

CSEP 564: Computer Security and Privacy

# Authentication [finish]

## (Web) Privacy and Anonymity

Fall 2022

David Kohlbrenner


dkohlbre@cs

Thanks to Franzi Roesner, Dan Boneh, Dieter Gollmann, Dan Halperin, David Kohlbrenner, Yoshi Kohno, Ada Lerner, John Manferdelli, John Mitchell, Vitaly Shmatikov, Bennet Yee, and many others for sample slides and materials ...

# Logistics

- Lab 2 active
  - Apologies about the downtime last night.
- Lab 1 grades are up, sploit8 EC being manually graded
- The last paper (Dec 7<sup>th</sup>) has a different writeup format
  - See Canvas for the rubric, we're asking for you to reflect on the content

# Improving(?) Passwords

- Add biometrics
  - For example, keystroke dynamics or voiceprint
- Graphical passwords
  - Goal: easier to remember? no need to write down?
- Password managers 
  - Examples: LastPass, KeePass, 1password, built into browsers/OS
- Two-factor authentication
  - Leverage phone (or other device) for authentication

# Multi-Factor Authentication

1. Sign in with your Google Account

Email: hikingfan@gmail.com  
ex: pat@example.com

Password: .....

☒ Stay signed in

[Can't access your account?](#)

2. Google accounts

**Enter verification code**

To verify your identity on this computer, enter the verification code generated by your mobile application.

Enter code: 466453

☒ Remember verification for this computer for 30 days.

[Other ways to get a verification code »](#)

Google Authenticator

966286  
wileyc@acme.com

001322

Turn on Login Approvals

What is Login Approvals?

Login Approvals is a security feature that requires you to enter a code that we text to your phone when you log in from an unrecognized computer. You can enable this feature in a few simple steps.

If you ever lose access to your phone, you can always return to a previously-recognized computer to regain access to your account.

Note: You'll need to have your mobile phone with you to complete this process.

# FIDO + Hardware Two Factors



Questions:

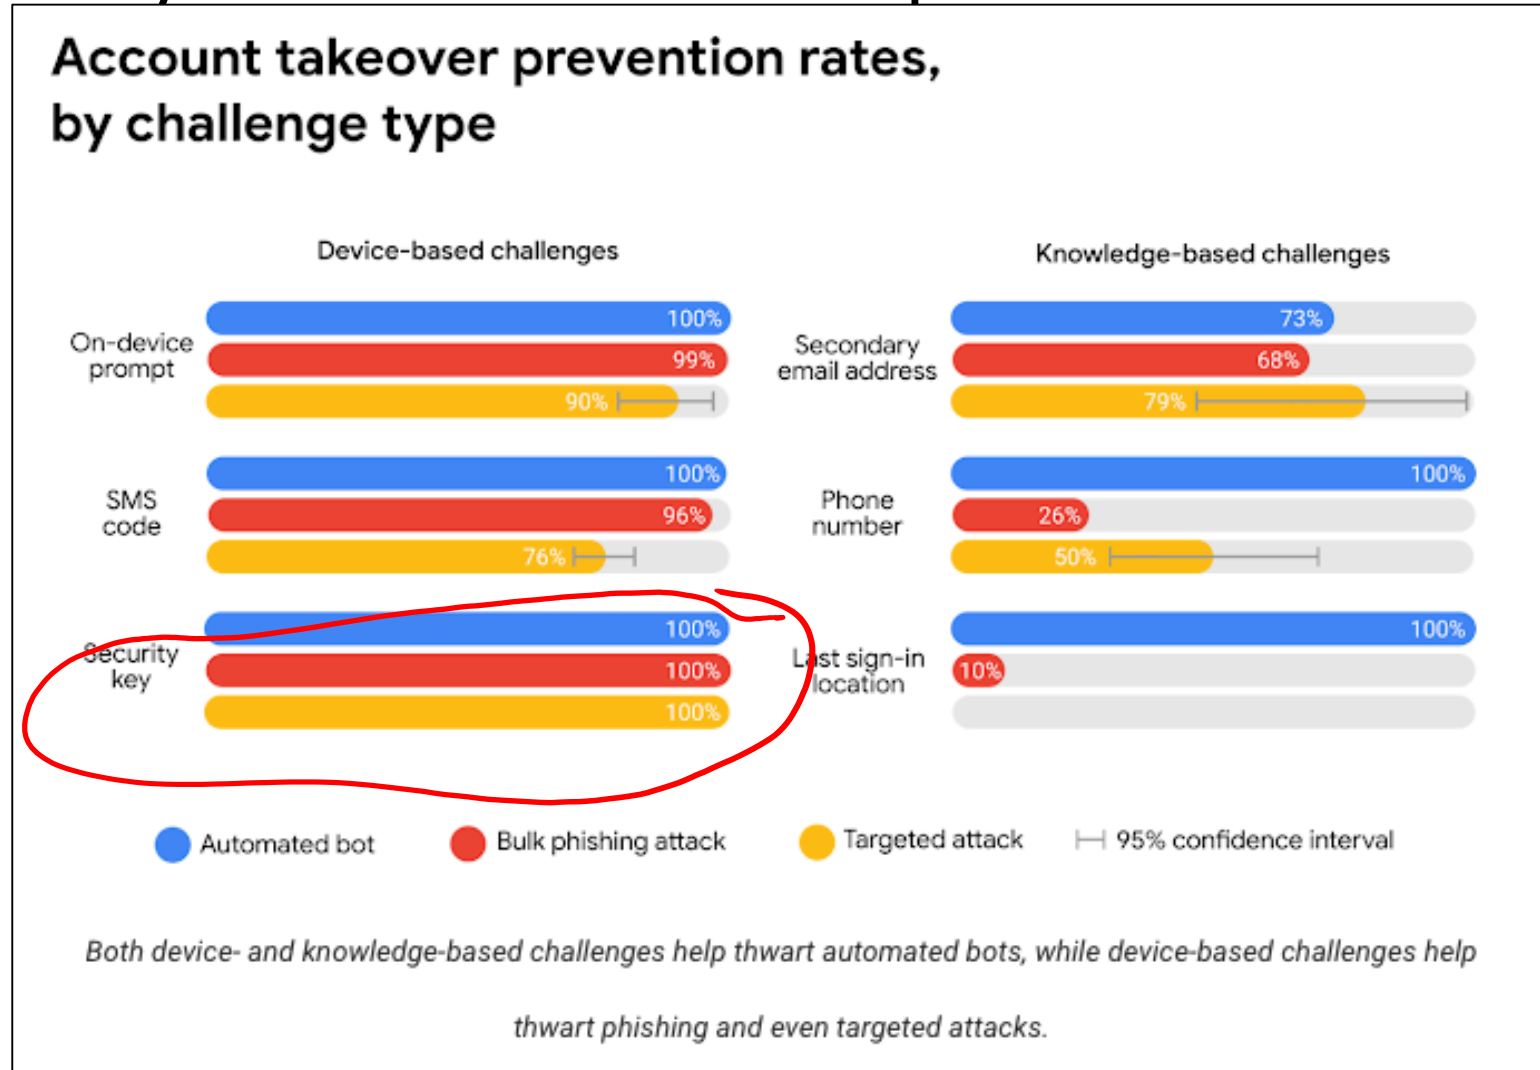
Do you use 2-factor auth?

Do you use a password manager?

Why or why not?


How to compromise account protected  
with hardware second factor?

# Secondary Factors Do Help!



# Graphical Passwords

- Many variants... one example: Passfaces
  - Assumption: easy to recall faces





# Graphical Passwords

- Another variant: draw on the image (Windows 8)




- Problem: users choose predictable points/lines

# Unlock Patterns



- Problems:
  - Predictable patterns (familiar pattern by now)
  - Smear patterns
  - Side channels: apps can use accelerometer and gyroscope to extract pattern!

