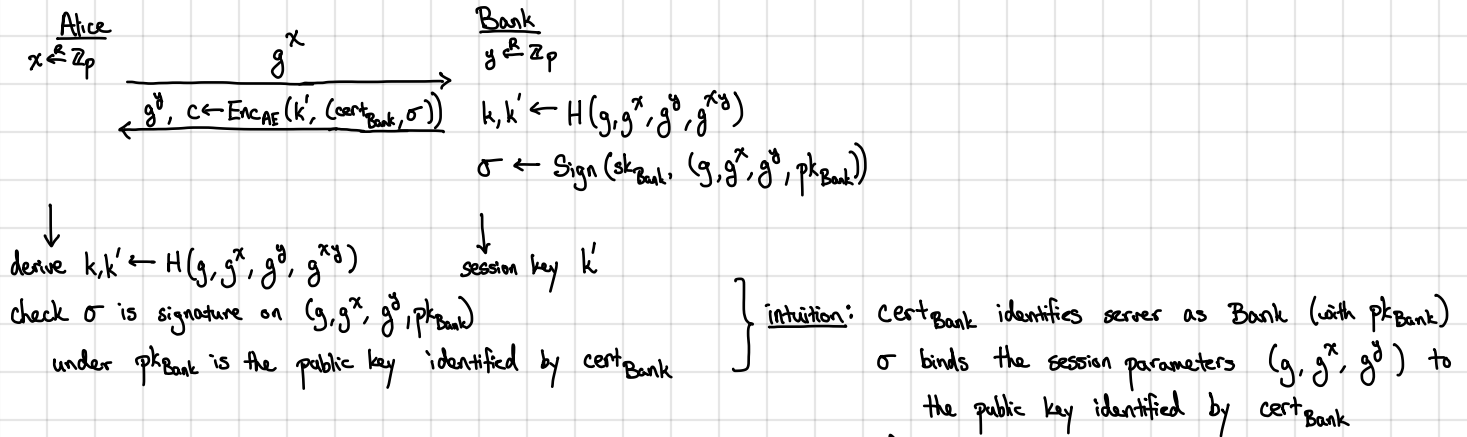


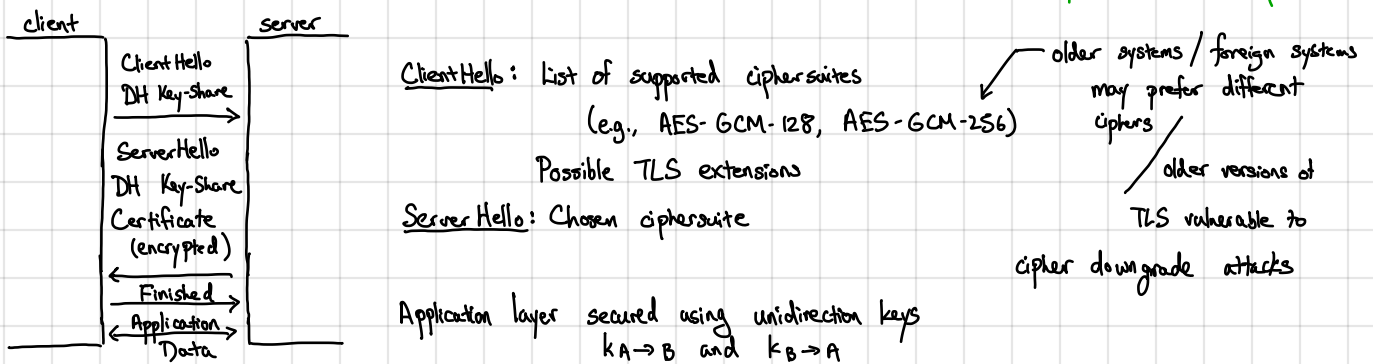
Basic flow of Diffie-Hellman based AKE:



End of protocol: Alice knows she is talking to Bank (but not vice versa!)

"one-sided AKE" - most common mode on the web

↳ Basis of TLS 1.3 handshake ("one-sided" AKE) **ALWAYS USE TLS 1.3 - Don't invent your own AKE protocol!**



In TLS 1.3, the only long-term secret on the server is a signing key. This is critical for achieving forward secrecy.

