

CSE508 Network Security

3/2/2016 **Authentication**

Michalis Polychronakis
Stony Brook University

Authentication

The process of reliably verifying the identity of someone (or something)

What is identity?

Which characteristics uniquely identify an entity?


Authentication is a critical service, as many other security mechanisms are based on it

Entity authentication is the security service that enables communicating parties to verify the identity of their peers

Two main types

Human to computer

Computer to computer



Credentials

Evidence used to prove an identity

User Authentication: credentials supplied by the user

- Something you know

- Something you have

- Something you are

Computer authentication: crypto, location

- Computers (in contrast to humans) can “remember” large secrets (keys) and perform complex cryptographic operations

- Location: evidence that an entity is at a specific place (e.g., IP address/subnet)

Authentication can be delegated

- The verifying entity accepts that a trusted third party has already established authentication

Something You Know: Password-based Authentication

Passwords, passphrases, pins, key-phrases, access codes, ...

Say the magic word

Good passwords are easy to remember and hard to guess

Easy to remember → easy to crack

Hard to crack → hard to remember

Bad ideas: DOB, SSN, zip code, favorite team name, ...

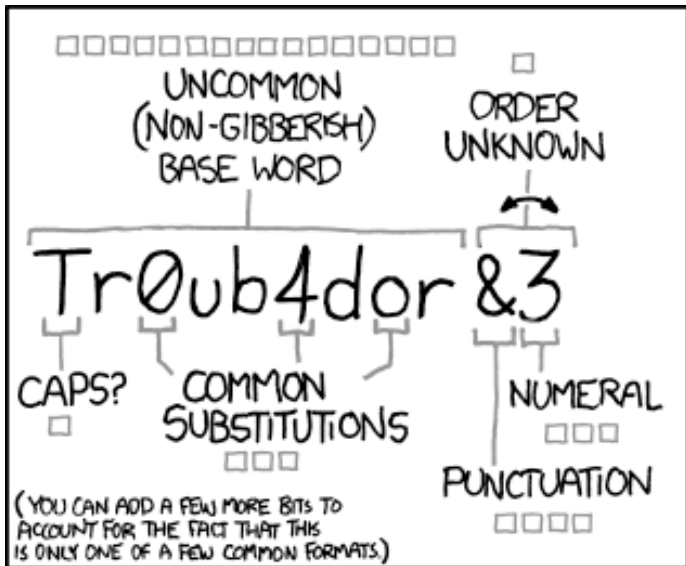
Password space (bits) depends on:

Password length

Character set

Better way to think about strong passwords

Long passphrases, combined with custom variations, symbols, numbers, capitalization, ...



~28 BITS OF ENTROPY

$2^{28} = 3 \text{ DAYS AT } 1000 \text{ GUESSES/SEC}$


(PLAUSIBLE ATTACK ON A WEAK REMOTE WEB SERVICE. YES, CRACKING A STOLEN HASH IS FASTER, BUT IT'S NOT WHAT THE AVERAGE USER SHOULD WORRY ABOUT.)

DIFFICULTY TO GUESS: **EASY**

WAS IT TROMBONE? NO, TROUBADOR. AND ONE OF THE 0s WAS A ZERO?

AND THERE WAS SOME SYMBOL...

DIFFICULTY TO REMEMBER: **HARD**



~44 BITS OF ENTROPY

$2^{44} = 550 \text{ YEARS AT } 1000 \text{ GUESSES/SEC}$

DIFFICULTY TO GUESS: **HARD**

THAT'S A BATTERY STAPLE.

CORRECT!

DIFFICULTY TO REMEMBER: YOU'VE ALREADY MEMORIZED IT

THROUGH 20 YEARS OF EFFORT, WE'VE SUCCESSFULLY TRAINED EVERYONE TO USE PASSWORDS THAT ARE HARD FOR HUMANS TO REMEMBER, BUT EASY FOR COMPUTERS TO GUESS.

Attacking Passwords

Offline cracking

Online guessing

Eavesdropping

Capturing



Brute force attacks

Password Storage

Storing passwords as plaintext is disastrous

Better way: store a cryptographic hash of the password

Even better: store a “salted” version of the password

Prevent precomputation of hash values (wordlists of popular passwords, rainbow tables, ...)

