# last time (1)

DNS time-to-lives and limited caching

mapping MAC addresses to IP addresses if know IP should be on network, ask everyone "is this yours?"

DHCP: broadcast "what IP/etc. to use"

network address translation large set of inside IP addresses map to small set of outside IP addresses table: (remote IP+port, outside IP+port)  $\rightarrow$  inside IP+port

# last time (2)

security goals: confidentiality, authenticity

encryption

E(key, data) = ciphertext cannot learn about data from ciphertext + E without key

```
message authentication codes

MAC(key, data) = tag

cannot find MAC(key, f(data)) without key

even with MAC(key, data), data
```

## quiz Q1

- 1:  $A{\rightarrow}$  A wifi router
- 2: A wifi router  $\rightarrow$  B wifi router (via wired network)
- 3: B wifi router  $\rightarrow$  B
- + again in reverse
- = 6 messages



frame addresses: machine on current network

packet address: ultimate (multi-step) destination



connections at transport layer

unless we change transport layer relies on being able to use network layer in same way

so need router-changing solution

## quiz Q5

best scheme would be encrypt+auth with key 2 (not an option)

B allows confirming 'guesses' about message using MAC

 $C{+}D$  allow reading decrypted message

## quiz Q6

best scheme would be encrypt+auth with key 2 (not an option)

- B allows confirming 'guesses' about message using MAC
- $C{+}D$  allow reading decrypted message

#### anon feedback

"This course has the best TAs of any CS class I have taken so far, they help you with the actual implementation as well as making sure you fully understand the overarching concepts the homeworks and coursework are supposed to teach you. They're really doing a great job!"

"I have really enjoyed the labs in this course thus far. ...unlike all other labs so far, the openmp lab left me feeling frustrated and confused. The crux of the issue is the write-up's nebulous connection between the parallel mapping and reduction stages. Like, do you do them sequentially or together? I still don't know. I think a fully worked-out example showing..."

hoped that pseudocode under first reduce strategy would help... providing worked out examples probably means more complex problem

#### exercise

suppose A, B have shared keys  $K_1, K_2$  assume attackers do not have keys

 $\mathsf{E}/\mathsf{D}=\mathsf{encrypt}/\mathsf{decrypt}\ \mathsf{function}$ 

A asks B to pay Sue \$100 by sending message with these parts: "2023-11-03: pay \$100"  $E(K_1, "2023-11-03 \text{ Sue"})$  $MAC(K_2, "2023-11-03 \text{ $100"})$ 

1. can eavesdropper learn: (a) who is being paid, (b) how much?

2. can machine-in-middle change: (a) who is being paid, (b) how much?

#### shared secrets impractical

problem: shared secrets usually aren't practical

need secure communication before I can do secure communication?

scaling problems

millions of websites  $\times$  billions of browsers = how many keys? hard to talk to new people

#### shared secrets impractical

problem: shared secrets usually aren't practical

need secure communication before I can do secure communication?

scaling problems millions of websites  $\times$  billions of browsers = how many keys? hard to talk to new people

#### shared secrets impractical

problem: shared secrets usually aren't practical

need secure communication before I can do secure communication?

#### scaling problems

millions of websites  $\times$  billions of browsers = how many keys? hard to talk to new people

will still need to have some sort of secure communication to setup!

because we need some way to know we aren't talking to attacker

will still need to have some sort of secure communication to setup!

because we need some way to know we aren't talking to attacker

but...

will still need to have some sort of secure communication to setup!

because we need some way to know we aren't talking to attacker but...