# What About Biometrics?

- Authentication: **What you are**
- Unique identifying characteristics to authenticate user or create credentials
  - Biological and physiological: Fingerprints, iris scan
  - Behaviors characteristics - how perform actions: Handwriting, typing, gait
- Advantages:
  - Nothing to remember
  - Passive
  - Can't share (generally)
  - With perfect accuracy, could be fairly unique

# Issues with Biometrics

- Private, but not secret
  - Maybe encoded on the back of an ID card?
  - Maybe encoded on your glass, door handle, ...
  - Sharing between multiple systems?
- Revocation is difficult (impossible?)
  - Sorry, your iris has been compromised, please create a new one...
- Physically identifying
  - Soda machine to cross-reference fingerprint with DMV?
- Birthday paradox
  - With false accept rate of 1 in a million, probability of false match is above 50% with only 1609 samples

# Shifting Threat Models...

**BBC NEWS**

[OPEN](#) **The News in 2 minutes**

**News services**  
Your news when  
want it

**News Front Page**



**Africa**

**Americas**

**Asia-Pacific**

**Europe**

**Middle East**

**South Asia**

**UK**

**Business**

**Health**

**Science/Nature**

**Technology**

**Entertainment**

Last Updated: Thursday, 31 March, 2005, 10:37 GMT 11:37 UK

[E-mail this to a friend](#) [Printable version](#)

## Malaysia car thieves steal finger

**By Jonathan Kent**  
BBC News, Kuala Lumpur

**Police in Malaysia are hunting for members of a violent gang who chopped off a car owner's finger to get round the vehicle's hi-tech security system.**

The car, a Mercedes S-class, was protected by a fingerprint recognition system.

Accountant K Kumaran's ordeal began when he was run down by four men in a small car as he was about to get into his Mercedes in a Kuala Lumpur suburb.

**SEE ALSO:**

[Malaysia to act against pirates](#)  
16 Mar 05 | As

**RELATED INTEREST**

[Malaysian police](#)

The BBC is not responsible for the content of internet sites


**TOP ASIA-PACIFIC STORIES**

[Australians warn of cuts](#)

[Taiwan campus](#)

# Attacking Biometrics

- An adversary might try to steal biometric info
  - Malicious fingerprint reader
    - Consider when biometric is used to derive a cryptographic key
  - Residual fingerprint on a glass






# Attacking Biometrics



# Attacking Biometrics





# Attacking Biometrics



# Attacking Biometrics



# Privacy and web tracking

# A topic in flux

- Tracking via cookies
- Tracking via other methods
- Fingerprinting

# Ads That Follow You

Advertisers (and others) track your browsing behaviors for the purposes of **targeted ads**, **website analytics**, and **personalized content**.


The collage includes the following elements:

- A Zappos browser window showing a product page for "Converse Chuck Taylor All Star Core Ox - Black" with a promotional banner: "Order before 1pm PST for FREE Next Business Day shipping on all Clo".
- A CNN browser window showing the "Breaking News" section.
- A The Onion browser window showing a news article with a temperature of 92° and a headline about a Christian Academy.
- A social media bar for "the ONION" with YouTube and Facebook links.
- A video player with a "Click to play" button.
- A targeted advertisement for "Chuck Taylor All Star Core Ox Classic Shoes - White" and "Solarsoft Mule Men's Shoes - Black" with a price of \$65 and a "SHOP NOW" button. The ad includes a link to "Learn more" and a question "Why am I seeing this ad?".

# Third-Party Web Tracking

**Browsing profile for user 123:**

- cnn.com
- theonion.com
- adult-site.com
- political-site.com



These ads allow **criteo.com** to link your visits between sites, **even if you never click on the ads.**



# Marketing Technology Landscape

## The Martech 5000

Total Solutions 8,000

Advertising & Promotion 922

Content & Experience 1,936

Social & Relationships 1,969

Commerce & Sales 1,314

Data 1,258

Management 601

Access all the data of this landscape & more at [martech5000.com](https://martech5000.com)

2019  
7,040 solutions


2018  
6,829 solutions

2017  
5,381 solutions

2016  
3,874 solutions

2015  
1,876 solutions

2014  
947 solutions





### Advertising & Promotion

#### Display & Programmatic Advertising



#### Mobile Marketing



#### Native/Content Advertising



#### PR



#### Print



#### Search & Social Advertising



#### Video Advertising



### Content & Experience

#### Content Marketing



#### CMS & Web Experience Management



#### DAM & MRM & PIM



#### Email Marketing



#### Interactive Content



#### Marketing Automation & Campaign/Lead Management



#### Mobile Apps



#### Optimization Personalization & Testing



#### SEO



#### Video Marketing



### Social & Relationships

#### ABM



#### Call Analytics & Management



#### Customer Experience Service & Success



#### Influencers



#### CRM



#### Advocacy Loyalty & Referrals



#### Community & Reviews



#### Events, Meetings & Webinars



#### Social Media Marketing & Monitoring



#### Live Chat & Chatbots



### Commerce & Sales

#### Retail, Proximity & IOT



#### Affiliate Marketing & Management




#### Sales Automation Enablement & Intelligence



### Data

#### Marketing Analytics Performance & Attribution



#### Audience/Marketing Data & Data Enhancement



#### Ipaas Cloud/Data Integration & Tag Management



#### Business/Customer Intelligence & Data Science



#### DMP



#### Dashboards & Data Visualization



#### Governance Compliance And Privacy



#### Mobile & Web Analytics



#### Customer Data Platform



### Management

#### Agile & Lean Management



#### Collaboration



#### Talent Management



#### Budgeting & Finance



#### Projects & Workflow



#### Product Management




#### Vendor Analysis & Management





# 6,521% growth 2011 to 2022



<https://chiefmartec.com/2022/05/marketing-technology-landscape-2022-search-9932-solutions-on-martechmap-com/>

# Concerns About Privacy

**THE WALL STREET JOURNAL.**  
WHAT THEY KNOW | JULY 30, 2010  
The Wall Street Journal  
A Journal investigating business

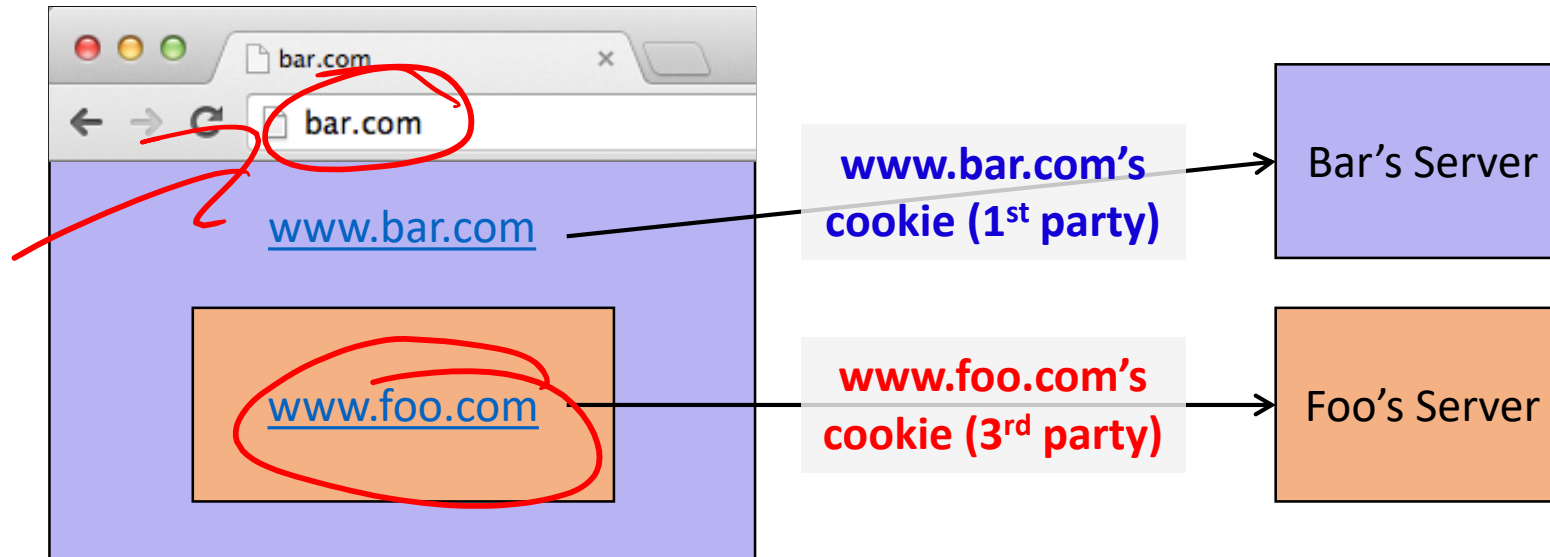
**The New York Times**  
May 6, 2011, 5:01 pm | 3 Comments  
**'Do Not Track' Privacy Bill Appears in Congress**  
By TANZINA VEGA  
And the privacy legislation just keeps on coming.  
On Friday, two bills were introduced in Washington in support of a Do Not Track mechanism that would give users control over how much of their data was collected by advertisers and other online companies.

