

CSC358 Intro. to Computer Networks

Lecture 11: VLAN, MPLS, Network Security

Amir H. Chinaei, Winter 2016

ahchinaei@cs.toronto.edu
<http://www.cs.toronto.edu/~ahchinaei/>

Many slides are (inspired/adapted) from the above source
 © all material copyright; all rights reserved for the authors




Office Hours: T 17:00–18:00 R 9:00–10:00 BA4222

TA Office Hours: W 16:00–17:00 BA3201 R 10:00–11:00 BA7172

csc358ta@cdf.toronto.edu
<http://www.cs.toronto.edu/~ahchinaei/teaching/2016jan/csc358/>

VLANs: motivation



consider:

- ❖ CS user moves office to EE, but wants connect to CS switch?
- ❖ single broadcast domain:
 - all layer-2 broadcast traffic (ARP, DHCP, unknown location of destination MAC address) must cross entire LAN
 - security/privacy, efficiency issues


Link Layer 5-2

VLANs

Virtual Local Area Network

switch(es) supporting VLAN capabilities can be configured to define multiple **virtual** LANS over single physical LAN infrastructure.

port-based VLAN: switch ports grouped (by switch management software) so that **single** physical switch



Link Layer 5-3

Port-based VLAN


❖ **traffic isolation:** frames to/from ports 1-8 can **only** reach ports 1-8

- can also define VLAN based on MAC addresses of endpoints, rather than switch port

❖ **dynamic membership:** ports can be dynamically assigned among VLANs


❖ **forwarding between VLANs:** done via routing (just as with separate switches)

- in practice vendors sell combined switches plus routers



Link Layer 5-4

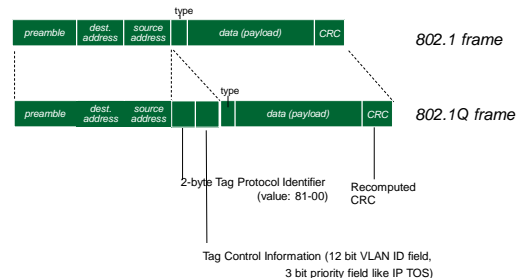
VLANs spanning multiple switches



- ❖ **trunk port:** carries frames between VLANs defined over multiple physical switches
 - frames forwarded within VLAN between switches can't be vanilla 802.1 frames (must carry VLAN ID info)
 - 802.1q protocol adds/removed additional header fields for frames forwarded between trunk ports

Link Layer 5-5


802.1Q VLAN frame format



Link Layer 5-6

Multiprotocol label switching (MPLS)

- ❖ initial goal: high-speed IP forwarding using fixed length label (instead of IP address)
 - fast lookup using fixed length identifier (rather than longest prefix matching)
 - borrowing ideas from Virtual Circuit (VC) approach
 - but IP datagram still keeps IP address!




Link Layer 5-7

MPLS capable routers

- ❖ a.k.a. label-switched router
- ❖ forward packets to outgoing interface based only on label value (*don't inspect IP address*)
 - MPLS forwarding table distinct from IP forwarding tables
- ❖ **flexibility**: MPLS forwarding decisions can *differ* from those of IP
 - use destination *and* source addresses to route flows to same destination differently (traffic engineering)
 - re-route flows quickly if link fails: pre-computed backup paths (useful for VoIP)

Link Layer 5-8

MPLS versus IP paths




- ❖ **IP routing**: path to destination determined by destination address alone



Link Layer 5-9

MPLS versus IP paths




- ❖ **IP routing**: path to destination determined by destination address alone
- ❖ **MPLS routing**: path to destination can be based on source *and* dest. address
 - **fast reroute**: precompute backup routes in case of link failure



Link Layer 5-10

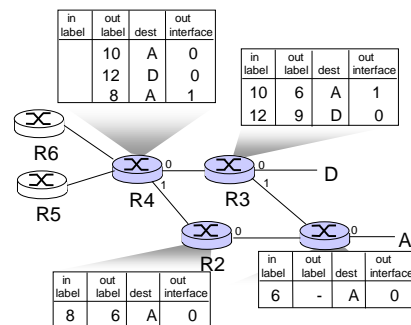
MPLS signaling

- ❖ modify OSPF, IS-IS link-state protocols to carry info used by MPLS routing
 - e.g., link bandwidth, amount of "reserved" link bandwidth
- ❖ entry MPLS router uses **RSVP-TE signaling protocol** to set up MPLS forwarding at downstream routers