#### can be broadcast communication

don't need full new sets of keys for each web browser

will still need to have some sort of secure communication to setup!

because we need some way to know we aren't talking to attacker but...

can be broadcast communication don't need full new sets of keys for each web browser

only with smaller number of trusted authorities don't need to have keys for every website in advance

#### asymmetric encryption

we'll have two functions:

encrypt: PE(public key, message) = ciphertextdecrypt: PD(private key, ciphertext) = message

(public key, private key) = "key pair"

## key pairs

'private key' = kept secret usually not shared with *anyone* 

'public key' = safe to give to everyone usually some hard-to-reverse function of public key

concept will appear in some other cryptographic primitives

#### asymmetric encryption properties

functions:

encrypt: PE(public key, message) = ciphertext decrypt: PD(private key, ciphertext) = message

should have:

knowing PE, PD, the public key, and ciphertext shouldn't make it too easy to find message knowing PE, PD, the public key, ciphertext, and message shouldn't help in finding private key

#### secrecy properties with asymmetric

not going to be able to make things as hard as "try every possibly private key"

but going to make it impractical

like with symmetric encryption want to prevent recovery of *any info about message* 

also have some other attacks to worry about:

e.g. no info about key should be revealed based on our reactions to decrypting maliciously chosen ciphertexts

#### using asymmetric v symmetric

both:

```
use secret data to generate key(s)
```

asymmetric (AKA public-key) encryption one "keypair" per recipient private key kept by recipient public key sent to all potential senders encryption is one-way without private key

symmetric encryption

one key per (recipient + sender) secret key kept by recipient + sender if you can encrypt, you can decrypt



in advance: B generates private key + public key

in advance: B sends public key to A (and maybe others) securely

A computes PE(public key, 'The secret formula is...') = \*\*\*\*\*\*

send on network: A  $\rightarrow$  B: \*\*\*\*\*\*\*

B computes PD(private key, \*\*\*\*\*\*) = 'The secret formula is ...'

#### digital signatures

symmetric encryption : asymetric encryption :: message authentication codes : digital signatures

## digital signatures

pair of functions:

```
sign: S(\text{private key}, \text{message}) = \text{signature}
verify: V(\text{public key}, \text{signature}, \text{message}) = 1 ("yes, correct signature")
```

(public key, private key) = key pair (similar to asymmetric encryption)

public key can be shared with everyone knowing  $S,\,V,$  public key, message, signature doesn't make it too easy to find another message + signature so that  $V({\rm public key},{\rm other message},{\rm other signature})=1$ 



in advance: A generates private key + public key

in advance: A sends public key to B (and maybe others) securely

A computes S(private key, 'Please pay ...') = \*\*\*\*\*\*\*

```
send on network: A \rightarrow B: 'Please pay ...', *******
```

B computes V(public key, 'Please pay ...', \*\*\*\*\*\*) = 1

#### tools, but...

have building blocks, but less than straightforward to use

lots of issues from using building blocks poorly

start of art solution: formal proof sytems
 mathematical proof that attacker doing X implies encryption/MAC/etc.
 broken
 ideally a somewhat machine-checkable proof

(we aren't going to be that formal...)

#### replay attacks

...

- $\begin{array}{l} \mathsf{A}{\rightarrow}\mathsf{B}: \mbox{ Did you order lunch? [signature 1 by A]} \\ & \mbox{ signature 1 by } \mathsf{A} = \mbox{ Sign}(\mathsf{A}\text{'s private signing key, "Did you order lunch?")} \\ & \mbox{ will check with Verify}(\mathsf{A}\text{'s public key, signature 1 by A, "Did you order lunch?")} \end{array}$
- $\begin{array}{l} B{\rightarrow}A: \mbox{ Yes. [signature 1 by B]} \\ signature 1 by B = \mbox{ Sign(B's private key, "Yes.")} \\ will check with \mbox{ Verify(B's public key, signature 1 by B, "Yes.")} \end{array}$
- $\begin{array}{l} A{\rightarrow}B: \mbox{ Vegetarian}? \ [signature \ 2 \ by \ A] \\ B{\rightarrow}A: \ No, \ not \ this \ time. \ [signature \ 2 \ by \ B] \end{array}$

 $A \rightarrow B$ : There's a guy at the door, says he's here to repair the AC. Should I let him in? [signature N by A]

#### replay attacks

 $A \rightarrow B$ : Did you order lunch? [signature 1 by A]  $B \rightarrow A$ : Yes. [signature 1 by B]  $A \rightarrow B$ : Vegetarian? [signature 2 by A]  $B \rightarrow A$ : No, not this time. [signature 2 by B]

•••

 $A \rightarrow B$ : There's a guy at the door, says he's here to repair the AC. Should I let him in? [signature ? by A]

how can attacker hijack the reponse to A's inquiry?