Even if two users have the same password, their hash values will be different -> need to be cracked separately

Salting *does not* make guessing a given password harder!

Username	Salt	Password hash
Bobbie	4238	$h(4238, \text{\$uperman})$
Tony	2918	$h(2918, 63\%TaeFF)$
Mitsos	6902	$h(6902, \text{zour1da})$
Mark	1694	$h(1694, \text{Rockybrook\#1})$

Still, password databases are getting leaked...

Password Cracking

Exhaustive search → infeasible for large password spaces

Dictionary attacks

- Language words

- List of previously leaked real user passwords

Variations and common patterns

- Prepend/append symbols/numbers/dates, weird capitalization, l33tspeak, visually similar characters, intended misspellings, ...

Target-specific information

- DOB, family names, favorite team, pets, hobbies, anniversaries, language, slang, ...

- Many ease to acquire from social networking services

- Particularly effective against “security questions”


Advanced techniques

- Probabilistic context-free grammars, Markov models, ...

Combination of all the above

25 Most-used (Worse) Passwords

password	letmein	2000
123456	monkey	jordan
12345678	696969	superman
1234	abc123	harley
qwerty	mustang	1234567
12345	michael	
dragon	shadow	
pussy	master	
baseball	jennifer	
football	111111	



Distribution of 4-digit sequences within RockYou passwords

Wordlists

ce#ebc.dk	4637324	gea8mw4yz	fujinshan	masich	gothpunksk8er	20081010
goddess5	bugger825	kukumbike	counter	pengaiwei	rftaeo48	leelou44
20071002	marmaris	260888	N8mr0n	coalesce	8d7R0K	8UfjeGb0
271075711	jinjin111	jordi10	520057	56402768	5172032	200358808
zs3cu7za	170383gp	lexusis	adc123	thesis	aics07	dellede
scoopn	3484427	kj011a039	bmaster	aabbcc894	34mariah	liang123.
frygas1411	fl33321	c84bwlrb	qbjh04zg	marion&maxime	dongqinwei	captainettekt
SL123456sl	zwqrfg	priyanka05	ueldaa79	614850	samarica	kwiki-mart
12345687ee123	67070857	loveneverdies	EMANUELLI	ydz220105	cap1014	mdovydas
xuexi2010	432106969	u8Aqebj576	yanjing	584521584521	0167387943	tigmys2001
daigoro	6856	FGYfgy77	assynt	txudecp	AE86Trueno	denial
12345614	704870704870	659397	62157173	84410545	19700913	678ad5251
DICK4080	pv041886	327296	0704224950753	pietro.chiara	mcsuap	woaiuwai
567891234	20060814	74748585	6903293	jman1514	bu56mpbu	1591591591212
tilg80	512881535	19720919	axaaxa	heryarma	danbee	hNbDGN
6z08c861	milanimilani	050769585	hilal1	39joinmam	passw<>	cardcap
:zark:	472619	nicopa	30091983	timelapse	money521	13985039393
ravishsneha	dbyxw888	2232566	2510618981	mwinkar	conan83	001104
150571611369	85717221	bearss	soukuokpan	251422	nxfjpl	desare11
661189	cc841215	n0tpublic	tosecondlife	willrock	rateg143	412724198
passme	ariana19321	isitreal00	p4os8m6q	YHrtfgDK	kojyihen	nibh1kab
trolovinasveta	bbnnn	ashraf19760	015614117	xys96exq	058336257	asferg
abdulkhaleque	ang34hehiu	48144	acw71790	mercadotecnia	sarah4444	hqb555
007816	wj112358	22471015	lsyljm2	8s5sBEx7	7363437	xgames7
xLDSX	Brenda85	antyzhou115	2xgialdl	0125040344	freindship	muckerlee
Florida2011	786525pb	0167005246	gaybar9	margitka	JytmvW0848	choqui67
037037	shi461988	ec13kag	88203009	omaopa	sb inbau	12130911
WestC0untry	pingu	226226226226	MKltyh87	dfTi6nh	30907891	lierwei120
hitsugaiya	yeybozip	6767537/33	quiggle	1314520521	0515043111	skytdvn
955998126	71477nak	mimilebrock	2063775206	pixma760	1973@ati	milena1995
3n3rmax	stokurew	gueis8850	fr3iH3it	pearpear	wlxgjf	kambala11

Password Hashing Functions

Problem: hash functions are very fast to evaluate → enable fast guessing attacks

Solution: slow down the guessing process (password “stretching”)

E.g., 10-100ms per check → significant cost for the server if it must handle many users (!)