Link Layer 5-11

MPLS forwarding tables



Link Layer 5-12

Data center networks

- ❖ 10's to 100's of thousands of hosts, often closely coupled, in close proximity:
 - e-business (e.g. Amazon)
 - content-servers (e.g., YouTube, Akamai, Apple, Microsoft)
 - search engines, data mining (e.g., Google)
- ❖ challenges:
 - multiple applications, each serving massive numbers of clients
 - managing/balancing load, avoiding processing, networking, data bottlenecks




Inside a 40-ft Microsoft container, Chicago data center

Link Layer 5-13

Data center networks

load balancer: application-layer routing


- receives external client requests
- directs workload within data center
- returns results to external client (hiding data center internals from client)



Link Layer 5-14

Data center networks

- ❖ rich interconnection among switches, racks:
 - increased throughput between racks (multiple routing paths possible)
 - increased reliability via redundancy



Link Layer 5-16

Link layer, LANs: outline

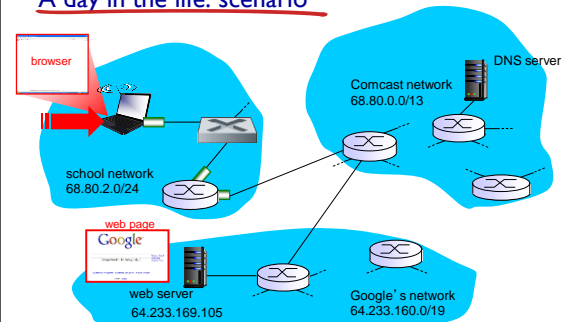
- 5.1 introduction, services
- 5.2 error detection, correction
- 5.3 multiple access protocols
- 5.4 LANs
 - addressing, ARP
 - Ethernet
 - switches
 - VLANs
- 5.5 link virtualization: MPLS
- 5.6 data center networking
- 5.7 a day in the life of a web request

Synthesis: a day in the life of a web request

- ❖ journey down protocol stack complete!
 - application, transport, network, link
- ❖ putting-it-all-together: synthesis!
 - **goal:** identify, review, understand protocols (at all layers) involved in seemingly simple scenario: requesting www page
 - **scenario:** student attaches laptop to campus network, requests/receives www.google.com


Link Layer 5-17

A day in the life: scenario



Link Layer 5-18


A day in the life... connecting to the Internet



- ❖ connecting laptop needs to get its own IP address, addr of first-hop router, addr of DNS server: use **DHCP**
- ❖ DHCP request **encapsulated** in **UDP**, encapsulated in **IP**, encapsulated in **802.3** Ethernet
- ❖ Ethernet frame **broadcast** (dest: FFFFFFFFFF) on LAN, received at router running **DHCP** server
- ❖ Ethernet **demuxed** to IP demuxed, UDP demuxed to DHCP

Link Layer 5-19

A day in the life... connecting to the Internet




- ❖ DHCP server formulates **DHCP ACK** containing client's IP address, IP address of first-hop router for client, name & IP address of DNS server
- ❖ encapsulation at DHCP server, frame forwarded (**switch learning**) through LAN, demultiplexing at client
- ❖ DHCP client receives DHCP ACK reply

Client now has IP address, knows name & addr of DNS server, IP address of its first-hop router

Link Layer 5-20

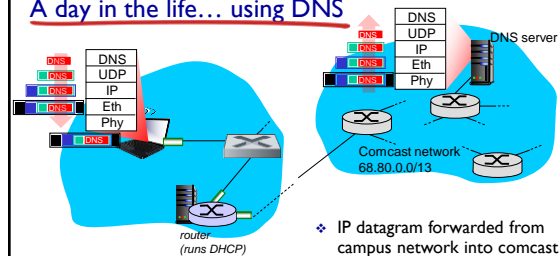
A day in the life... ARP (before DNS, before HTTP)



