

CSci 4271W
Development of Secure Software Systems
Day 18: Web Security 2

Stephen McCamant
University of Minnesota, Computer Science & Engineering

Outline

- Cross-site scripting, cont'd
- More cross-site risks
- Announcements intermission
- Confidentiality and privacy
- Even more web risks
- Crypto basics

No string-free solution

- For server-side XSS, no way to avoid string concatenation
- Web page will be sent as text in the end
- XSS especially hard kind of injection

Don't deny-list

- Browser capabilities continue to evolve
- Attempts to list all bad constructs inevitably incomplete
- Even worse for XSS than other injection attacks

Filter failure: one-pass delete

- Simple idea: remove all occurrences of `<script>`
- What happens to `<scr<script>ipt>`?

Filter failure: UTF-7

- You may have heard of UTF-8
 - Encode Unicode as 8-bit bytes
- UTF-7 is similar but uses only ASCII
- Encoding can be specified in a `<meta>` tag, or some browsers will guess
- `+ADw-script+AD4-`

Filter failure: event handlers

```
<IMG onmouseover="alert('xss')">
```

- Put this on something the user will be tempted to click on
- There are more than 100 handlers like this recognized by various browsers

Use good libraries

- Coding your own defenses will never work
- Take advantage of known good implementations
- Best case: already built into your framework
 - Disappointingly rare

Content Security Policy

- Added HTTP header, W3C recommendation
- Lets site opt-in to stricter treatment of embedded content, such as:
 - No inline JS, only loaded from separate URLs
 - Disable JS `eval` et al.
- Has an interesting violation-reporting mode

Outline

Cross-site scripting, cont'd
More cross-site risks
Announcements intermission
Confidentiality and privacy
Even more web risks
Crypto basics

HTTP header injection

- Untrusted data included in response headers
- Can include CRLF and new headers, or premature end to headers
- AKA "response splitting"

Content sniffing

- Browsers determine file type from headers, extension, and content-based guessing
 - Latter two for ~1% server errors
- Many sites host "untrusted" images and media
- Inconsistencies in guessing lead to a kind of XSS
 - E.g., "chimera" PNG-HTML document

Cross-site request forgery

- Certain web form on `bank.com` used to wire money
- Link or script on `evil.com` loads it with certain parameters
 - Linking is exception to same-origin
- If I'm logged in, money sent automatically

CSRF prevention

- Give site's forms random-nonce tokens
 - E.g., in POST hidden fields
 - Not in a cookie, that's the whole point
- Reject requests without proper token
 - Or, ask user to re-authenticate
- XSS can be used to steal CSRF tokens

Open redirects

- Common for one page to redirect clients to another
- Target should be validated
 - With authentication check if appropriate
- Open redirect*: target supplied in parameter with no checks
 - Doesn't directly hurt the hosting site
 - But reputation risk, say if used in phishing
 - We teach users to trust by site

Outline

Cross-site scripting, cont'd
More cross-site risks
Announcements intermission
Confidentiality and privacy
Even more web risks
Crypto basics

Web security reading

- The OWASP Top Ten is a web page enumerating the most important web security threats, with advice about what to do about them
- Reading quiz will be due a week from today, Thursday the 28th

Outline

- Cross-site scripting, cont'd
- More cross-site risks
- Announcements intermission
- Confidentiality and privacy
- Even more web risks
- Crypto basics

Site perspective

- Protect confidentiality of authenticators
 - Passwords, session cookies, CSRF tokens
- Duty to protect some customer info
 - Personally identifying info ("identity theft")
 - Credit-card info (Payment Card Industry Data Security Standards)
 - Health care (HIPAA), education (FERPA)
 - Whatever customers reasonably expect

You need to use SSL/TLS

- We have come around to view that more sites need to support HTTPS
 - Special thanks to WiFi, NSA
- If you take credit cards (of course)
- If you ask users to log in
 - Must be protecting something, right?
 - Also important for users of Tor, proxies, etc.

Server-side encryption

- Also consider encrypting data "at rest"
- (Or, avoid storing it at all)
- Provides defense in depth
 - Reduce damage after another attack
- May be hard to truly separate keys
 - OWASP example: public key for website → backend credit card info

Adjusting client behavior

- HTTPS and password fields are basic hints
- Consider disabling autocomplete
 - Usability tradeoff, save users from themselves
 - Finally standardized in HTML5
- Consider disabling caching
 - Performance tradeoff
 - Better not to have this on user's disk
 - Or proxy? You need SSL/TLS

User vs. site perspective

- User privacy goals can be opposed to site goals
- Such as in tracking for advertisements
- Browser makers can find themselves in the middle
 - Of course, differ in institutional pressures

Third party content / web bugs

- Much tracking involves sites other than the one in the URL bar
 - For fun, check where your cookies are coming from
- Various levels of cooperation
- *Web bugs* are typically 1x1 images used only for tracking

Cookies arms race

- Privacy-sensitive users like to block and/or delete cookies
- Sites have various reasons to retain identification
- Various workarounds:
 - Similar features in Flash and HTML5
 - Various channels related to the cache
 - *Evercookie*: store in n places, regenerate if subset are deleted

Browser fingerprinting

- Combine various server or JS-visible attributes passively
 - User agent string (10 bits)
 - Window/screen size (4.83 bits)
 - Available fonts (13.9 bits)
 - Plugin versions (15.4 bits)

(Data from panopticlick.eff.org, far from exhaustive)