Make heavy use of available resources

Computation should be fast enough to validate honest users, but render password guessing infeasible

Adaptable: flexible cost (time/memory complexity) parameters

Bcrypt [Provos and Mazières, 1999]

Cost-parameterized, modified version of the Blowfish encryption algorithm

Tunable cost parameter (exponential number of loop iterations)

Alternatives: Scrypt (memory-hard), PBKDF2 (PKCS standard)

Online Guessing

Similar strategy to offline guessing, but rate-limited

Connect, try a few passwords, get disconnected, repeat...

Prerequisite: know a valid user name

Many failed attempts can lead to a system reaction

Introduce delay before accepting future attempts (exponential backoff)

Shut off completely (e.g., ATM capturing/disabling a card after 3 tries)

Ask user to solve a CAPTCHA

Very common against publicly accessible SSH, VPN, RDP, and other servers

Main reason people move sshd to a non-default port

Fail2Ban: block IP address after many failed attempts → may allow an attacker to lock you out of the server (!)

Better: disable password auth and use a key pair → cumbersome if having to log in from many/others' computers

LOGIN: mitch
PASSWORD: FooBar!-7
SUCCESSFUL LOGIN

(a)

LOGIN: carol
INVALID LOGIN NAME
LOGIN:

(b)

LOGIN: carol
PASSWORD: Idunno
INVALID LOGIN
LOGIN:

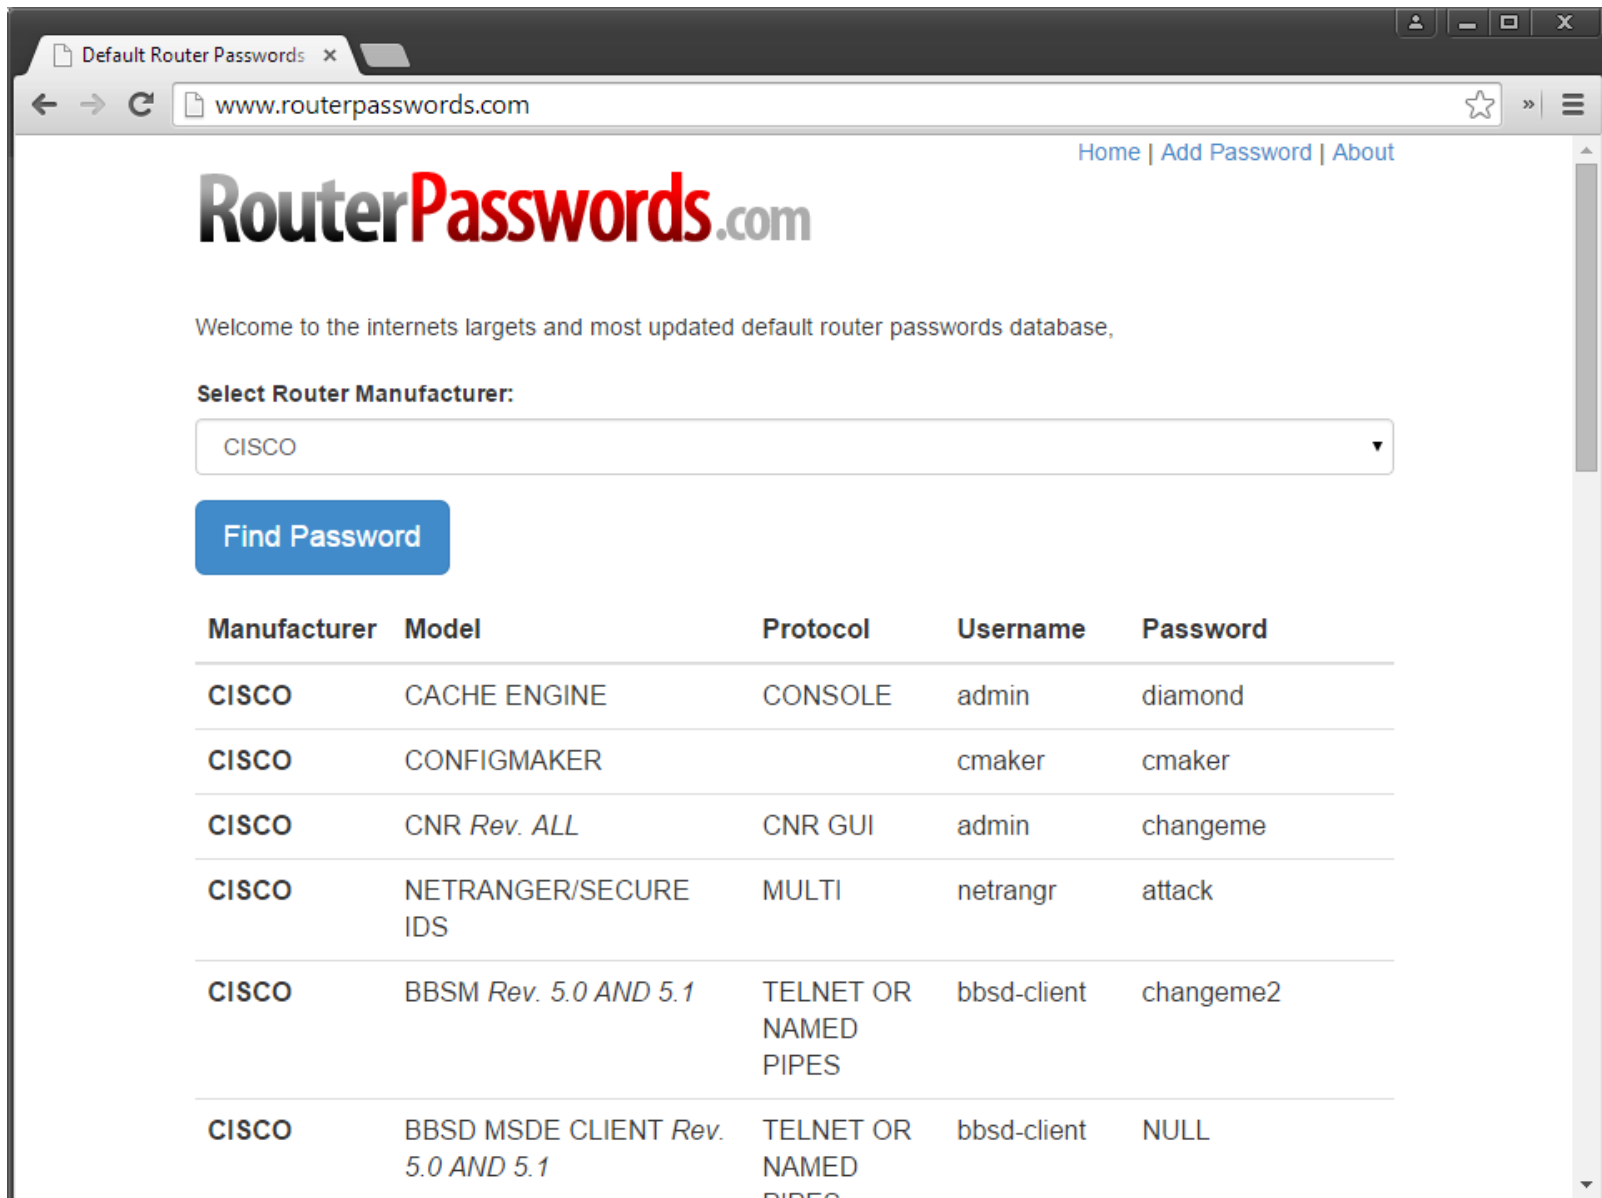
(c)

(a) A successful login

(b) Login rejected after name is entered

(c) Login rejected after name and password are typed

Try the Default First



The screenshot shows a web browser window with the address bar displaying "www.routerpasswords.com". The page features the "RouterPasswords.com" logo and a navigation menu with links for "Home", "Add Password", and "About". A welcome message reads: "Welcome to the internet's targets and most updated default router passwords database,". Below this is a section titled "Select Router Manufacturer:" with a dropdown menu currently set to "CISCO". A blue "Find Password" button is positioned below the dropdown. The main content area displays a table of default router passwords for Cisco devices.