- ❖ before sending **HTTP** request, need IP address of **www.google.com**: **DNS**
- ❖ DNS query created, encapsulated in UDP, encapsulated in IP, encapsulated in Eth. To send frame to router, need MAC address of router interface: **ARP**
- ❖ **ARP query** broadcast, received by router, which replies with **ARP reply** giving MAC address of router interface
- ❖ client now knows MAC address of first hop router, so can now send frame containing DNS query

Link Layer 5-21

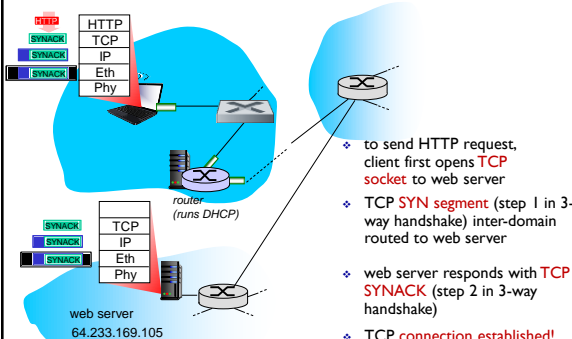
A day in the life... using DNS



- ❖ IP datagram containing DNS query forwarded via LAN switch from client to 1st hop router
- ❖ IP datagram forwarded from campus network into comcast network, routed (tables created by **RIP, OSPF, IS-IS** and/or **BGP** routing protocols) to DNS server
- ❖ demux'ed to DNS server
- ❖ DNS server replies to client with IP address of **www.google.com**

Link Layer 5-22

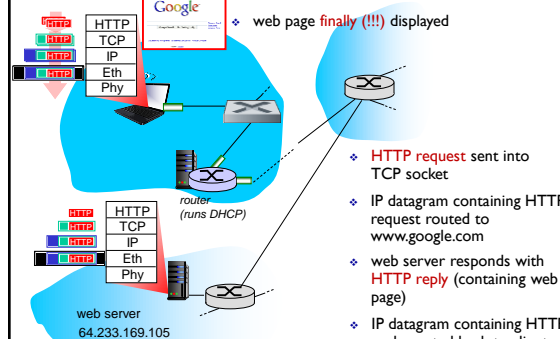
A day in the life... TCP connection carrying HTTP



- ❖ to send HTTP request, client first opens **TCP socket** to web server
- ❖ **TCP SYN segment** (step 1 in 3-way handshake) inter-domain routed to web server
- ❖ web server responds with **TCP SYNACK** (step 2 in 3-way handshake)
- ❖ **TCP connection established!**

Link Layer 5-23

A day in the life... HTTP request/reply



❖ web page **finally (!!!)** displayed

- ❖ **HTTP request** sent into TCP socket
- ❖ IP datagram containing HTTP request routed to **www.google.com**
- ❖ web server responds with **HTTP reply** (containing web page)
- ❖ IP datagram containing HTTP reply routed back to client

Link Layer 5-24

Chapter 8: Network Security

Chapter goals:

- ❖ understand principles of network security:
 - cryptography and its *many* uses beyond “confidentiality”
 - authentication
 - message integrity
- ❖ security in practice:
 - firewalls and intrusion detection systems
 - security in application, transport, network, link layers

Network Security

8-25

What is network security?

confidentiality: only sender, intended receiver should “understand” message contents

- sender encrypts message
- receiver decrypts message

authentication: sender, receiver want to confirm identity of each other

message integrity: sender, receiver want to ensure message not altered (in transit, or afterwards) without detection


access and availability: services must be accessible and available to users

Network Security

8-26

Friends and enemies: Alice, Bob, Trudy

- ❖ well-known in network security world
- ❖ Bob, Alice (lovers!) want to communicate “securely”
- ❖ Trudy (intruder) may intercept, delete, add messages



Network Security

Who might Bob, Alice be?

- ❖ ... well, *real-life* Bobs and Alices!
- ❖ Web browser/server for electronic transactions (e.g., on-line purchases)
- ❖ on-line banking client/server
- ❖ DNS servers
- ❖ routers exchanging routing table updates
- ❖ other examples?

Network Security

8-28

There are bad guys (and girls) out there!

Q: What can a “bad guy” do?


