Servers (37 HTTP)
- ❑ Updated January 2006

// QueSO v980922

- ❑ Jordi Murgó of Apostols.org
- ❑ "¿Que Sistema Operativo?"
- ❑ Open source (GPL)
- ❑ Uses 7 Sendcases
- ❑ Database: 96 OSes
- ❑ Updated: September 1998




The first OS
fingerprinter

// Xprobe2 v0.3

- ❑ Ofir Arkin of Sys-Security
 - Fyodor Yarochkin & Meder Kydyraliev




- ❑ Open source (GPL)
- ❑ Database: 224 OSes
- ❑ Updated: July 2005



// Current Tools [2006 / 2007]

- nmap – OS and Application Identification
- SinFP – Active OS Identification
- p0f – Passive OS Identification
- fpdns – DNS Identification

// nmap v4.23RC1



- ❑ Fyodor of Insecure.Org

- ❑ Open source (GPL)


- ❑ 45 Sendcases (2 HTTP)

- ❑ Database: 3871 Applications (1458 HTTP)


- ❑ Updated: November 2007



// Modern nmap



- nmap has incorporated many of the techniques used by past tools



// What to Test?

- IP Stack Operating Systems
- SMTP Mail Servers
- FTP File Servers
- NTP Time Servers
- **HTTP** **Web Servers**
- DNS Name Servers
- Client Apps Web Browsers

// Why Focus on HTTP?

- ❑ Number of available targets online
- ❑ Management interfaces
- ❑ Single Packet transactions
- ❑ The variety of Web Server software available

“The interweb was born and poof we lost 65,000 ports” “Port 80 and 443 and most people think that is what the Internet is”

- Bruce Potter (DefCon 15)

// Fingerprint Shootout Targets



v7.0



v2.0.61



v0.6.17



v1.4.18



v2.25b

Welcome to the year 2007

// Best Tools for the Job

□ Amap v5.2




Jan 2006

□ httpprint 0.301



Dec 2005

□ nmap 4.23RC1



Nov 2007

// Microsoft IIS 7.0



[Server Header]

HTTP/1.1 200 OK

Content-Length: 689

Content-Type: text/html

Last-Modified: Thu, 08 Nov 2007 19:52:52 GMT

Accept-Ranges: bytes

ETag: "2cd9df24022c81:0"

Server: Microsoft-IIS/7.0

X-Powered-By: ASP.NET

Date: Tue, 20 Nov 2007 02:32:22 GMT

Connection: close


// Microsoft IIS 7.0



http-apache-2 / http-iis / webmin



Microsoft-IIS/6.0 (93%)



Microsoft IIS webserver 7.0

// Apache 2.0.61 [ServerTokens Full]



[httpd.conf]

ServerTokens Full

[Server Header]

HTTP/1.1 200 OK

Date: Wed, 14 Nov 2007 12:05:41 GMT

Server: Apache/2.0.61 (Unix)

Content-Location: index.html.en

Vary: negotiate,accept-language,accept-charset

TCN: choice

Last-Modified: Sun, 21 Nov 2004 14:35:21 GMT

ETag: "33072-5b0-a64a7c40;33088-961-a64a7c40"

Accept-Ranges: bytes

Content-Length: 1456

Connection: close

Content-Type: text/html

Content-Language: en

Expires: Wed, 14 Nov 2007 12:05:41 GM

// Apache 2.0.61 [ServerTokens Prod]



[httpd.conf]

ServerTokens Prod

[Server Header]

HTTP/1.1 200 OK

Date: Wed, 14 Nov 2007 12:14:51 GMT

Server: Apache

Content-Location: index.html.en

Vary: negotiate,accept-language,accept-charset

TCN: choice

Last-Modified: Sun, 21 Nov 2004 14:35:21 GMT

ETag: "33072-5b0-a64a7c40;33088-961-a64a7c40"

Accept-Ranges: bytes

Content-Length: 1456

Connection: close

Content-Type: text/html

Content-Language: en

Expires: Wed, 14 Nov 2007 12:14:51 GMT


// Apache 2.0.61 [ServerTokens Prod]



□ http-apache-2 / webmin



□ Apache/2.0.x (84%)



□ Apache httpd

// lighttpd v1.4.18



[lighttpd.conf]

#server.tag = "lighttpd"

[Server Header]

HTTP/1.1 404 Not Found

Content-Type: text/html

Content-Length: 345

Date: Wed, 14 Nov 2007 12:31:51 GMT

Server: lighttpd/1.4.18

```
// lighttpd v1.4.18
```

[lighttpd.conf]

```
server.tag = "lighttpd"
```

[Server Header]

```
HTTP/1.1 404 Not Found
```

```
Content-Type: text/html
```

```
Content-Length: 345
```

```
Date: Wed, 14 Nov 2007 12:31:51 GMT
```

```
Server: lighttpd
```