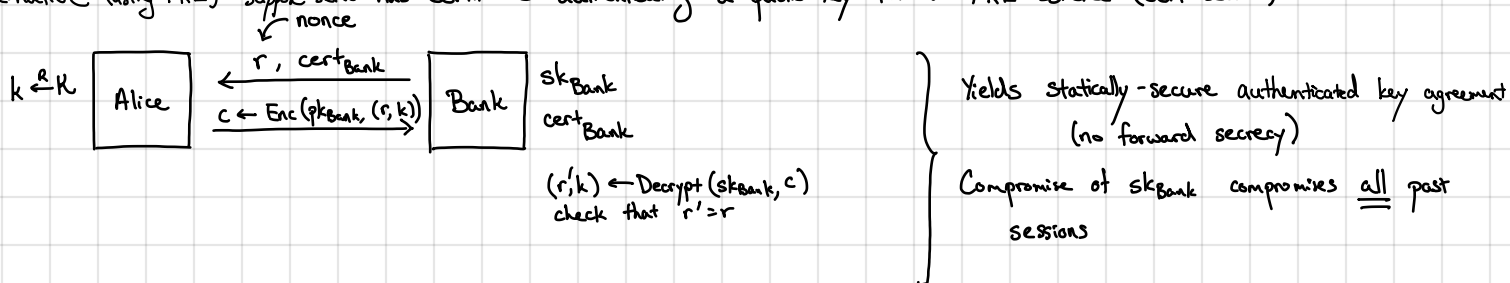
Forward secrecy: compromise of server in the future cannot affect secrecy of sessions in the past

↳ In TLS, server secret is a signing key - fresh Diffie-Hellman secret used for each session is fresh ("ephemeral")

Compromising signing key allows impersonation of server, but does not break secrecy of past sessions

↳ As we will see, not all AKE protocols provide forward secrecy

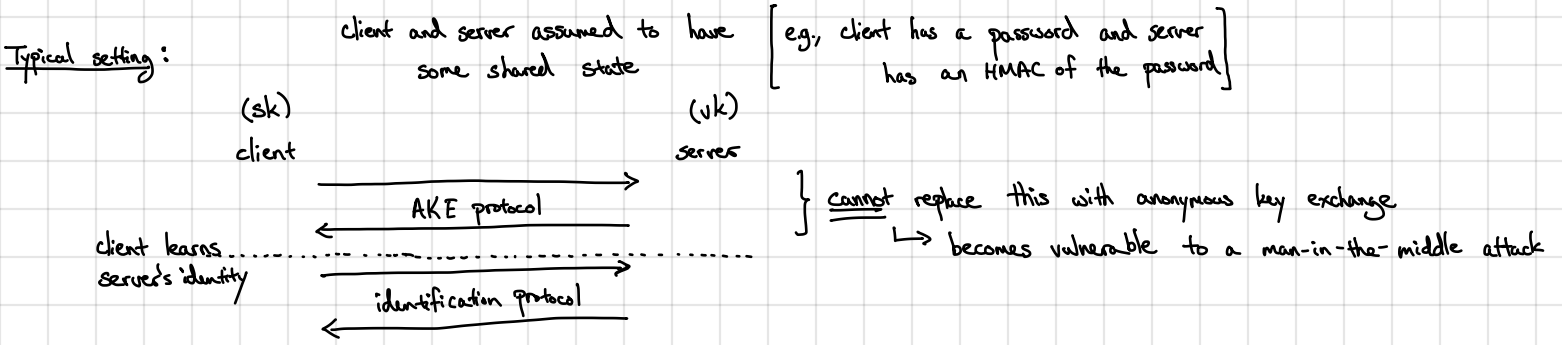
Alternative (using PKE): suppose server has certificate authenticating a public key for a PKE scheme (CCA-secure):



TLS 1.3 and authenticated key-exchange protocols on the Internet typically provide one-sided authentication (i.e., client learns id of the server, but not vice versa)

Question: how does the client authenticate to the server (without providing a certificate)


↳ e.g., how does client login to a web service?