A: A lot! See section 1.6

- **eavesdrop:** intercept messages
- actively **insert** messages into connection
- **impersonation:** can fake (spoof) source address in packet (or any field in packet)
- **hijacking:** “take over” ongoing connection by removing sender or receiver, inserting himself in place
- **denial of service:** prevent service from being used by others (e.g., by overloading resources)

Network Security

8-29

The language of cryptography



m plaintext message

$K_A(m)$ ciphertext, encrypted with key K_A

$m = K_B(K_A(m))$


Network Security

Breaking an encryption scheme

- ❖ **cipher-text only attack:** Trudy has ciphertext she can analyze
- ❖ **two approaches:**
 - brute force: search through all keys
 - statistical analysis
- ❖ **known-plaintext attack:** Trudy has plaintext corresponding to ciphertext
 - e.g., in monoalphabetic cipher, Trudy determines pairings for a,l,i,c,e,b,o.
- ❖ **chosen-plaintext attack:** Trudy can get ciphertext for chosen plaintext

Network Security

Symmetric key cryptography



symmetric key crypto: Bob and Alice share same (symmetric) key: K_S

- ❖ e.g., key is knowing substitution pattern in mono alphabetic substitution cipher

Q: how do Bob and Alice agree on key value?

Network Security

Simple encryption scheme

substitution cipher: substituting one thing for another

- monoalphabetic cipher: substitute one letter for another

plaintext: abcdefghijklmnopqrstuvwxyz
 ↓ ↓
 ciphertext: mnbvcxzasdfghjklpoiuytrewq

e.g.: Plaintext: bob. i love you. alice
 ciphertext: nkn. s gktc wky. mgsbc

Encryption key: mapping from set of 26 letters to set of 26 letters

Network Security

A more sophisticated encryption approach

- ❖ n substitution ciphers, M_1, M_2, \dots, M_n
 - ❖ cycling pattern:
 - e.g., $n=4$: M_1, M_3, M_4, M_3, M_2 ; $M_1, M_3, M_4, M_3, M_2, \dots$
 - ❖ for each new plaintext symbol, use subsequent substitution pattern in cyclic pattern
 - dog: d from M_1 , o from M_3 , g from M_4
- Encryption key:** n substitution ciphers, and cyclic pattern
- key need not be just n-bit pattern

Network Security

Symmetric key crypto: DES

DES: Data Encryption Standard


- ❖ US encryption standard [NIST 1993]
- ❖ 56-bit symmetric key, 64-bit plaintext input
- ❖ block cipher with cipher block chaining
- ❖ how secure is DES?
 - DES Challenge: 56-bit-key-encrypted phrase decrypted (brute force) in less than a day
 - no known good analytic attack
- ❖ making DES more secure:
 - 3DES: encrypt 3 times with 3 different keys

Network Security

Symmetric key crypto: DES

DES operation

initial permutation
 16 identical "rounds" of function application, each using different 48 bits of key
 final permutation



Network Security

AES: Advanced Encryption Standard

- ❖ symmetric-key NIST standard, replaced DES (Nov 2001)
- ❖ processes data in 128 bit blocks
- ❖ 128, 192, or 256 bit keys
- ❖ brute force decryption (try each key) taking 1 sec on DES, takes 149 trillion years for AES

Network Security

Public Key Cryptography

symmetric key crypto

- ❖ requires sender, receiver know shared secret key
- ❖ Q: how to agree on key in first place (particularly if never "met")?


public key crypto

- ❖ radically different approach [Diffie-Hellman76, RSA78]
- ❖ sender, receiver do *not* share secret key
- ❖ *public* encryption key known to *all*
- ❖ *private* decryption key known only to receiver



Network Security

Public key cryptography



Network Security

Public key encryption algorithms

requirements:

- ① need $K_B^+(\cdot)$ and $K_B^-(\cdot)$ such that $K_B^-(K_B^+(m)) = m$
- ② given public key K_B^+ , it should be impossible to compute private key K_B^-

RSA: Rivest, Shamir, Adelson algorithm

Network Security

Prerequisite: modular arithmetic

- ❖ $x \bmod n$ = remainder of x when divide by n
- ❖ facts:
 - $[(a \bmod n) + (b \bmod n)] \bmod n = (a+b) \bmod n$
 - $[(a \bmod n) - (b \bmod n)] \bmod n = (a-b) \bmod n$
 - $[(a \bmod n) * (b \bmod n)] \bmod n = (a*b) \bmod n$
- ❖ thus $(a \bmod n)^d \bmod n = a^d \bmod n$
- ❖ example: $x=14, n=10, d=2$:
 - $(x \bmod n)^d \bmod n = 4^2 \bmod 10 = 6$
 - $x^d = 14^2 = 196 \quad x^d \bmod 10 = 6$