By JENNIFER VALENTINO-DEVRIES,  
JEREMY SINGER-VINE and ASHKAN SOLTANI  
December 24, 2012

als  
ion

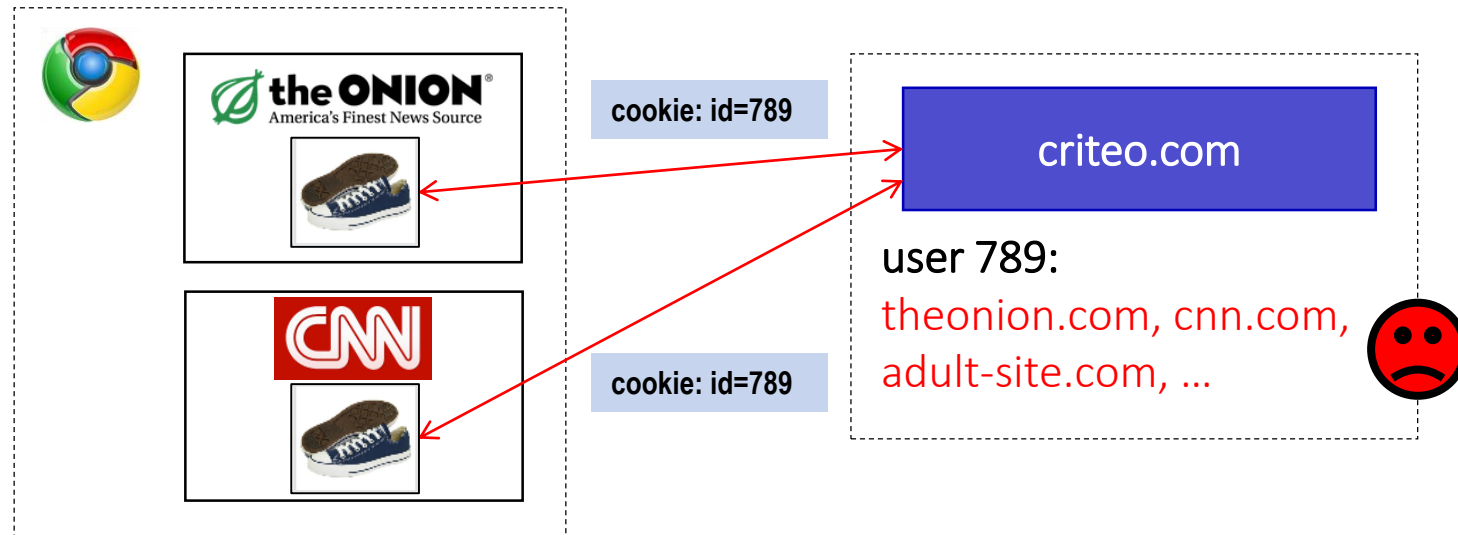
# First and Third Parties

- **First-party cookie:** belongs to top-level domain.
- **Third-party cookie:** belongs to domain of embedded content (such as image, iframe).



# Anonymous Tracking

Trackers **included in other sites** use **third-party cookies** containing unique **identifiers** to create browsing profiles.



# Basic Tracking Mechanisms

- Tracking requires:
  - (1) re-identifying a user.
  - (2) communicating id + visited site back to tracker.

## ▼ Hypertext Transfer Protocol

▶ GET /pixel/p-3aud4J6uA4Z6Y.gif?labels=InvisibleBox&bust=2710 HTTP/1.1\r\n

Host: pixel.quantserve.com\r\n

Connection: keep-alive\r\n

Accept: image/webp,\*/\*;q=0.8\r\n

User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_2) AppleWebKit/537.36

Referer: http://www.theonion.com/\r\n

Accept-Encoding: gzip, deflate, sdch\r\n

Accept-Language: en-US,en;q=0.8\r\n

Cookie: mc=52a65386-f1de1-00ade-0b26e; d=ENkBRgGHD4GYEA35MMIL74MKiyDs1A2MQI1Q

# Tracking Technologies

- HTTP Cookies

- HTTP Auth

- HTTP Etags

- Content cache

- IE userData

- HTML5 protocol and content handlers

- HTML5 storage

- Flash cookies

- Silverlight storage

- TLS session ID & resume

- Browsing history

- window.name

- HTTP STS

- DNS cache

- “Zombie” cookies that respawn

(<http://samy.pl/evercookie>)

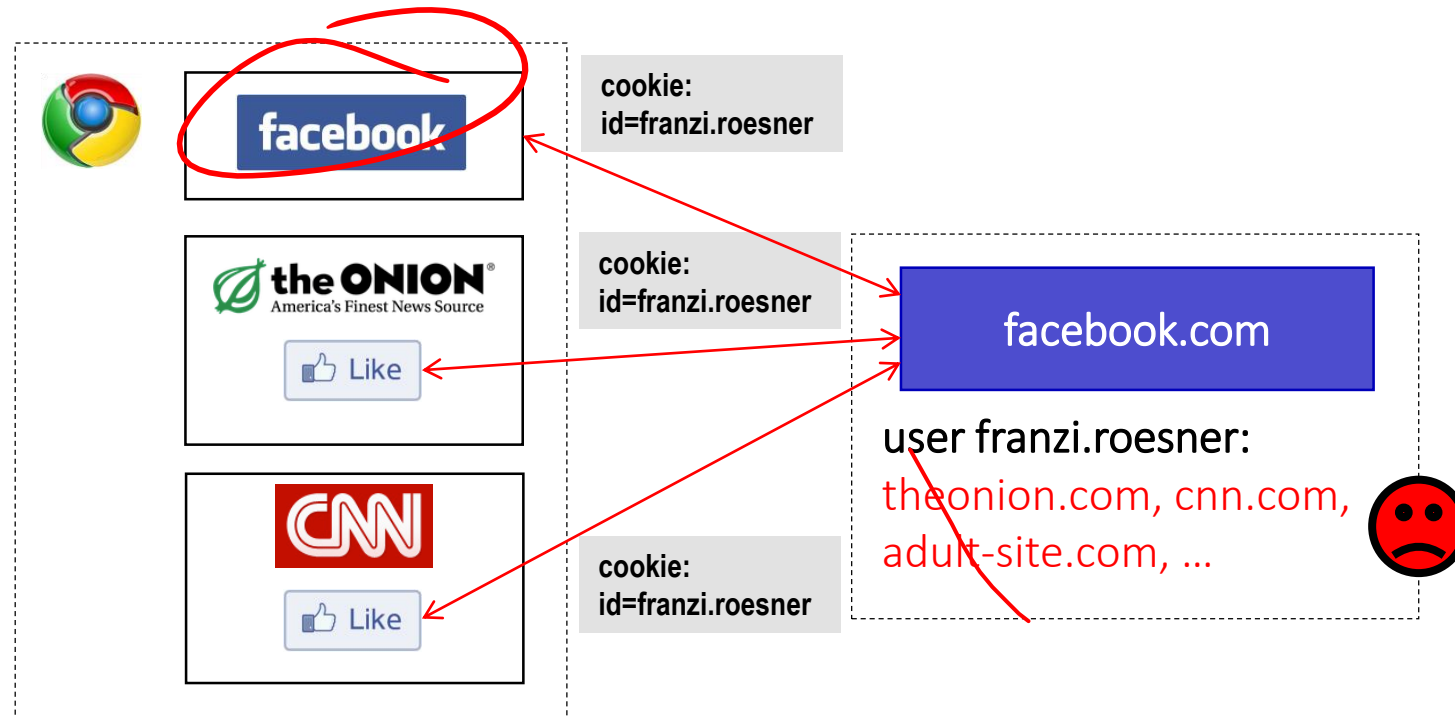
# Other Trackers?