// lighttpd v1.4.18



http-apache-2 / webmin



Apache-Tomcat/4.1.29 (58%)



lighttpd


// nginx v0.6.17



webmin



Microsoft-IIS/6.0 (48%)



nginx http proxy 0.6.17


// thttpd v2.25b 29dec2003



webmin



Microsoft-IIS/6.0 (51%)



thttpd 2.25b 29dec2003

// Fingerprinting Tools EOL

- ❑ Database Aging
- ❑ Static Sendcases
- ❑ Fingerprint / Database Corruption
- ❑ Smoke, Mirrors & Banners
- ❑ Rise of Obfuscation

// Database Aging

- ❑ Database of known Servers stops being updated – no longer has the latest releases
- ❑ Within 2 years, the accuracy of the fingerprinting tool will noticeably suffer

// Static Sends

- ❑ Tools which use a “strict matching” strategy prevent any changes to the sendcases
- ❑ Only new *results* can be added
- ❑ No new questions can be asked without invalidating all previous results

// Static Sendcases

- ❑ New sendcases would need to be tested against all historical servers in database
- ❑ IIS 7 and Apache 2.2.6 are out - do you have CERN or NCSA httpd around?

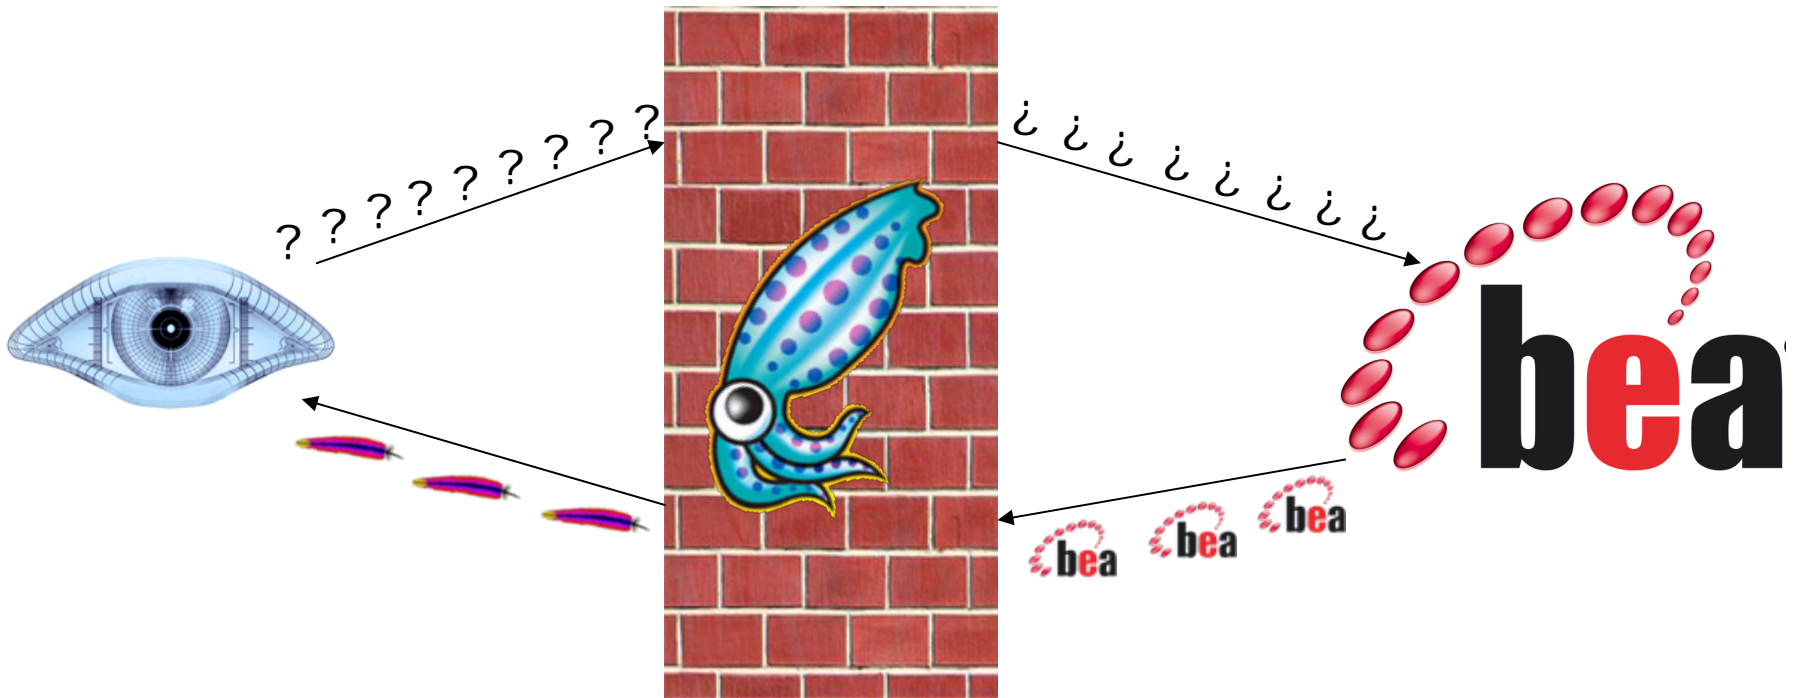


NCSA
HTTPd



// Fingerprint Corruption

- ❑ Submitted data is flawed
- ❑ Network normalization impacts a tool's accuracy and fingerprint collection



// Database Corruption

- ❑ nmap's open nature makes it vulnerable to incorrect information being added
- ❑ Unless you independently verify new fingerprints against several confirmed targets you can not 100% trust them as accurate
- ❑ Configuration Issues and Network Conditions multiply valid fingerprints
- ❑ Section 2.12 of "Present and Future of Xprobe2" - O. Arkin

// Smoke Mirrors & Banners

- ❑ Much of the Accuracy you see is little more than blind faith
- ❑ Banners are becoming highly mutable
- ❑ Several Linux Distros now ship with light obfuscation by default
- ❑ ServerMask can obfuscate IIS

// Obfu What?

ob'fus·ca'tion (n.)

- 1. to make obscure or unclear*
- 2. confusion resulting from failure to understand*

// In simple terms

□ Hiding in plain sight



// Why Hide?