Network Security

RSA: getting ready

- ❖ message: just a bit pattern
- ❖ bit pattern can be uniquely represented by an integer number
- ❖ thus, encrypting a message is equivalent to encrypting a number.

example:

- ❖ $m=10010001$. This message is uniquely represented by the decimal number 145.
- ❖ to encrypt m , we encrypt the corresponding number, which gives a new number (the ciphertext).

Network Security

RSA: Creating public/private key pair

1. choose two large prime numbers p, q .
(e.g., 1024 bits each)
2. compute $n = pq$, $z = (p-1)(q-1)$
3. choose e (with $e < n$) that has no common factors with z (e, z are “relatively prime”).
4. choose d such that $ed-1$ is exactly divisible by z .
(in other words: $ed \bmod z = 1$).

5. public key is (n, e) , private key is (n, d) .
 K_B^+ K_B^-

Network Security

RSA: encryption, decryption

0. given (n, e) and (n, d) as computed above
1. to encrypt message m ($< n$), compute
 $c = m^e \bmod n$
2. to decrypt received bit pattern, c , compute
 $m = c^d \bmod n$

$$\text{magic happens! } m = (\underbrace{m^e \bmod n}_c)^d \bmod n$$

Network Security

RSA example:

Bob chooses $p=5, q=7$. Then $n=35, z=24$.
 $e=5$ (so e, z relatively prime).
 $d=29$ (so $ed-1$ exactly divisible by z).

encrypting 8-bit messages.

encrypt: bit pattern m m^e $c = m^e \bmod n$
 00001100 12 24832 17

decrypt: c c^d $m = c^d \bmod n$
 17 481968572106750915091411825223071697 12

Network Security

Why does RSA work?

- ❖ must show that $c^d \bmod n = m$
where $c = m^e \bmod n$
- ❖ fact: for any x and y : $x^y \bmod n = x^{(y \bmod z)} \bmod n$
 - where $n = pq$ and $z = (p-1)(q-1)$
- ❖ thus,

$$\begin{aligned} c^d \bmod n &= (m^e \bmod n)^d \bmod n \\ &= m^{ed} \bmod n \\ &= m^{(ed \bmod z)} \bmod n \\ &= m^1 \bmod n \\ &= m \end{aligned}$$

Network Security

RSA: another important property

The following property will be *very* useful later:

$$K_B^-(K_B^+(m)) = m = K_B^+(K_B^-(m))$$

use public key first,
followed by
private key

use private key
first, followed by
public key

result is the same!

Network Security

Why $K_B^-(K_B^+(m)) = m = K_B^+(K_B^-(m))$?

follows directly from modular arithmetic:

$$\begin{aligned} (m^e \bmod n)^d \bmod n &= m^{ed} \bmod n \\ &= m^{de} \bmod n \\ &= (m^d \bmod n)^e \bmod n \end{aligned}$$

Network Security

Why is RSA secure?

- ❖ suppose you know Bob's public key (n, e) . How hard is it to determine d ?
- ❖ essentially need to find factors of n without knowing the two factors p and q
 - fact: factoring a big number is hard

Network Security

RSA in practice: session keys

- ❖ exponentiation in RSA is computationally intensive
- ❖ DES is at least 100 times faster than RSA
- ❖ use public key crypto to establish secure connection, then establish second key – symmetric session key – for encrypting data

session key, K_S

- ❖ Bob and Alice use RSA to exchange a symmetric key K_S
- ❖ once both have K_S , they use symmetric key cryptography

Network Security

Chapter 8 roadmap


- 8.1 What is network security?
- 8.2 Principles of cryptography
- 8.3 Message integrity, *authentication*
- 8.4 Securing e-mail
- 8.5 Securing TCP connections: SSL
- 8.6 Network layer security: IPsec
- 8.7 Securing wireless LANs
- 8.8 Operational security: firewalls and IDS

Network Security

Authentication

Goal: Bob wants Alice to “prove” her identity to him

Protocol ap1.0: Alice says “I am Alice”




Failure scenario??