## “Personal” Trackers



# Personal Tracking




- Tracking is not anonymous (linked to accounts).
- Users directly visit tracker's site → evades some defenses.

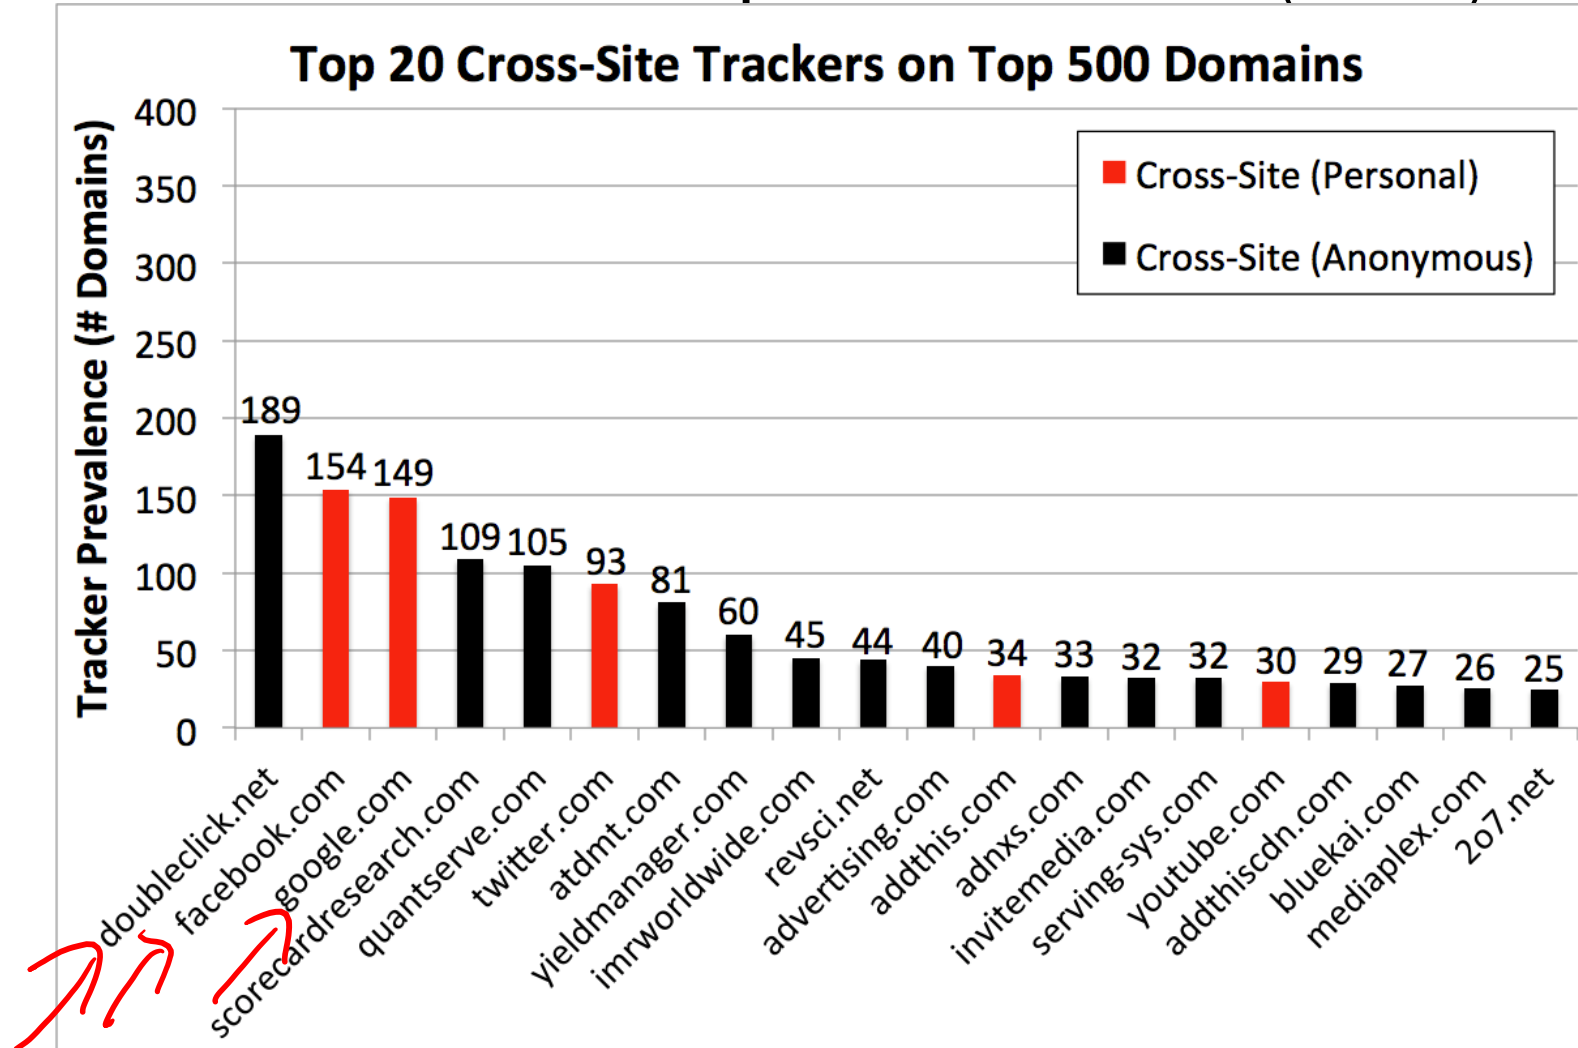


# How prevalent is tracking? (2011)

524 unique trackers on Alexa top 500 websites (homepages



# Who/what are the top trackers? (2011)




# How has this changed over time?

- The web has existed for a while now...
  - What about tracking before 2011?
  - What about tracking before 2009?
- Solution: time travel!



# The Wayback Machine to the Rescue




Time travel for web tracking: <http://trackingexcavator.cs.washington.edu>



# 1996-2016: More & More Tracking

- More trackers of more types, more per site, **more coverage**



# Defenses to Reduce Tracking

- Do Not Track?



Send a 'Do Not Track' request with your browsing traffic

Do Not Track is not a technical defense:  
trackers must honor the request.

# Defenses to Reduce Tracking

- Do Not Track proposal?
- Private browsing mode?

Private browsing mode protects against local, not network, attackers.

## You've gone incognito

Now you can browse privately, and other people who use this device won't see your activity. However, downloads and bookmarks will be saved. [Learn more](#)

Chrome won't save the following information:

- Your browsing history
- Cookies and site data
- Information entered in forms

Your activity might still be visible to:

- Websites you visit
- Your employer or school
- Your internet service provider

# Defenses to Reduce Tracking

- Do Not Track proposal?
- Private browsing mode?
- Third-party cookie blocking?





# 3<sup>rd</sup> party cookies

- Safari and FF (mostly) now block 3<sup>rd</sup> party cookies
  - <https://webkit.org/blog/10218/full-third-party-cookie-blocking-and-more/>
  - <https://blog.mozilla.org/blog/2019/09/03/todays-firefox-blocks-third-party-tracking-cookies-and-cryptomining-by-default/>

- Chrome...

“By undermining the business model of many ad-supported websites, blunt approaches to cookies encourage the use of opaque techniques such as fingerprinting (an invasive workaround to replace cookies), which can actually reduce user privacy and control. We believe that we as a community can, and must, do better.”