#### replay attacks

...

 $A \rightarrow B$ : Did you order lunch? [signature 1 by A]  $B \rightarrow A$ : Yes. [signature 1 by B]  $A \rightarrow B$ : Vegetarian? [signature 2 by A]  $B \rightarrow A$ : No, not this time. [signature 2 by B]

 $A \rightarrow B$ : There's a guy at the door, says he's here to repair the AC. Should I let him in? [signature ? by A]

how can attacker hijack the reponse to A's inquiry?

as an attacker, I can copy/paste B's earlier message! just keep the same signature, so it can be verified! Verify(B's public key, "Yes.", signature 2 from B) = 1

# nonces (1)

one solution to replay attacks:

 $A \rightarrow B$ : #1 Did you order lunch? [signature 1 from A] signature from A = Sign(A's private key, "#1 Did you order lunch?")

 $\begin{array}{l} B{\rightarrow}A{:}~\#1~Yes.~[signature~1~from~B]\\ A{\rightarrow}B{:}~\#2~Vegetarian?~[signature~2~from~A]\\ B{\rightarrow}A{:}~\#2~No,~not~this~time.~[signature~2~from~B] \end{array}$ 

•••

 $A \rightarrow B$ : #54 There's a guy at the door, says he's here to repair the AC. Should I let him in? [signature ? from A]

(assuming A actually checks the numbers)

# nonces (2)

...

another solution to replay attacks:

 $B \rightarrow A$ : [next number #91523] [signature from B]  $A \rightarrow B$ : #91523 Did you order lunch? [next number #90382] [signature from A]  $B \rightarrow A$ : #90382 Yes. [next number #14578] [signature from B]

 $A \rightarrow B$ : #6824 There's a guy at the door, says he's here to repair the AC. Should I let him in? [next number #36129][signature from A]

(assuming A actually checks the numbers)

## replay attacks (alt)

 $M \rightarrow B: \#50 \text{ Did you order lunch? [signature by M]} B \rightarrow M: \#50 \text{ Yes. [signature intended for M by B]}$ 

 $A \rightarrow B$ : #50 There's a guy at the door, says he's here to repair the AC. Should I let him in? [signature ? by A]

how can M hijack the reponse to A's inquiry?

## replay attacks (alt)

 $M \rightarrow B: \#50 \text{ Did you order lunch? [signature by M]} B \rightarrow M: \#50 \text{ Yes. [signature intended for M by B]}$ 

 $A \rightarrow B$ : #50 There's a guy at the door, says he's here to repair the AC. Should I let him in? [signature ? by A]

how can M hijack the reponse to A's inquiry?

as an attacker, I can copy/paste B's earlier message! just keep the same signature, so it can be verified! Verify(B's public key, "#50 Yes.", signature intended for M by B) = 1

#### confusion about who's sending?

in addition to nonces, either

write down more who is sending + other context so message can't be reused and/or

use unique set of keys for each principal you're talking to

with symmetric encryption, also "reflection attacks" A sends message to B, attacker sends A's message back to A as if it's from B

#### other attacks without breaking math

#### **TLS** state machine attack

from https://mitls.org/pages/attacks/SMACK

protocol:

step 1: verify server identity
step 2: receive messages from server

attack:

if server sends "here's your next message", instead of "here's my identity" then broken client ignores verifying server's identity

#### **Matrix vulnerabilties**

one example from https://nebuchadnezzar-megolm.
github.io/static/paper.pdf

system for confidential multi-user chat

```
protocol + goals:
```

each device (my phone, my desktop) has public key to talk to me, you verify one of my public keys to add devices, my client can forward my other devices' public keys

bug:

when receiving new keys, clients did not check who they were forwarded from correctly  $% \left( {{{\mathbf{r}}_{\mathbf{r}}}_{\mathbf{r}}} \right)$ 

#### on the lab

## getting public keys?

browser talking to websites needs public keys of every single website?

not really feasible, but...