Network Security

Authentication

Goal: Bob wants Alice to “prove” her identity to him

Protocol ap1.0: Alice says “I am Alice”

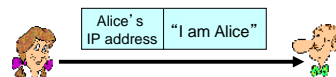


in a network,
Bob can not “see” Alice,
so Trudy simply declares
herself to be Alice

Network Security

Authentication: another try

Protocol ap2.0: Alice says “I am Alice” in an IP packet containing her source IP address




Failure scenario??

Network Security


Authentication: another try

Protocol ap2.0: Alice says "I am Alice" in an IP packet containing her source IP address



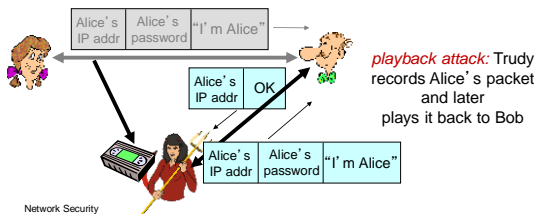
Authentication: another try

Protocol ap3.0: Alice says "I am Alice" and sends her secret password to "prove" it.



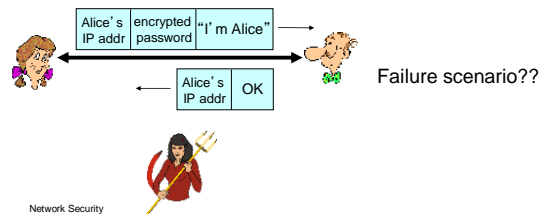
Authentication: another try

Protocol ap3.0: Alice says "I am Alice" and sends her secret password to "prove" it.



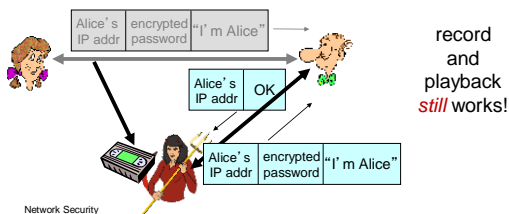
Authentication: yet another try

Protocol ap3.1: Alice says "I am Alice" and sends her **encrypted** secret password to "prove" it.




Authentication: yet another try

Protocol ap3.1: Alice says "I am Alice" and sends her **encrypted** secret password to "prove" it.



Authentication: yet another try

Goal: avoid playback attack
nonce: number (R) used only *once-in-a-lifetime*
ap4.0: to prove Alice "live", Bob sends Alice **nonce**, R. Alice must return R, encrypted with shared secret key




Authentication: ap5.0

ap4.0 requires shared symmetric key

❖ can we authenticate using public key techniques?

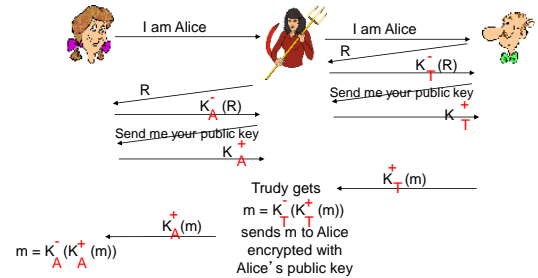
ap5.0: use nonce, public key cryptography



Network Security

ap5.0: security hole

man (or woman) in the middle attack: Trudy poses as Alice (to Bob) and as Bob (to Alice)



Network Security

ap5.0: security hole

man (or woman) in the middle attack: Trudy poses as Alice (to Bob) and as Bob (to Alice)



difficult to detect:

- ❖ Bob receives everything that Alice sends, and vice versa. (e.g., so Bob, Alice can meet one week later and recall conversation!)
- ❖ problem is that Trudy receives all messages as well!