- ❑ Throw off script kiddies and botz
- ❑ Throw off PenTesters and Auditors
- ❑ Sweep Vulns under the rug
- ❑ Added level of security through obscurity
- ❑ Why Not???

// How to Hide

□ Apache

- Modify `ap_release.h` / `httpd-defaults.conf` (compile time)
- `ServerTokens` in `httpd.conf` (runtime)

□ lighttpd

- `server-tag` option in `lighttpd.conf` (runtime)

□ Bannerless BEA WebLogic 7.1 SP6/8.1 SP4

□ ServerMask has lots of options for IIS

// Fingerprinting Obfuscated Targets



□ lighttpd pretending to be

Server: Apache/2.0.52 (Red Hat)

mod_perl/1.99_16 Perl/v5.8.5 DAV/2

PHP/4.3.9 mod_python/3.1.3 Python/2.3.4

mod_ssl/2.0.52 OpenSSL/0.9.7a



□ Apache with

Server: AAAAAA/0.0.00

// Fingerprinting Obfuscated Targets



- thttpd pretending to be
Server: Microsoft-IIS/7.0



- nginx pretending to be
Server: GFE/1.3


// Fakepache/2.0.52



http-apache-2 / webmin



Microsoft-IIS/6.0 (39%)



Apache httpd 2.0.52

// Fakepache/2.0.52




Wrong



Wrong

(39%)



Wrong

// AAAAAA 0.0.00



http-apache-2 / webmin



Apache/2.0.x (84%)



Unknown

// AAAAAA 0.0.00




Good



Good

(84%)



? Nope


// Microsofthttpd-IIS/7.0



http-iis / webmin



Microsoft-IIS/6.0 (51%)




Microsoft IIS webserver 7.0



Wrong



Wrong (51%)



Wrong


// ngFE/1.3



webmin



Microsoft-IIS/6.0 (48%)



Google httpd 1.3 (GFE)



Wrong



Wrong (48%)



Wrong

// Web Servers Unmasked

- ❑ Surely we can do better than that
- ❑ Obfuscation Detection is possible when Banners and highly mutable options are not considered when collecting fingerprints
- ❑ If not banners? Then what?

// Modern Ideas on Fingerprinting

- Dynamic Sendcases
- Response Analysis
- Decision Trees/Graphs
 - Be smart!

// Ideas: Dynamic Sendcases

- Why try all sendcases every time?
- How does one deal with sendcase growth?
 - You can't. It doesn't scale!
 - Also has more points of failure
- Need to dynamically determine what should be sent next
 - But..... How??????



// Ideas: Response Analysis

□ Idea!