#### certificate idea

let's say A has B's public key already.

if C wants B's public key and knows A's already:

A can generate "certificate" for B: "B's public key is XXX" AND Sign(A's private key, "B's public key is XXX")

B send copy of their "certificate" to C (most common idea)

if C trusts A, now C has B's public key if C does not trust A, well, can't trust this either

#### certificate idea

let's say A has B's public key already.

if C wants B's public key and knows A's already:

A can generate "certificate" for B: "B's public key is XXX" AND Sign(A's private key, "B's public key is XXX")

B send copy of their "certificate" to C (most common idea)

if C trusts A, now C has B's public key if C does not trust A, well, can't trust this either

#### certificate idea

let's say A has B's public key already.

if C wants B's public key and knows A's already:

A can generate "certificate" for B: "B's public key is XXX" AND Sign(A's private key, "B's public key is XXX")

B send copy of their "certificate" to C (most common idea)

if C trusts A, now C has B's public key if C does not trust A, well, can't trust this either

#### certificate authorities

websites (and others) go to *certificates authorities* (CA) with their public key

certificate authorities sign messages like: "The public key for foo.com is XXX."

signed message called certificate

send certificates to browsers to verify identity website can forward certificate instead of browser contacting CA directly

#### certificate authorities

websites (and others) go to *certificates authorities* (CA) with their public key

certificate authorities sign messages like: "The public key for foo.com is XXX."

signed message called certificate

send certificates to browsers to verify identity website can forward certificate instead of browser contacting CA directly

## example web certificate (1)

Version: 3 (0x2)
Serial Number: 7b:df:f6:ae:2e:d7:db:74:d3:c5:77:ac:bc:44:bf:1b
Signature Algorithm: sha256WithRSAEncryption
Issuer:

countryName	=	US
stateOrProvinceName	=	MI
localityName	=	Ann Arbor
organizationName	=	Internet2
organizationalUnitName	=	InCommon
commonName	=	InCommon RSA Server CA
Validity		
Not Before: Apr 25 00:00:0	00	2023 GMT
Not After : Apr 24 23:59:5	59	2024 GMT
Subject:		
countryName	=	US
stateOrProvinceName	=	Virginia
organizationName	=	University of Virginia
commonName	=	canvas.its.virginia.edu

• • • •

X509v3 extensions:

. . . .

X509v3 Subject Alternative Name: DNS:canvas.its.virginia.edu

# example web certificate (2)

```
. . . .
    Subject Public Key Info:
        Public Key Algorithm: rsaEncryption
            RSA Public-Key: (2048 bit)
            Modulus:
                00:a2:fb:5a:fb:2d:d2:a7:75:7e:eb:f4:e4:d4:6c:
                94:be:91:a8:6a:21:43:b2:d5:9a:48:b0:64:d9:f7:
                f1:88:fa:50:cf:d0:f3:3d:8b:cc:95:f6:46:4b:42:
. . . .
Signature Algorithm: sha256WithRSAEncryption
Signature Value:
    24:3a:67:c8:0d:ef:eb:8c:eb:ba:8f:d5:11:d2:1e:ea:44:eb:
    fe:af:93:7d:d9:4a:2b:44:a3:7f:47:50:aa:d1:b3:9c:a8:a8:
```

```
. . . .
```

#### certificate chains

- That certificate signed by "InCommon RSA Server CA"
- $\mathsf{C}\mathsf{A}=\mathsf{certificate} \text{ authority}$
- so their public key, comes with my OS/browser? not exactly...
- they have their own certificate signed by "USERTrust RSA Certification Authority"
- and their public key comes with your OS/browser?