Aug 2022: Remove 3<sup>rd</sup> party cookies by 2024

# How should Google respond?

7:48

- [Pollev.com/dkohlbre](https://pollev.com/dkohlbre)
- Pretend someone fired all the ad/chrome execs and hired your group instead
- Safari and Firefox have removed ad's ability to track users via 3rd party cookies, and Google has committed to the same in Chrome by 2024.
- How should google respond to 3rd party cookies being removed?
- Think about the technical solutions, policy solutions, and even business model solutions available to you!

# Cookie Ghostwriting

- flickr.com
  - `<script src=siftscience.com/s.js />`
- S.js runs
  - Fp = fingerprintjs2
  - Setcookie(fp)
  - Hexagon-analytics.com/cookierecoiever?cookie=fp
- Every time you load flickr.com what happens?

Journey to the Center of the Cookie Ecosystem: Unraveling Actors' Roles and Relationship

# Cookie Ghostwriting

- ~~flickr.com~~ patreon.com
  - `<script src=siftscience.com/s.js />`
- S.js runs
  - Fp = fingerprintjs2
  - Setcookie(fp)
  - Hexagon-analytics.com/cookierecoiever?cookie=fp
- Every time you load ~~flickr.com~~ patreon.com what happens?

Journey to the Center of the Cookie Ecosystem: Unraveling Actors' Roles and Relationship

# Fingerprinting is out there

- Better than a 'voluntary' cookie: involuntary, unchangeable id!
  - "Fingerprint"
- Idea: Measure 'behavior' of browser
  - Smash into unique ID

# Fingerprinting Web Browsers

- User agent
- HTTP ACCEPT headers
- Browser plug-ins
- MIME support
- Clock skew
- Installed fonts
- Cookies enabled?
- Browser add-ons
- Screen resolution
- HTML5 canvas  
(differences in graphics SW/HW!)

# HTML5 Canvas Fingerprinting - Text




Figure 7: Difference maps for a group on `text_arial`

Mowery and Shacham, 2012

# HTML5 Canvas Fingerprinting - Image

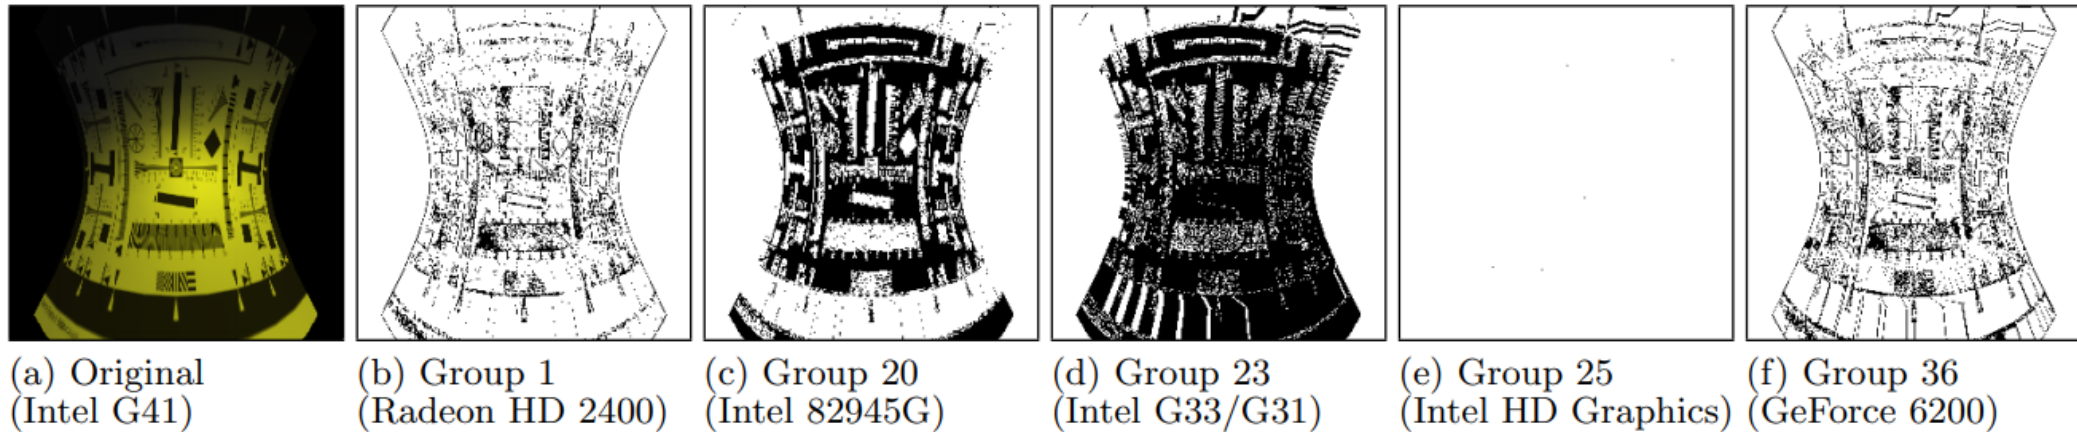


Figure 10: Original render and difference maps for Group 24

Mowery and Shacham, 2012



# And its out there!



Figure 4: Different images printed to canvas by fingerprinting scripts. Note that the phrase “*Cwm fjordbank glyphs vext quiz*” in the top image is a *perfect pangram*, that is, it contains all the letters of the English alphabet only once to maximize diversity of the outcomes with the shortest possible string.

# COVER YOUR TRACKS

See how trackers view your browser

[Learn](#)

[About](#)

## HOW TO READ YOUR REPORT

You will see a summary of your overall tracking protection. The first section gives you a general idea of what your browser configuration is blocking (or not blocking). Below that is a list of specific browser characteristics in the format that a tracker would view them. We also provide descriptions of how they are incorporated into your fingerprint.

## HOW CAN TRACKERS TRACK YOU?

Trackers use a variety of methods to identify and track users. Most often, this includes tracking cookies, but it can also include browser fingerprinting. Fingerprinting is a sneakier way to track users and makes it harder for users to regain control of their browsers. This report measures how easily trackers might be able to fingerprint your browser.

## HOW CAN I USE MY RESULTS TO BE MORE ANONYMOUS?

Knowing how identifiable you are, or whether you are blocking trackers, can help you take steps to better protect your privacy. Browser add-ons or protection mechanisms built into the browser can help. Even so, the sneakiest trackers have ways around even the strongest security.

Here are your Cover Your Tracks results. They include an overview of how visible you are to trackers, with an index (and glossary) of all the metrics we measure below.

**Our tests indicate that you have strong protection against Web tracking.**

### IS YOUR BROWSER:

Blocking tracking ads?	<u>Yes</u>
Blocking invisible trackers?	<u>Yes</u>
Protecting you from <u>fingerprinting</u> ?	<u>Your browser has a nearly-unique fingerprint</u>

Still wondering how fingerprinting works?

[LEARN MORE](#)

Note: because tracking and p

Your R

One in 145,235 browsers have the same fingerprint

measure all forms of

Within our dataset of several hundred thousand visitors tested in the past 45 days, only **one in 145,235.5 browsers have the same fingerprint as yours.**

# Fingerprinting as a security measure

- Blocking bots (e.g. reCAPTCHA)



- Validating users over-time



# How should we view tracking and fingerprinting efforts?

# “Privacy preserving” personalized ads aka FLoC

- <https://github.com/WICG/turtledove>
  - The browser, not the advertiser, holds the information about what the advertiser thinks a person is interested in.
  - Advertisers can serve ads based on an interest, but cannot combine that interest with other information about the person — in particular, with who they are or what page they are visiting.
  - Web sites the person visits, and the ad networks those sites use, cannot learn about their visitors' ad interests.

# “Privacy preserving” personalized ads aka Topics

- <https://github.com/patcg-individual-drafts/topics>
  - The browser, not the advertiser, holds the information about what the advertiser thinks a person is interested in.
  - Advertisers can serve ads based on an interest, but cannot combine that interest with other information about the person — in particular, with who they are or what page they are visiting.
  - Web sites the person visits, and the ad networks those sites use, cannot learn about their visitors' ad interests.

<https://support.google.com/google-ads/answer/11899856?hl=en>

# Privacy is far more than web tracking

- We've only started talking about it, in only 1 context.