- Each response says something not already known
- Analyze each response to decide what should be done next

□ Essentially, it's a Choose Your Own Adventure™ - The Web Server Edition

// Ideas: Trees / Graphs

- Decision Trees seem like the best solution to the problem
 - Have logical path-finding/decision making
 - Can be easily added to (grafting) as new things are found and profiled

- Tree structures do have limitations
 - There is only 1 path

- Therefore, need a structure that is tree-like

// Building a Smart Fingerprinter

- ❑ Nice and all but...How?
- ❑ Let's try an example...
- ❑ How would you describe this pear?



// Building a Smart Fingerprinter

- ❑ Easy enough - Now let's work backwards
- ❑ Assume, you are trying to guess an object (the pear)
- ❑ You know the characteristics of the following objects:
 - Pear
 - Car
 - Cup
 - House
 - Pizza

// Building a Smart Fingerprinter

- What kind of questions would you ask now?

- Remembering the list.
 - Pear
 - Car
 - Cup
 - House
 - Pizza

- How about...
 - Are you edible?

// Building a Smart Fingerprinter

- ❑ So, the only foods on the list that match that criteria are **pizza** and **pear**

- ❑ Then you have the following list:
 - **Pear**
 - Car
 - Cup
 - House
 - **Pizza**

// Building a Smart Fingerprinter

- ❑ The possibilities have been narrowed down
 - Based on the response to the first question

- ❑ The next question would never be:
 - Can a person sit in it?


- ❑ You have already eliminated all objects that would have fallen into this category
 - The answer does not narrow down the result set

// Building a Smart Fingerprinter

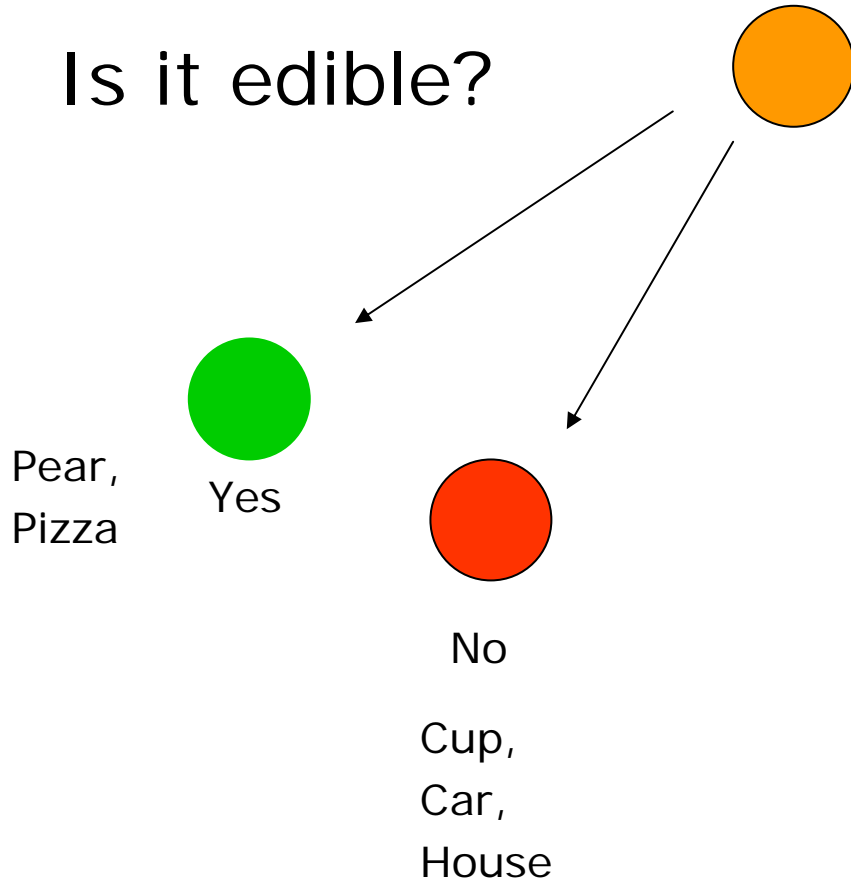
- Essentially, this is the 20-Questions algorithm
 - Eliminate all that it *cannot* be - you are left with what it *is*
 - In reality, it's Information Theory
- All present fingerprinters do not take this into account
 - They will ask all the questions first and then try to decide what it is
 - Even if there is only one possibility left!

// In Visual Form (Graphical)


Is it edible?