Network Security

Digital signatures

cryptographic technique analogous to hand-written signatures:


- ❖ sender (Bob) digitally signs document, establishing he is document owner/creator.
- ❖ *verifiable, nonforgeable*: recipient (Alice) can prove to someone that Bob, and no one else (including Alice), must have signed document

Network Security

Digital signatures

simple digital signature for message m:

- ❖ Bob signs m by encrypting with his private key K_B^- , creating "signed" message, $K_B^-(m)$



Network Security

Digital signatures

- ❖ suppose Alice receives msg m, with signature: $m, K_B^-(m)$
- ❖ Alice verifies m signed by Bob by applying Bob's public key K_B^+ to $K_B^-(m)$ then checks $K_B^+(K_B^-(m)) = m$.
- ❖ If $K_B^+(K_B^-(m)) = m$, whoever signed m must have used Bob's private key.

Alice thus verifies that:

- ✓ Bob signed m
- ✓ no one else signed m
- ✓ Bob signed m and not m'

non-repudiation:

- ✓ Alice can take m, and signature $K_B^-(m)$ to court and prove that Bob signed m

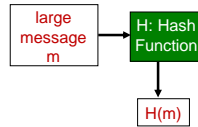
Network Security

Message digests

computationally expensive to public-key-encrypt long messages

goal: fixed-length, easy- to-compute digital "fingerprint"

❖ apply hash function H to m , get fixed size message digest, $H(m)$.



Hash function properties:

- ❖ many-to-1
- ❖ produces fixed-size msg digest (fingerprint)
- ❖ given message digest x , computationally infeasible to find m such that $x = H(m)$

Network Security

Internet checksum: poor crypto hash function

Internet checksum has some properties of hash function:

- ✓ produces fixed length digest (16-bit sum) of message
- ✓ is many-to-one

But given message with given hash value, it is easy to find another message with same hash value:


message	ASCII format	message	ASCII format
I O U 1	49 4F 55 31	I O U 9	49 4F 55 39
0 0 . 9	30 30 2E 39	0 0 . 1	30 30 2E 31
9 B O B	39 42 D2 42	9 B O B	39 42 D2 42
	B2 C1 D2 AC		B2 C1 D2 AC

different messages but identical checksums!

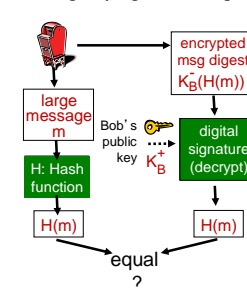
Network Security

Digital signature = signed message digest

Bob sends digitally signed message:



Alice verifies signature, integrity of digitally signed message:



Network Security

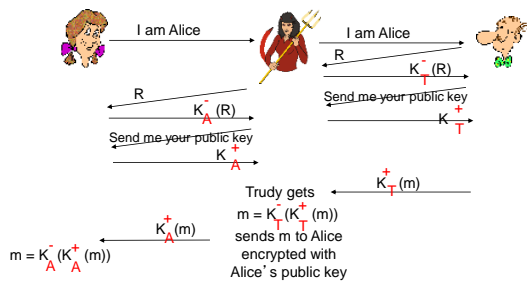
Hash function algorithms

- ❖ MD5 hash function widely used (RFC 1321)
 - computes 128-bit message digest in 4-step process.
 - arbitrary 128-bit string x , appears difficult to construct msg m whose MD5 hash is equal to x
- ❖ SHA-1 is also used
 - US standard [NIST, FIPS PUB 180-1]
 - 160-bit message digest

Network Security

Recall: ap5.0 security hole

man (or woman) in the middle attack: Trudy poses as Alice (to Bob) and as Bob (to Alice)



Network Security


Public-key certification

- ❖ motivation: Trudy plays pizza prank on Bob
 - Trudy creates e-mail order: Dear Pizza Store, Please deliver to me four pepperoni pizzas. Thank you, Bob
 - Trudy signs order with her private key
 - Trudy sends order to Pizza Store
 - Pizza Store verifies signature; then delivers four pepperoni pizzas to Bob
 - Bob doesn't even like pepperoni

Network Security

Certification authorities


- ❖ **certification authority (CA)**: binds public key to particular entity, E.
- ❖ E (person, router) registers its public key with CA.
 - E provides "proof of identity" to CA.
 - CA creates certificate binding E to its public key.
 - certificate containing E's public key digitally signed by CA – CA says "this is E's public key"



Network Security

Certification authorities

- ❖ when Alice wants Bob's public key:
 - gets Bob's certificate (Bob or elsewhere).
 - apply CA's public key to Bob's certificate, get Bob's public key



Network Security