# Anonymity



# Paper Discussion Time!

8.28

“Tor: The Second-Generation Onion Router”

Roger Dingledine, Nick Mathewson, Paul Syverson

- Choose one/more and discuss with neighbors:
  - What adversary is TOR intended to protect against?
  - What is a ‘hidden service’?
  - Why would someone want a hidden service?
  - Who runs an exit node, and why?
  - What are some of the weak points (as described here) of TOR?
  - Who uses TOR?



*The New Yorker,*  
1993


*"On the Internet, nobody knows you're a dog."*

# Privacy on Public Networks

- Internet is designed as a public network
  - Machines on your LAN may see your traffic, network routers see all traffic that passes through them
- Routing information is public
  - IP packet headers identify source and destination
  - Even a passive observer can figure out who is talking to whom
- Encryption does not hide identities
  - Encryption hides payload, but not routing information
  - Even IP-level encryption (tunnel-mode IPSec/ESP) reveals IP addresses of IPSec gateways
- Modern web: Accounts, web tracking, etc. ...



# What is Anonymity?

- Anonymity is the state of being not identifiable within a **set of subjects**
  - You cannot be anonymous by yourself! 
  - Big difference between anonymity and confidentiality
  - Hide your activities among others' similar activities
- Unlinkability of action and identity
  - For example, sender and email they send are no more related after observing communication than before
- Unobservability (hard to achieve)
  - Observer cannot even tell whether a certain action took place or not

# Questions



**Q1:** Why might we **want** people to have anonymity on the Internet?

**Q2:** Why might we **not want** people to have anonymity on the Internet?

# Applications of Anonymity (I)

- Privacy
  - Hide online transactions, Web browsing, etc. from intrusive governments, marketers and archivists
- Untraceable electronic mail
  - Corporate whistle-blowers
  - Political dissidents
  - Socially sensitive communications (online AA meeting)
  - Confidential business negotiations
- Law enforcement and intelligence
  - Sting operations and honeypots
  - Secret communications on a public network

# Applications of Anonymity (II)

- Digital cash
  - Electronic currency with properties of paper money (online purchases unlinkable to buyer's identity)
- Anonymous electronic voting 
- Censorship-resistant publishing 

# Part 1: Anonymity in Datasets



# How to release an anonymous dataset?

## A Face Is Exposed for AOL Searcher No. 4417749


By MICHAEL BARBARO and TOM ZELLER Jr.; Saul Hansell contributed reporting for this article.  
Published: August 9, 2006

Buried in a list of 20 million Web search queries collected by AOL and recently released on the Internet is user No. 4417749. The number was assigned by the company to protect the searcher's anonymity, but it was not much of a shield.

No. 4417749 conducted hundreds of searches over a three-month period on topics ranging from "numb fingers" to "60 single men" to "dog that urinates on everything."

And search by search, click by click, the identity of AOL user No. 4417749 became easier to discern. There are queries for "landscapers in Lilburn, Ga," several people with the last name Arnold and "homes sold in shadow lake subdivision gwinnett county georgia."

It did not take much investigating to follow that data trail to Thelma Arnold, a 62-year-old widow who lives in Lilburn, Ga., frequently researches her friends' medical ailments and loves her three dogs. "Those are my searches," she said, after a reporter read part of the list to her.

 FACEBOOK

 TWITTER

 GOOGLE+

 EMAIL

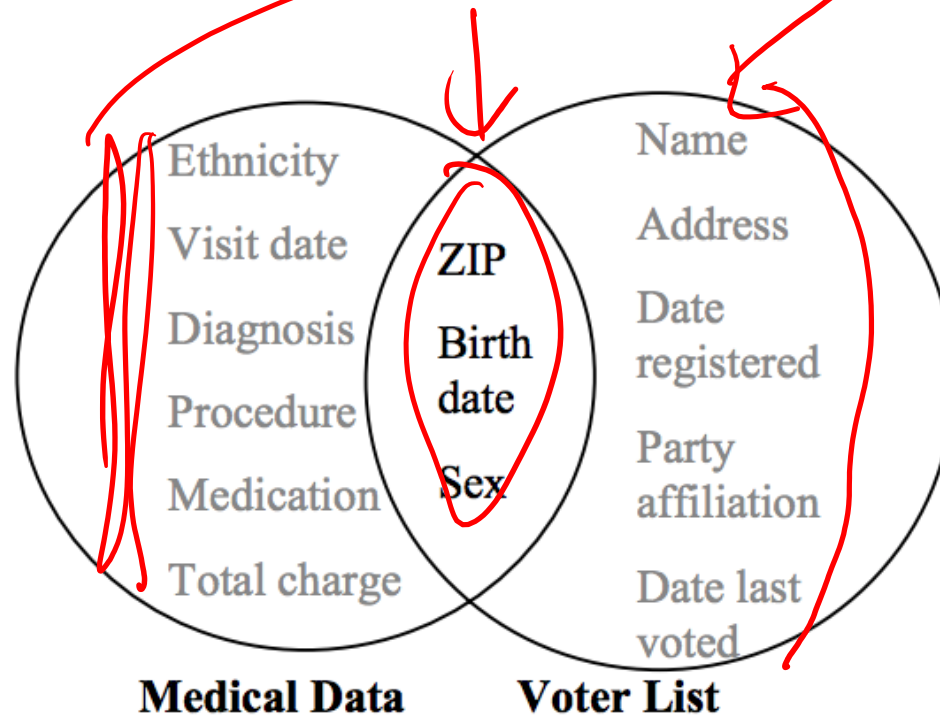
 SHARE

 PRINT

 REPRINTS

# How to release an anonymous dataset?

- Possible approach: **remove identifying information from datasets?**



Massachusetts  
medical+voter data  
[Sweeney 1997]

**Figure 1 Linking to re-identify data**

# k-Anonymity

- Each person contained in the dataset cannot be distinguished from at least  $k-1$  others in the data.

Name	Age	Gender	State of domicile	Religion	Disease
Ramsha	29	Female	Tamil Nadu	Hindu	Cancer
Yadu	24	Female	Kerala	Hindu	Viral infection
Salima	28	Female	Tamil Nadu	Muslim	TB
Kaker	27	Male	Karnataka	Parsi	No illness
Joan	24	Female	Kerala	Christian	Heart-related
Bahuksana	23	Male	Karnataka	Buddhist	TB
Rambha	19	Male	Kerala	Hindu	Cancer
Kishor	29	Male	Karnataka	Hindu	Heart-related
John	17	Male	Kerala	Christian	Heart-related
John	19	Male	Kerala	Christian	Viral infection

# k-Anonymity

- Each person contained in the dataset cannot be distinguished from at least  $k-1$  others in the data.

Name	Age	Gender	State of domicile	Religion	Disease
*	$20 < \text{Age} \leq 30$	Female	Tamil Nadu	*	Cancer
*	$20 < \text{Age} \leq 30$	Female	Kerala	*	Viral infection
*	$20 < \text{Age} \leq 30$	Female	Tamil Nadu	*	TB
*	$20 < \text{Age} \leq 30$	Male	Karnataka	*	No illness
*	$20 < \text{Age} \leq 30$	Female	Kerala	*	Heart-related
*	$20 < \text{Age} \leq 30$	Male	<b>Robust De-anonymization of Large Sparse Datasets</b>  Arvind Narayanan and Vitaly Shmatikov The University of Texas at Austin		
*	$\text{Age} \leq 20$	Male			
*	$20 < \text{Age} \leq 30$	Male			
*	$\text{Age} \leq 20$	Male			
*	$\text{Age} \leq 20$	Male			
*	$\text{Age} \leq 20$	Male	Kerala	*	Viral infection

Doesn't work for  
high-dimensional  
datasets (which  
tend to be **sparse**)

# Netflix Challenge:

- Netflix released a (non-uniform) random sample of user's movie ratings
- Challenge was to build a better recommendation system
- Data was 'anonymous'
  - ID # only
  - Random selection of a given user's ratings
  - "noise" added (appears that there was no noise)



# Result: No real anonymity

- Cross-correlate with IMBD ratings
- A handful (6 or fewer) ratings of non-top 500 movies is enough!