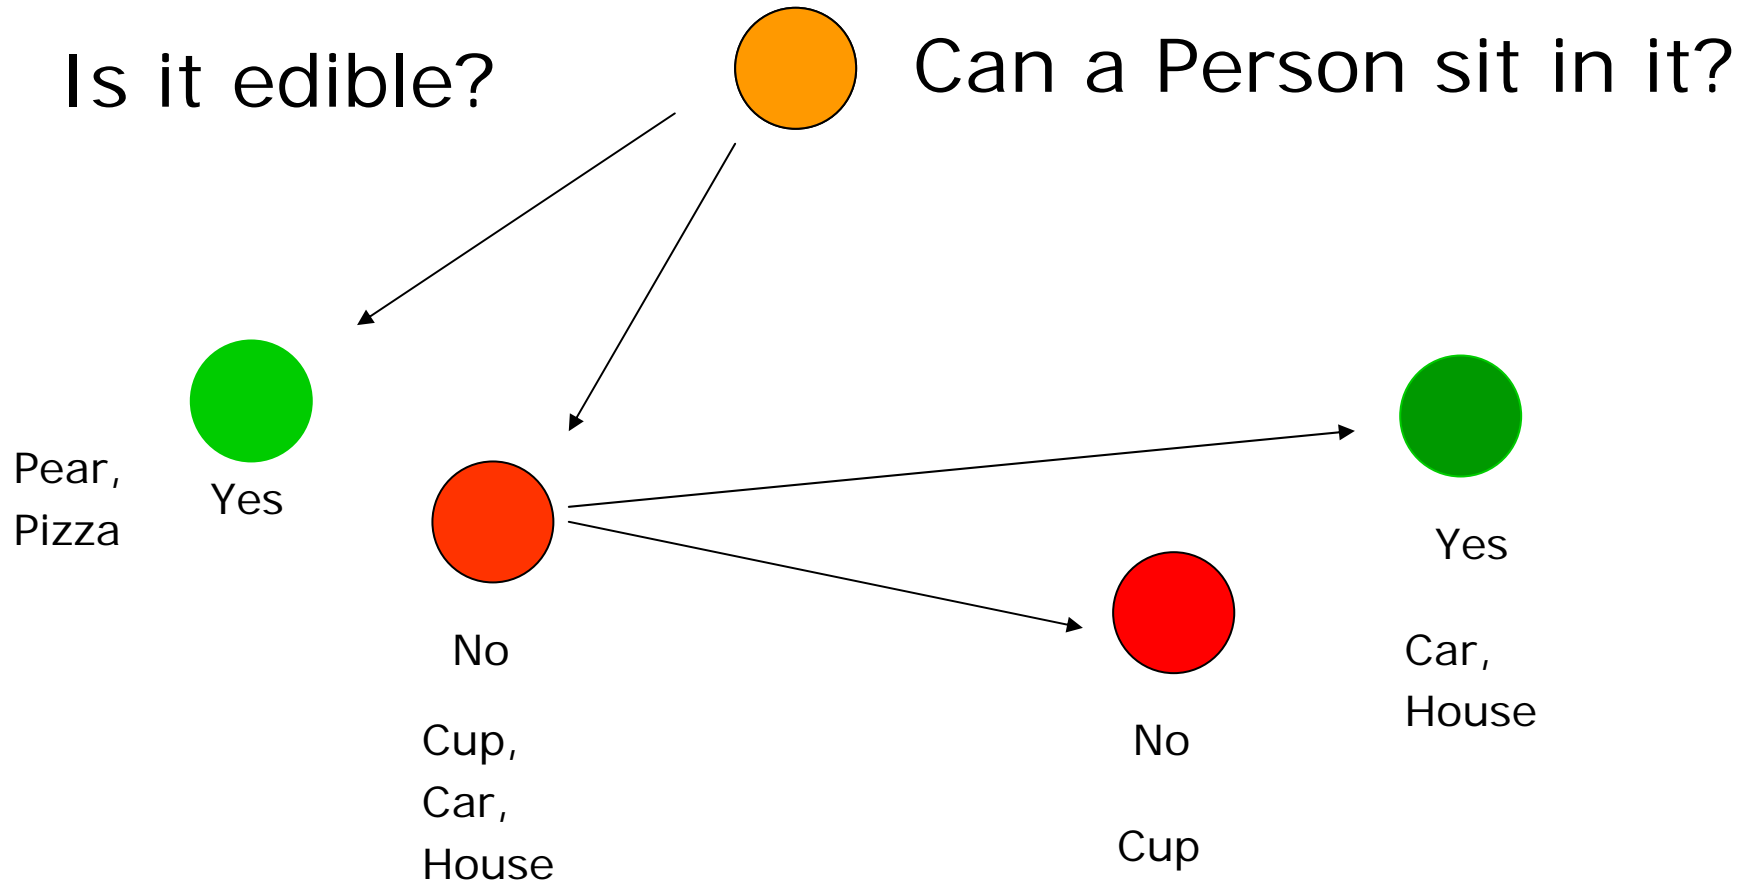
// In Visual Form (Graphical)



// In Visual Form (Graphical)



// In Visual Form (Graphical)



// Building a Smart Fingerprinter

- This method works for any server:
 1. Send a request.
 2. Use the response to eliminate all the servers it cannot possibly be.
 3. If there is only one server left, stop - you know what it is.
 4. Otherwise, choose the next request based on the servers that are left.
 5. Repeat.