Manufacturer	Model	Protocol	Username	Password
CISCO	CACHE ENGINE	CONSOLE	admin	diamond
CISCO	CONFIGMAKER		cmaker	cmaker
CISCO	CNR Rev. ALL	CNR GUI	admin	changeme
CISCO	NETRANGER/SECURE IDS	MULTI	netrangr	attack
CISCO	BBSM Rev. 5.0 AND 5.1	TELNET OR NAMED PIPES	bbsd-client	changeme2
CISCO	BBSD MSDE CLIENT Rev. 5.0 AND 5.1	TELNET OR NAMED PIPES	bbsd-client	NULL

Eavesdropping and Replay

Physical world

- Watch user type password (shoulder surfing)

- Cameras (ATMs skimmers)

- Lift fingerprints (iPhone)

- Post-it notes

Network makes things easier

- Sniffing (LAN, WiFi, ...)

- Man-in-the-Middle attacks

Defenses

- Encryption

- One-time password schemes

Kerberos

Long-lived vs. session keys

Use long-lived key for authentication and negotiating session keys

Use “fresh,” ephemeral session keys (prevent replay, cryptanalysis, old compromised keys) for encrypted communication, MACs, ...

Kerberos: most widely used (non-web) single sign-on system

Originally developed at MIT, now used in Unix, Windows, ...

Authenticate users to services: using their password as the initial key, without having to retype it for every interaction

A Key Distribution Center (KDC) acts as a trusted third party for key distribution

Online authentication: Variant of Needham-Schroeder protocol

Assumes a non-trusted network: prevents eavesdropping

Assumes that the Kerberos server and user workstations are secure...

Use cases: workstation login, remote share access, printers, ...

Password Capture

Hardware bugs/keyloggers

Software keyloggers/malware

Phishing

Social engineering

Welcome to Windows




Microsoft®
Windows
Professional

Copyright © 1985-2001
Microsoft Corporation




Press Ctrl-Alt-Delete to begin.

Requiring this key combination at startup helps keep computer secure. For more information, click [Help](#).



(a)



(b)

- (a) Correct login screen
- (b) Phony login screen

Something You Have: Authentication Tokens

One-time passcode tokens

- Time-based

- Counter-based

Other authentication tokens: store certificates, encryption keys, challenge-response, ...

Smartcards (contact or contactless)

- Identification, authentication, data storage, limited processing

- Magnetic stripe cards, EMV (chip-n-pin credit cards), SIM cards, RFID tags, ...

USB/bluetooth tokens, mobile phones, watches, ...

- Can be used as authentication devices

Multi-factor Authentication

Present several separate credentials of different types

Most common: 2-factor authentication (2FA)

Example: Password + hardware token, mobile phone, ...

Example: ATM card + pin

Motivation: a lost/guessed password is now not enough for attackers → not always true

Man-in-the-Middle: set up fake banking website, relay password to real website, let the user deal with the second factor...

Man-in-the-browser: hijack/manipulate an established session after authentication has completed (banking Trojans)

Dual infection: compromise both PC and mobile device

Implementation-dependent usability issues

In-flight WiFi, but cannot receive SMS...

Fallback: backup one-time-use passcodes (where to keep them?)

Biometrics

Fingerprint reader (iOS)

Face recognition (android)

Retina/Iris scanner

Voice recognition

...

Continuous authentication

Keystroke timing, usability patterns, ...

Crypto-based Authentication

Some way to use a cryptographic key to prove a user's identity

Basic idea: user performs a requested cryptographic operation on a value (a challenge) that the verifier supplies

- Usually based on knowledge of a key (secret key or private key)

- Can use symmetric (e.g., Kerberos) or public key schemes

How can we trust a key? Why is it authentic?

- Need to establish a level of trust

- Different approaches: TOFU, PKI, Web of trust

- Emerging approach: blockchain/ledger-based PKI

Trust on First Use

Use case: SSH

Performs *mutual authentication*

Server *always* authenticates the client

password, key pair, ...