# Differential Privacy


- **Setting:** Trusted party has a database
- **Goal:** allow queries on the database that are useful but preserve the privacy of individual records
- **Differential privacy intuition:** add noise so that an output is produced with similar probability whether any single input is included or not
- Privacy of the computation, not of the dataset

$S \mid S + i$

## Part 2: Anonymity in Communication

# Chaum's Mix


- Early proposal for anonymous email
  - David Chaum. “Untraceable electronic mail, return addresses, and digital pseudonyms”. Communications of the ACM, February 1981.



Before spam, people thought anonymous email was a good idea 😊

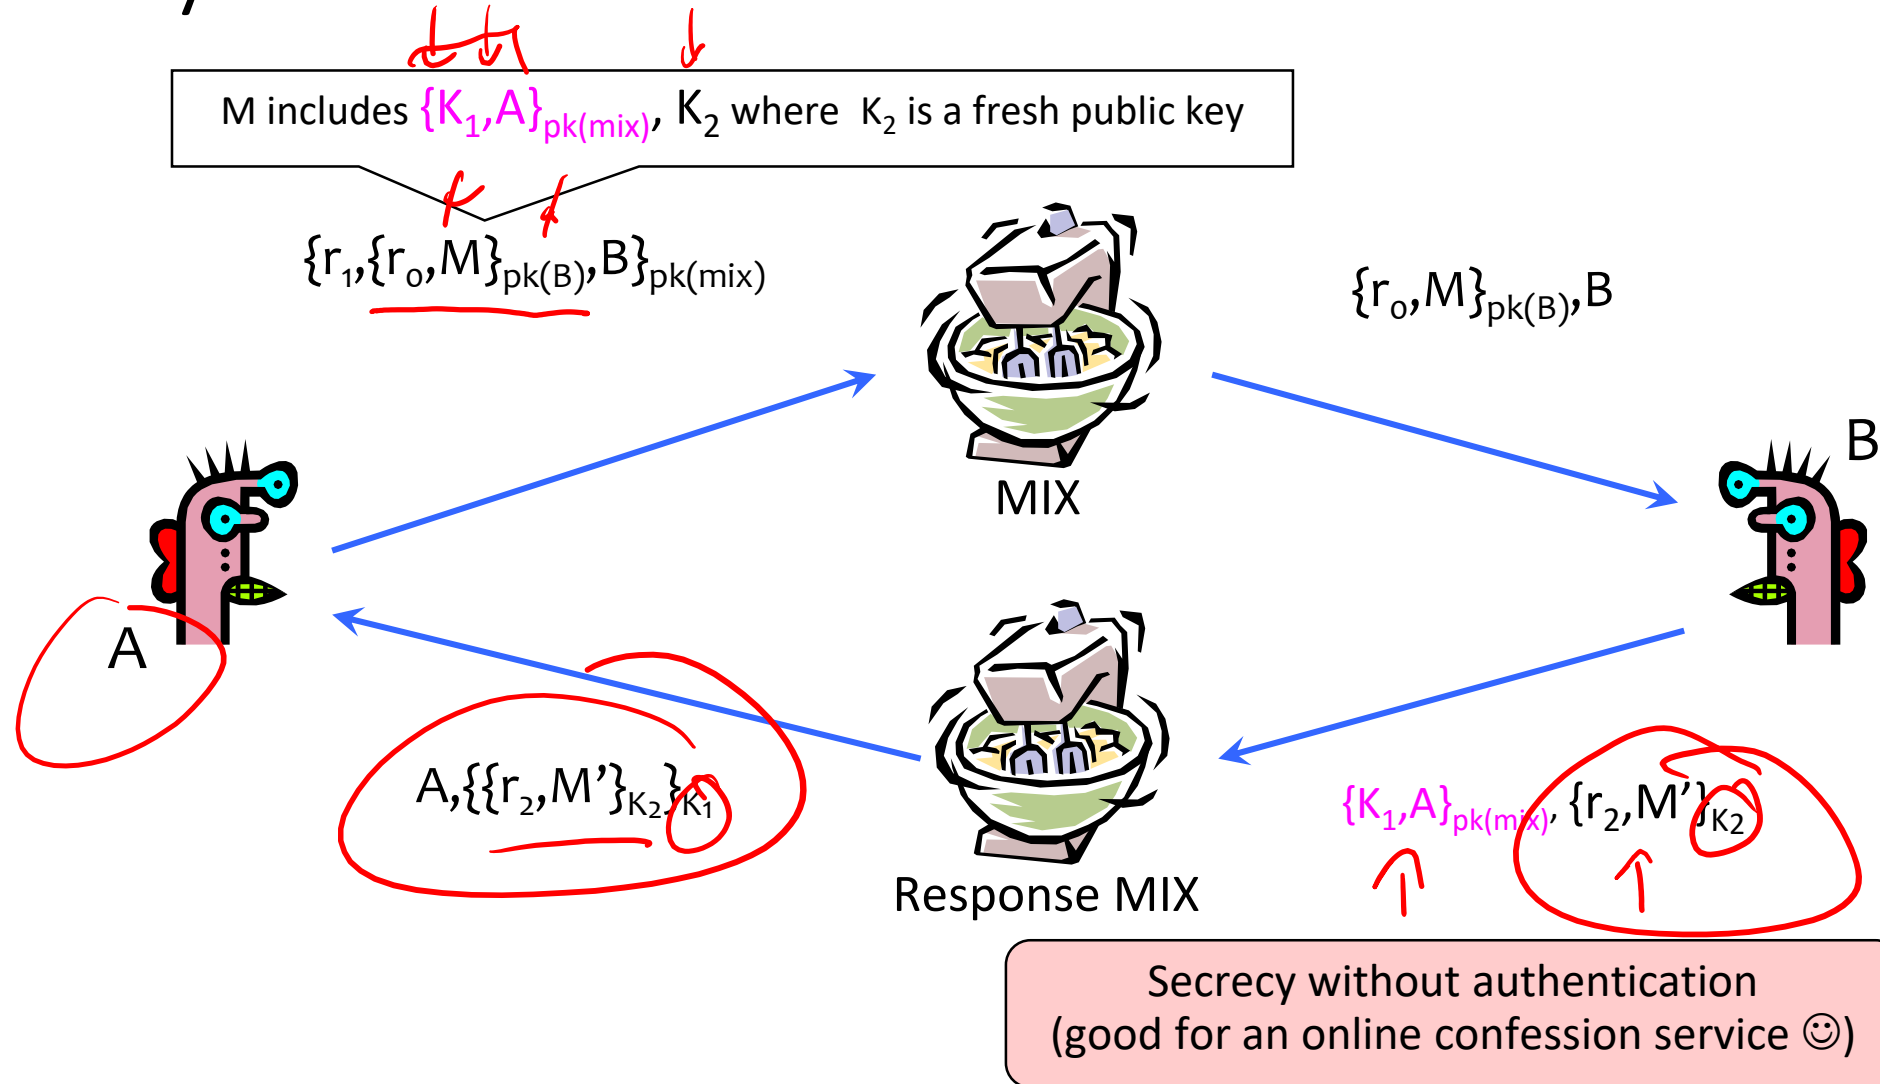
- Modern anonymity systems use Mix as the basic building block

# Basic Mix Design



Adversary knows all senders and all receivers, but cannot link a sent message with a received message

# Anonymous Return Addresses





# Mix Cascades and Mixnets




- Messages are sent through a **sequence of mixes**
  - Can also form an arbitrary network of mixes (“mixnet”)
- Some of the mixes may be controlled by attacker, but even a single good mix ensures anonymity
- Pad and buffer traffic to foil **correlation attacks**

# Disadvantages of Basic Mixnets

- Public-key encryption and decryption at each mix are **computationally expensive**
- Basic mixnets have **high latency**
  - OK for email, not OK for anonymous Web browsing
- Challenge: **low-latency anonymity network**


# Another Idea: Randomized Routing

e.g., Onion Routing



- Sender chooses a random sequence of routers
  - Some routers are honest, some controlled by attacker
  - Sender controls the length of the path

# Onion Routing




- Routing info for each link encrypted with router's public key
- Each router learns only the identity of the next router

# Tor

- Second-generation onion routing network
  - <http://tor.eff.org>
  - Developed by Roger Dingledine, Nick Mathewson and Paul Syverson
  - Specifically designed for **low-latency** anonymous Internet communications
- Running since October 2003
- “Easy-to-use” client proxy
  - Freely available, can use it for anonymous browsing


# Tor Circuit Setup (1)

- Client proxy establishes a symmetric session key and circuit with Onion Router #1



# Tor Circuit Setup (2)


- Client proxy extends the circuit by establishing a symmetric session key with Onion Router #2
  - Tunnel through Onion Router #1






# Tor Circuit Setup (3)

- Client proxy extends the circuit by establishing a symmetric session key with Onion Router #3
  - Tunnel through Onion Routers #1 and #2



# Using a Tor Circuit

- Client applications connect and communicate over the established Tor circuit.