(but both CAs now operated by UK-based Sectigo)

#### certificate hierarchy



#### certificate hierarchy **USERTrust RSA** GlobalSign Root CA Certification Authority operated by GlobalSign nv-sa originally operated by USERTrust, Inc. subsid. of GMO Internet Group since 2007 acquired by Comodo, Inc (2004) Comodo's CA division renamed Sectigo (2018) GTS Root R1 - - operated by Google Trust Services LLC InCommon RSA Server CA ... operated by Sectigo GTS CA 1C3 on behalf of the Internet2 (not-for-profit) ... some "trust anchors" included with browsers and OSes (for GTS Root R1, only more recent browsers/OSes)

#### how many trust anchors?

Mozilla Firefox (as of 27 Feb 2023) 155 trust anchors operated by 55 distinct entities

#### Microsoft Windows (as of 27 Feb 2023) 237 trust anchors operated by 86 distinct entities

## public-key infrastructure

ecosystem with certificate authorities and certificates for everyone

called "public-key infrastructure"

several of these:

for verifying identity of websites for verifying origin of domain name records (kind-of) for verifying origin of applications in some OSes/app stores/etc. for encrypted email in some organizations

•••



exercise: how should website certificates verify identity?

#### how do certificate authorities verify

for web sites, set by CA/Browser Forum

organization of:

everyone who ships code with list of valid certificate authorities Apple, Google, Microsoft, Mozilla, Opera, Cisco, Qihoo 360, Brave, ... certificate authorities

decide on rules ("baseline requirements") for what CAs do

#### BR domain name identity validation

options involve CA choosing random value and:

sending it to domain contact (with domain registrar) and receive response with it, or

observing it placed in DNS or website or sent from server in other specific way

exercise: problems this doesn't deal with?

keep their private keys in tamper-resistant hardware

maintain publicly-accessible database of *revoked* certificates some browsers check these, sometimes

certificate transparency

public logs of every certificate issued some browsers reject non-logged certificates so you can tell if bad certificate exists for your website

'CAA' records in the domain name system can indicate which CAs are allowed to issue certificates in DNS (but CAs apparently not required to use DNSSEC (certificate infrastructure for signing domain name records) when looking this up)

keep their private keys in tamper-resistant hardware

maintain publicly-accessible database of *revoked* certificates some browsers check these, sometimes

certificate transparency

public logs of every certificate issued some browsers reject non-logged certificates so you can tell if bad certificate exists for your website

'CAA' records in the domain name system can indicate which CAs are allowed to issue certificates in DNS (but CAs apparently not required to use DNSSEC (certificate infrastructure for signing domain name records) when looking this up)

keep their private keys in tamper-resistant hardware

maintain publicly-accessible database of *revoked* certificates some browsers check these, sometimes

#### certificate transparency

public logs of every certificate issued some browsers reject non-logged certificates so you can tell if bad certificate exists for your website

'CAA' records in the domain name system can indicate which CAs are allowed to issue certificates in DNS (but CAs apparently not required to use DNSSEC (certificate infrastructure for signing domain name records) when looking this up)

keep their private keys in tamper-resistant hardware

maintain publicly-accessible database of *revoked* certificates some browsers check these, sometimes

certificate transparency

public logs of every certificate issued some browsers reject non-logged certificates so you can tell if bad certificate exists for your website

'CAA' records in the domain name system

can indicate which CAs are allowed to issue certificates in DNS (but CAs apparently not required to use DNSSEC (certificate infrastructure for signing domain name records) when looking this up)

## backup slides

#### cryptographic hash uses

find shorter 'summary' to substitute for data what hashtables use them for, but... we care that adversaries can't cause collisions!

#### cryptographic hash uses

find shorter 'summary' to substitute for data what hashtables use them for, but... we care that adversaries can't cause collisions!

deal with message limits in signatures/etc.

password hashing — but be careful! [next slide]

constructing message authentication codes hash message + secret info (+ some other details)