// Advantages

- ❑ Minimize the number of questions
 - Theoretically possible to fingerprint in only one request
- ❑ Speed
- ❑ Dynamic data structure
 - The tree like structure is regenerated each time
 - There is no set tree
- ❑ Ability to guess!
 - If something looks like a tree, talks like a tree but smells like a flower?
- ❑ Self-Learning?

// Disadvantages

□ Data

- Acquiring data – *lots* of data

□ Dealing with applications/servers that are unknown

□ What to do if you are wrong?

// Ways to Fingerprint HTTP

□ HTTP Headers

- Existence of Fields (field names)
- Values (tokens)
- Formatting
- Uniqueness
- Ordering

□ HTTP Version support

□ Error Messages

□ Dynamic Webpage content

// Example

OPTIONS * HTTP/1.1 request

■ Apache 2.x:

```
HTTP/1.1 200 OK
Date: Sun, 18 Nov 2007 23:55:06 GMT
Server: Apache/2.0.61
Allow: GET,HEAD,POST,OPTIONS,TRACE
Content-Length: 0
Content-Type: text/plain
```

■ Netscape Enterprise:

```
HTTP/1.1 200 OK
Server: Netscape-Enterprise/3.6 SP3
Date: Sun, 18 Nov 2007 23:57:55 GMT
Content-length: 0
Public: HEAD, GET, PUT, POST
```

// HTTP Status Code

■ Apache 2.x:

HTTP/1.1 200 OK
Date: Sun, 18 Nov 2007 23:55:06 GMT
Server: Apache/2.0.61
Allow: GET,HEAD,POST,OPTIONS,TRACE
Content-Length: 0
Content-Type: text/plain

■ Netscape Enterprise:

HTTP/1.1 200 OK
Server: Netscape-Enterprise/3.6 SP3
Date: Sun, 18 Nov 2007 23:57:55 GMT
Content-length: 0
Public: HEAD, GET, PUT, POST

// HTTP Reason

- Apache 2.x:

```
HTTP/1.1 200 OK
Date: Sun, 18 Nov 2007 23:55:06 GMT
Server: Apache/2.0.61
Allow: GET,HEAD,POST,OPTIONS,TRACE
Content-Length: 0
Content-Type: text/plain
```

- Netscape Enterprise:

```
HTTP/1.1 200 OK
Server: Netscape-Enterprise/3.6 SP3
Date: Sun, 18 Nov 2007 23:57:55 GMT
Content-length: 0
Public: HEAD, GET, PUT, POST
```

- These values are not defined in the RFC, so they could be anything
- Microsoft IIS is the only server that supports internationalization

// HTTP Protocol

■ Apache 2.x:

HTTP/1.1 200 OK

Date: Sun, 18 Nov 2007 23:55:06 GMT

Server: Apache/2.0.61

Allow: GET,HEAD,POST,OPTIONS,TRACE

Content-Length: 0

Content-Type: text/plain

■ Netscape Enterprise :

HTTP/1.1 200 OK

Server: Netscape-Enterprise/3.6 SP3

Date: Sun, 18 Nov 2007 23:57:55 GMT

Content-length: 0

Public: HEAD, GET, PUT, POST

// HTTP Methods Supported

- Apache 2.x:

HTTP/1.1 200 OK

Date: Sun, 18 Nov 2007 23:55:06 GMT

Server: Apache/2.0.61

Allow: GET, HEAD, POST, OPTIONS, TRACE

Content-Length: 0

Content-Type: text/plain

- Netscape Enterprise:

HTTP/1.1 200 OK

Server: Netscape-Enterprise/3.6 SP3

Date: Sun, 18 Nov 2007 23:57:55 GMT

Content-length: 0

Public: HEAD, GET, PUT, POST

- Apache responds with the RFC 2616 Compliant Allow header field, while Netscape responds with Public
- Method ordering in the field AND the spacing between them

// Header Field Capitalization

■ Apache 2.x:

HTTP/1.1 200 OK

Date: Sun, 18 Nov 2007 23:55:06 GMT

Server: Apache/2.0.61

Allow: GET,HEAD,POST,OPTIONS,TRACE

Content-Length: 0

Content-Type: text/plain

■ Netscape Enterprise:

HTTP/1.1 200 OK

Server: Netscape-Enterprise/3.6 SP3

Date: Sun, 18 Nov 2007 23:57:55 GMT

Content-length: 0

Public: HEAD, GET, PUT, POST

- There are many other examples of comparison that you can use

// httpfp

- ❑ HTTP Server Fingerprinter
- ❑ Attempts to fingerprint in only 1 request
- ❑ Only attempts to fingerprint what the web server is
 - Not specifically which version it is, even if we know
 - ❑ ie: Apache instead of Apache 2.0.61

// httpfp Techniques

❑ Functions on a single sendcase

```
OPTIONS * HTTP/1.1\x0d\x0a
Host: %hostname%\x0d\x0a
Connection: Close\x0d\x0a\x0d\x0a
```

- ❑ Banners are provided to the user for informational purposes only
- ❑ Attempt to discover proxies/web application firewalls/load balancers when possible

// Designing httpfp

□ Dynamic Content?

- Could not deal with dynamic content
- Though surprising how useful it can be
 - Most admins do not (or cannot) change error pages
- Mostly thought of as a bad way to go
 - But not as bad as it is thought

□ Multiple Send Statements?

- Outside of the scope of the project

// What httpfp looks for

- ❑ Fingerprints are generated using 20 different criteria
- ❑ This includes:
 - The Existence of certain HTTP Header Fields
 - The Values of HTTP Headers Fields
 - HTTP Version support
 - Error Messages
 - Unique HTTP Header Fields
 - HTTP Header ordering

// Pros & Cons of the Design

- ❑ Single Sendcase
 - Limiting to what can be done
 - What happens when two things look alike?

- ❑ Content Handlers
 - Applications on the back end that can edit content and fields
 - Content Handlers that rewrite requests they are not supposed to handle
 - ❑ .NET, mod_rewrite & mod_forward do this

- ❑ Use of header ordering
 - Inaccurate science. Many things (proxies/content handlers) can manipulate the header order

- ❑ To Guess or Not to Guess?

- ❑ Not 100% accurate

// Why choose this approach?

- Why was it designed this way?
 - It was easier to do it this way
 - Minimizes scope
 - Single request can provide many different unique answers
 - To prove that you can fingerprint accurately with very limited information

// httpfp v1.0

- ❑ Fingerprint Database contains 1620 unique entries
- ❑ Accurately identifies 234 different Web Servers/proxies/web application firewalls

// Real Fingerprint Example

HTTP Status:

```
0:1:2:1:1:1:0:0:11:200:0:0:0:0:0:0:0:
Date,Content-length,Public,Connection
:OK:HEAD,GET,PUT,POST
```



HTTP/1.1 **200** OK

Server: Netscape-Enterprise/3.6 SP3

Date: Sun, 18 Nov 2007 23:57:55 GMT

Content-length: 0

Public: HEAD, GET, PUT, POST

// httpfp performance



□ lighttpd

192.168.2.178 **lighttpd** lighttpd
0:0:1:0:1:1:0:1:10:501:1:0:0:0:0:0:0: Connection,Date,Content-
Length,Accept-Ranges,Content-Type: Not Implemented:



□ lighttpd

192.168.2.178 **lighttpd** Apache/2.0.52 (Red Hat)
mod_perl/1.99_16 Perl/v5.8.5 DAV/2 PHP/4.3.9
mod_python/3.1.3 Python/2.3.4 mod_ssl/2.0.52 OpenSSL/0.9.7a
0:0:1:0:1:1:0:1:10:501:1:0:0:0:0:0:0: Connection,Date,Content-
Length,Accept-Ranges,Content-Type: Not Implemented:

// httpfp performance



□ Apache 2.2.6

192.168.216.126 False None **Apache** AAAAAA
1:0:1:1:1:1:0:0:11:200:1:1:0:0:0:0:0:0: Date,Allow,Content-
Length,Connection,Content-
Type:OK:GET,HEAD,POST,OPTIONS,TRACE



□ Apache

192.168.2.178 **Apache** AAAAAA/0.0.00 (Unix)
1:0:1:1:1:1:0:0:11:200:1:1:0:0:0:0:0:0: Date,Allow,Content-
Length,Connection,Content-
Type:OK:GET,HEAD,POST,OPTIONS,TRACE