# How do you know who to talk to?

- Directory servers
  - Maintain lists of active onion routers, their locations, current public keys, etc.
  - Control how new routers join the network
    - “Sybil attack”: attacker creates a large number of routers
  - Directory servers’ keys ship with Tor code

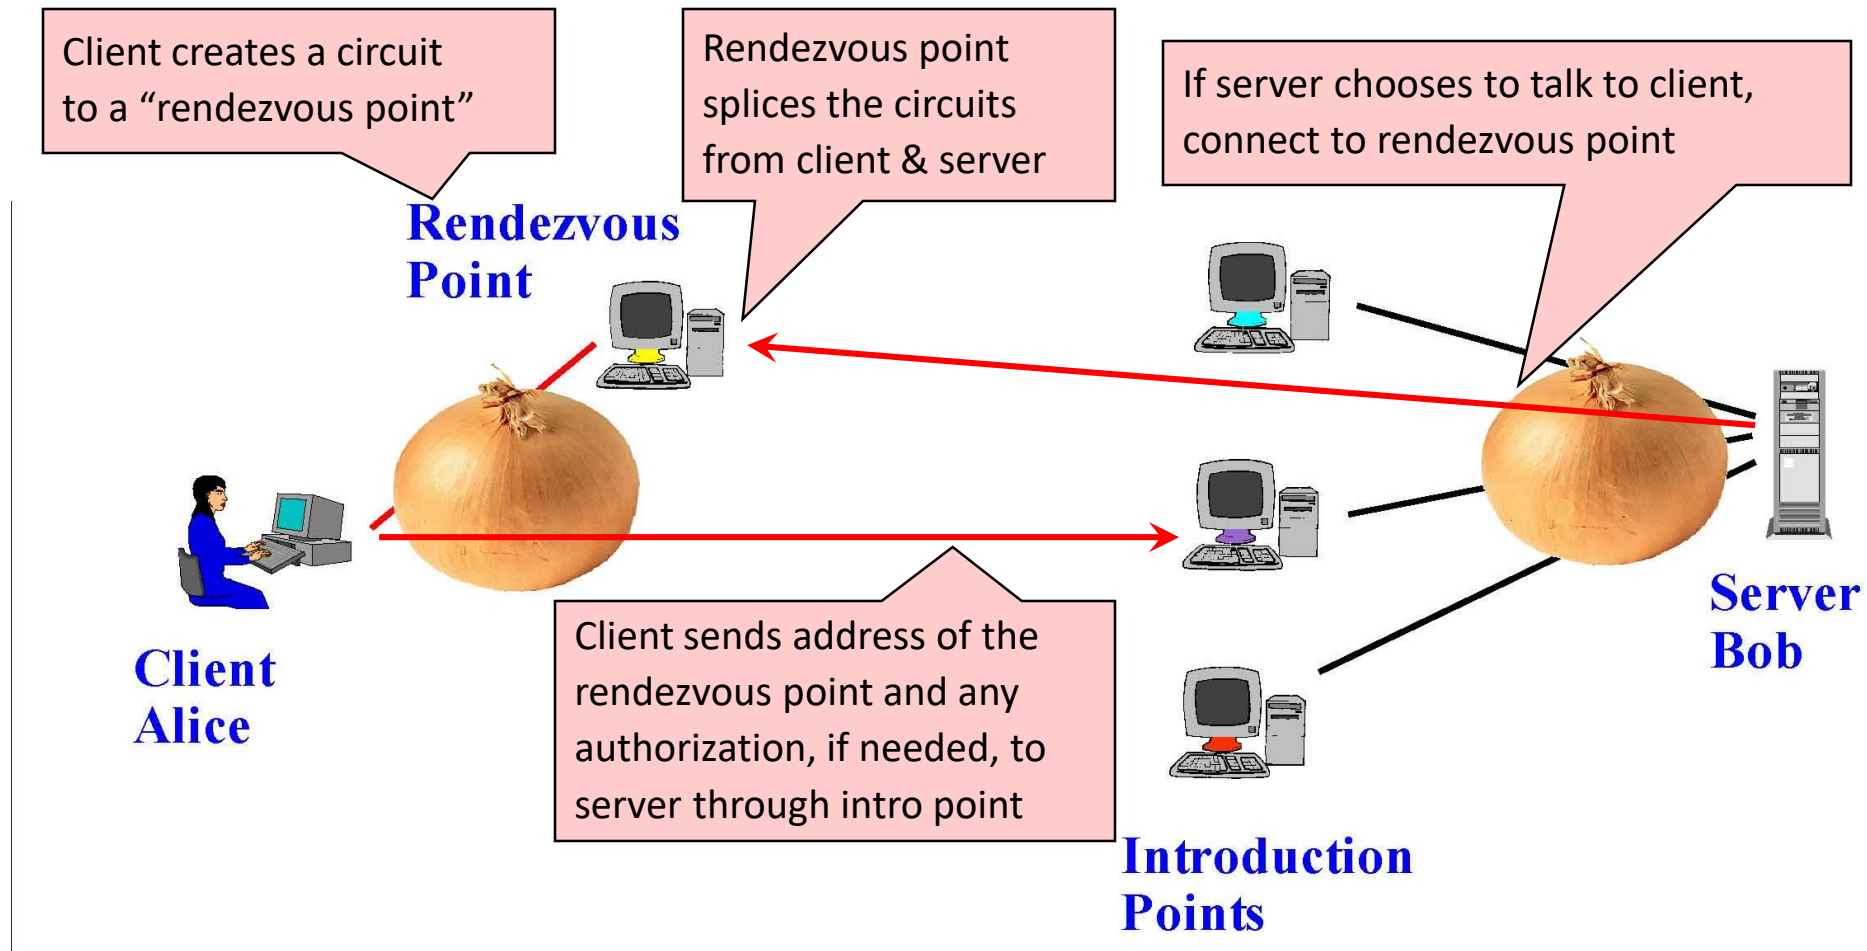
# Location Hidden Service

- **Goal:** deploy a server on the Internet that anyone can connect to **without knowing where it is or who runs it**
- Accessible from anywhere
- Resistant to censorship
- Can survive a full-blown DoS attack
- Resistant to physical attack
  - Can't find the physical server!

# Creating a Location Hidden Server



# Using a Location Hidden Server




# Issues and Notes of Caution

- Passive traffic analysis
  - Infer from network traffic who is talking to whom
  - To hide your traffic, must carry other people's traffic!
- Active traffic analysis
  - Inject packets or put a timing signature on packet flow
- Compromise of network nodes
  - Attacker may compromise some routers
    - Powerful adversaries may compromise “too many”
  - It is not obvious which nodes have been compromised
    - Attacker may be passively logging traffic
  - Better not to trust any individual router
    - Assume that some fraction of routers is good, don't know which



# Issues and Notes of Caution

- Tor isn't completely effective by itself
  - Tracking cookies, fingerprinting, etc.
  - Exit nodes can see everything!



# Issues and Notes of Caution

- The simple act of using Tor could make one a **target for additional surveillance**
- Hosting an exit node could result in **illegal activity coming from your machine**
- Tor not designed to protect against adversaries with the capabilities of a state (public statement by designers, at least in the past)