Client almost always authenticates the server – *except the first time!*

First connection: server presents its public key

No other option for the user but to accept it: MitM opportunity

Subsequent connections: client remembers server's key, and triggers an alert on key mismatch

Pragmatic solution, but shifts the burden to users

Users must determine the validity of the presented key

Assuming a key change is valid without verifying the new key offers no protection against MitM (unfortunately, that's what most users do)

```
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@   WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!   @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!
Someone could be eavesdropping on you right now (man-in-the-middle attack)!
It is also possible that the RSA host key has just been changed.
The fingerprint for the RSA key sent by the remote host is
df:c8:52:aa:cd:e3:da:8c:ec:50:46:db:4d:21:d9:c7.
Please contact your system administrator.
Add correct host key in /root/.ssh/known_hosts to get rid of this message.
Offending key in /root/.ssh/known_hosts:1
RSA host key for 192.168.2.5 has changed and you have requested strict checking.
Host key verification failed.
```

Certificates

How can we distribute “trusted” public keys?

Public directory -> risk of forgery and tampering

More practical solution: “certified” public keys

A certificate is a digitally signed message containing an identity and a public key

Makes an association between a user/entity and a private key

Valid until a certain period

Why trust a certificate?

Because it is signed by an “authority”

Third party’s signature prevents tampering



Public Key Infrastructures (PKI)

Facilitate the authentication and distribution of public keys based on identities

Set of roles, policies, and procedures to create, manage, distribute, use, store, and revoke certificates

An issuer signs certificates for subjects

Trust anchor

Methods of certification

Certificate authorities (hierarchical structure – root of trust)

Web of trust (decentralized, peer-to-peer structure)

Certificate Authorities

Trusted third-parties responsible for certifying public keys

Most CAs are tree-structured

Single point of failure: CAs can be compromised!

Why should we trust an authority?

How do you know the public key of the Certificate Authority (CA)?

CA's public key (trust anchor) must be provided out of band

Trust has to start somewhere

Operating systems and browsers are pre-configured with ~200 trusted root certificates

A public key for any website in the world will be accepted without warning if certified by any of these CAs

Web of Trust

Entirely decentralized authentication

- No single point of failure

- No need to buy certs from CAs

- Used in PGP

Users sign other users' keys

- Only if they deem them trustworthy

- Certificate signings can form an arbitrarily complex graph

- Users can verify path to as many trust anchors as they wish

Drawbacks

- Hard to use, requires in-person verification – key signing parties!

- Hard to know what trust level to assign transitively

WoT Alternative: Online Social "Tracking"

The screenshot shows a web browser window with the address bar displaying `https://keybase.io/mikepo`. The page header features the Keybase logo, a search bar, and navigation links for 'Join', 'Login', and help icons. The main content area displays the profile for 'mikepo', including a circular profile picture of Michalis Polychronakis, his name, and his public key `8EBD 8F30 8899 8AFF`. Social media links for Twitter and GitHub are also present. A green notification box states: "mikepo has an invitation available. If you know mikepo, you can ask them for an invitation to Keybase." Below the profile, there are two buttons: 'Encrypt' and 'Verify'. The bottom section is divided into two columns: 'Tracking (6)' and 'Trackers (6)'. The 'Tracking' column lists users: hargikas, mstamat, and gianluca_string. The 'Trackers' column lists users: hargikas, kontaxis, and mstamat. On the left, a code block shows instructions for using the 'keybase' command line tool.

Michalis Polychronakis

keybase.io/mikepo

8EBD 8F30 8899 8AFF

polychronakis tweet

polychronakis gist

miikepo has an invitation available
If you know mikepo, you can ask them for an invitation to Keybase.

Encrypt Verify

mikepo from the [command line](#)

```
# first
keybase join # if you're new, or
keybase login # if you're not.

# then
keybase push # if you already have a public key, or
keybase gen # if this is all new to you
```

Tracking (6)

- hargikas
- mstamat
- gianluca_string

Trackers (6)

- hargikas
- kontaxis
- mstamat

Best Practices

User education is important!

Pick long passwords (passphrases)

Never reuse the same password on different services

Never share passwords

No post-its

Use two-factor authentication when possible

Use a password manager

Not only for passwords! Also for “security” questions...

Use SSH keys instead of passwords