// httpfp performance



□ IIS

192.168.216.171 False None **IIS** Microsoft-IIS/5.0
1:1:1:1:1:1:0:1:11:200:0:0:1:0:0:0:0: Date, Connection, Content-Length, Accept-Ranges, Public, Allow: OK: OPTIONS, TRACE, GET, HEAD, DELETE, PUT, POST, COPY, MOVE, MKCOL, PROPFIND, PROPPATCH, LOCK, UNLOCK, SEARCH



□ thttpd

192.168.216.126 False None **thttpd** Microsoft-IIS/7.0
0:0:0:0:1:1:0:1:11:400:1:0:0:0:0:0:1: Content-Type, Date, Accept-Ranges, Connection: Bad Request:

// httpfp performance



192.168.2.178 False None **nginx** nginx/0.6.17
0:0:1:0:1:1:0:0:11:405:1:0:0:0:0:0:0: Date,Content-
Type,Content-Length,Connection: Not Allowed:



192.168.2.178 False None **nginx** GFE/1.3
0:0:1:0:1:1:0:0:11:405:1:0:0:0:0:0:0: Date,Content-
Type,Content-Length,Connection: Not Allowed:

// How Far Can 1 Sendcase Go?

- The amount of information retrieved from 1 sendcase is surprising
 - However, this approach has its limits
 - Possible that you can't decipher between N web servers with just one request
 - You might be left with 2 or 3 servers at the end
 - Certain servers are built to emulate other known web servers (e.g nginx)
 - Certain proxies/load balancers do not like the sendcase and will not allow you to bypass
- Dynamic sendcases will open new possibilities

// Discoveries

□ People still run OLD stuff

- We're not talking 2 years old here - More like 10 years!
- Of ~300K hosts:
 - 13 instances of Apache 1.1
 - 92 instances of Apache 1.2
 - More than the total servers found for WebLogic, Zope, AOLServer, thttpd and Roxen
 - 858! Instances of Apache 1.3 < 1.3.9
 - Mostly on old Cobalt Qube's
 - Vulnerable to a whole ton of stuff
 - 1 instance of CERN 3.0 (1995)

// Obfuscation Analysis

- ❑ Began to scour results for instances of obfuscation
- ❑ Focused on light obfuscation – banner modification or removal
- ❑ We discovered you can detect proxies and firewalls this way

// Rates of Obfuscation for Servers

Server	Total	Obfuscated	%
Apache	204933	5885	2.87%
IIS	85834	1150	1.34%
lighttpd	921	346	37.6%
Netscape	904	105	11.4%
Domino	700	0	0%
WebLogic	72	18	25.35%

// Rates of Obfuscation for Servers

- Why so high for lighttpd and WebLogic?
 - You can configure the server header in lighttpd in configuration file on runtime
 - BEA removed the header from most instances of WebLogic as of WebLogic 7.1 SP6/8.1 SP4 and above.
 - Therefore, if you find a WebLogic header, it's for an old version

- Therefore, if you give people the means to change the header easily, they will do it

- How accurate is Netcraft's Web Server Survey if it does not account for Obfuscated servers?
 - No one knows how they do it

- Caveat: No way to prove that the selection of hostnames is truly random
 - Our selection of 300K hostnames may be biased, but it's a fun comparison

// Netcraft

- Data is from Netcraft's October 2007 survey:

Server	Netcraft %	httpfp %
Apache	47.73	66.37
IIS	37.13	28.07
Sun	1.58	0.03

// Thank you

Questions?

Jay Graver jgraver@gmail.com

Ryan Poppa rpoppa@gmail.com

httpfp <http://www.ncircle.com/labs/>

// References

- ❑ Netcraft October Survey

http://news.netcraft.com/archives/2007/10/11/october_2007_web_server_survey.html

- ❑ Google Server Survey

<http://googleonlinesecurity.blogspot.com/2007/06/web-server-software-and-malware.html>

- ❑ J. Grossman – Identifying Web Servers: A first look into the future – BH 2002

http://www.whitehatsec.com/presentations/Black_Hat_Singapore_2002/BlackHat2002-Singapore.zip

- ❑ HTTPrint

<http://www.net-square.com/httpprint/>

- ❑ O. Arkin “Present and Future of Xprobe2”

http://www.sys-security.com/archive/papers/Present_and_Future_Xprobe2-v1.0.pdf

- ❑ B. Potter “Dirty Secrets of the Security Industry”

<http://video.google.com/videoplay?docid=-4408250627226363306&hl=en>

- ❑ THC’s Amap

<http://freeworld.thc.org/thc-amap/>

// Special Thanks

- Jeff Forristal – HP/SPI Dynamics
- Tyler Reguly – nCircle Network Security