History stealing

- History of what sites you've visited is not supposed to be JS-visible
- But, many side-channel attacks have been possible
 - Query link color
 - CSS style with external image for visited links
 - Slow-rendering timing channel
 - Harvesting bitmaps
 - User perception (e.g. fake CAPTCHA)

Browser and extension choices

- More aggressive privacy behavior lives in extensions
 - Disabling most JavaScript (NoScript)
 - HTTPS Everywhere (centralized list)
 - Tor Browser Bundle
- Default behavior is much more controversial
 - Concern not to kill advertising support as an economic model

Outline

Cross-site scripting, cont'd

More cross-site risks

Announcements intermission

Confidentiality and privacy

Even more web risks

Crypto basics

Misconfiguration problems

- Default accounts
- Unneeded features
- Framework behaviors
 - Don't automatically create variables from query fields

Openness tradeoffs

- Error reporting
 - Few benign users want to see a stack backtrace
- Directory listings
 - Hallmark of the old days
- Readable source code of scripts
 - Doesn't have your DB password in it, does it?

Using vulnerable components

- Large web apps can use a lot of third-party code
- Convenient for attackers too
 - OWASP: two popular vulnerable components downloaded 22m times
- Hiding doesn't work if it's popular
- Stay up to date on security announcements

Clickjacking

- Fool users about what they're clicking on
 - Circumvent security confirmations
 - Fabricate ad interest
- Example techniques:
 - Frame embedding
 - Transparency
 - Spoof cursor
 - Temporal "bait and switch"

Crawling and scraping

- A lot of web content is free-of-charge, but proprietary
 - Yours in a certain context, if you view ads, etc.
- Sites don't want it downloaded automatically (*web crawling*)
- Or parsed and user for another purpose (*screen scraping*)
- High-rate or honest access detectable

Outline

Cross-site scripting, cont'd
More cross-site risks
Announcements intermission
Confidentiality and privacy
Even more web risks
Crypto basics

-ography, -ology, -analysis

- Cryptography (narrow sense): designing encryption
- Cryptanalysis: breaking encryption
- Cryptology: both of the above
- Code (narrow sense): word-for-concept substitution
- Cipher: the "codes" we actually care about

Caesar cipher

- Advance three letters in alphabet:
 $A \rightarrow D, B \rightarrow E, \dots$
- Decrypt by going back three letters
- Internet-era variant: rot-13
- Easy to break if you know the principle

Keys and Kerckhoffs's principle

- The only secret part of the cipher is a *key*
- Security does not depend on anything else being secret
- Modern (esp. civilian, academic) crypto embraces openness quite strongly

Symmetric vs. public key

- Symmetric key (today's lecture): one key used by all participants
- Public key: one key kept secret, another published
 - Techniques invented in 1970s
 - Makes key distribution easier
 - Depends on fancier math

Goal: secure channel

- Leaks no content information
 - Not protected: size, timing
- Messages delivered intact and in order
 - Or not at all
- Even if an adversary can read, insert, and delete traffic

One-time pad

- Secret key is truly random data as long as message
- Encrypt by XOR (more generally addition mod alphabet size)
- Provides perfect, "information-theoretic" secrecy
- No way to get around key size requirement

Computational security

- More realistic: assume adversary has a limit on computing power
- Secure if breaking encryption is computationally infeasible
 - E.g., exponential-time brute-force search
- Ties cryptography to complexity theory

Key sizes and security levels

- Difficulty measured in powers of two, ignore small constant factors
- Power of attack measured by number of steps, aim for better than brute force
- 2^{32} definitely too easy, probably 2^{64} too
- Modern symmetric key size: at least 2^{128}

Crypto primitives

- Base complicated systems on a minimal number of simple operations
- Designed to be fast, secure in wide variety of uses
- Study those primitives very intensely

Attacks on encryption

- Known ciphertext
 - Weakest attack
- Known plaintext (and corresponding ciphertext)
- Chosen plaintext
- Chosen ciphertext (and plaintext)
 - Strongest version: adaptive

Certificational attacks

- Good primitive claims no attack more effective than brute force
- Any break is news, even if it's not yet practical
 - Canary in the coal mine
- E.g., $2^{126.1}$ attack against AES-128
- Also watched: attacks against simplified variants

Fundamental ignorance

- We don't really know that any computational cryptosystem is secure
- Security proof would be tantamount to proving $P \neq NP$
- Crypto is fundamentally more uncertain than other parts of security

Relative proofs

- Prove security under an unproved assumption
- In symmetric crypto, prove a construction is secure if the primitive is
 - Often the proof looks like: if the construction is insecure, so is the primitive
- Can also prove immunity against a particular kind of attack

Random oracle paradigm

- Assume ideal model of primitives: functions selected uniformly from a large space
 - Anderson: elves in boxes
- Not theoretically sound; assumption cannot be satisfied
- But seems to be safe in practice

Pseudorandomness and distinguishers

- Claim: primitive cannot be distinguished from a truly random counterpart
 - In polynomial time with non-negligible probability
- We can build a distinguisher algorithm to exploit any weakness
- Slightly too strong for most practical primitives, but a good goal

Open standards

- How can we get good primitives?
- Open-world best practice: run competition, invite experts to propose then attack
- Run by neutral experts, e.g. US NIST
- Recent good examples: AES, SHA-3

A certain three-letter agency

- National Security Agency (NSA): has primary responsibility for "signals intelligence"
- Dual-mission tension:
 - Break the encryption of everyone in the world
 - Help US encryption not be broken by foreign powers