Threat models: Adversary's goal is to authenticate to server

- Direct attack: adversary only sees vk and needs to authenticate (e.g., physical analogy: door lock - adversary can observe the lock, does not see the key sk)
- Eavesdropping attack: adversary gets to observe multiple interactions between honest client and the server (e.g., physical analogy: wireless car key - adversary observes communication between car key and car)
- Active attack: adversary can impersonate the server and interact with the honest client (e.g., physical analogy: fake ATM in the mall - honest clients interact directly with the adversary)


Simple (insecure) password-based protocol:



NEVER STORE PASSWORDS IN THE CLEAR!

Slightly better solution: hash the passwords before storing

server maintains mappings Alice $\mapsto H(\text{pwd}_{\text{Alice}})$
Bob $\mapsto H(\text{pwd}_{\text{Bob}})$
where H is a collision-resistant hash function



If passwords have high entropy, then hard to recover pwd from $H(\text{pwd})$ [by one-wayness of H]

↳ But not true in practice...

Users often choose weak passwords (e.g., 123456, password, 123456789, ...)

↳ With a dictionary of 360 million entries, can cover about 25% of user passwords
(3% choose 123456)

(10% choose among top 25 common passwords)

} Based on password hashes that have been leaked from compromised databases

Simple hashing vulnerable to "offline dictionary attack":

adversary computes table $(\text{pwd}, H(\text{pwd}))$ for common passwords — completely offline
given $H(\text{pwd})$, can now invert with a single lookup if pwd is contained in the database

for LinkedIn breach in 2012, attacker stole password file with ~6 million passwords

(all passwords hashed using single iteration of unsalted SHA-1) → 90% of passwords recovered in ~6 days!

Problem: One-time precomputation (computing the lookup table) can be reused to compromise many passwords

Overall cost of attack: $O(m+n)$ where m is the dictionary size and n is the number of passwords to attack

Defense #1: Salt passwords before hashing: namely when storing password pwd , sample salt $\overset{r}{\leftarrow} \{0,1\}^n$ and store
 $(\text{salt}, H(\text{salt} \parallel \text{pwd}))$ on the server

Note: Salt is a public value (needed for verification)

↑
typically, $n \geq 64$

Offline dictionary attack no longer effective since every salt value induces different set of hash values

Overall cost of dictionary attack: $O(mn)$ — need to re-hash dictionary for every salt

Defense #2: Use a slow hash function [SHA-1 is very fast — enables fast brute-force search]

- PBKDF2 (password-based key-derivation function): iterate a cryptographic hash function many times:

(or bcrypt)

$\text{PBKDF2}(\text{pwd}, \text{salt}) : H(H(\dots H(\text{salt} \parallel \text{pwd}) \dots))$

can use 100,000 or
1,000,000 iterations of SHA-256

honest user only needs to evaluate
hash function once per authentication;
adversary evaluates many times

Drawback: custom hardware can evaluate SHA-256 very fast

- scrypt (more recent: Argon2i): slow hash function that needs lots of memory (space) to evaluate

↳ custom hardware do not provide substantial savings (limiting factor is space, not compute)

Can also use a keyed hash function (e.g., HMAC with key stored in HSM)

↳ ensures adversary who does not know key cannot brute force at all!

Best practice: Always salt passwords

Always use a slow hash function (e.g., PBKDF2, scrypt) or keyed hash function or both!

$\$cur = \text{'password'}$

$\$cur = \text{md5}(\$cur)$ raw MD5 hash — not secure!

$\$salt = \text{randbytes}(20)$

$\$cur = \text{hmac_sha1}(\$cur, \$salt)$

$\$cur = \text{remote_hmac_sha256}(\$cur, \$secret)$

$\$cur = \text{scrypt}(\$cur, \$salt)$ slow hash function

$\$cur = \text{hmac_sha256}(\$cur, \$salt)$

Facebook password onion
(circa 2014)

↓
layers gradually added over time to
achieve better security
(and probably to avoid password
rehashing)

salted, keyed
hash function
(key on remote service)