Networking Overview

Today

- Speed run of 438
- Focus on topics that will be covered in papers

- Reminder: Project proposals due in 1 week

What is the Internet?

- To the layperson: useful services
  - Web, email, video, voice

- Technically: global system that lets *hosts* communicate
  - Physical infrastructure
    - switches, routers, links, radios
  - Protocols
    - WiFi, Ethernet, IP, TCP, HTTP

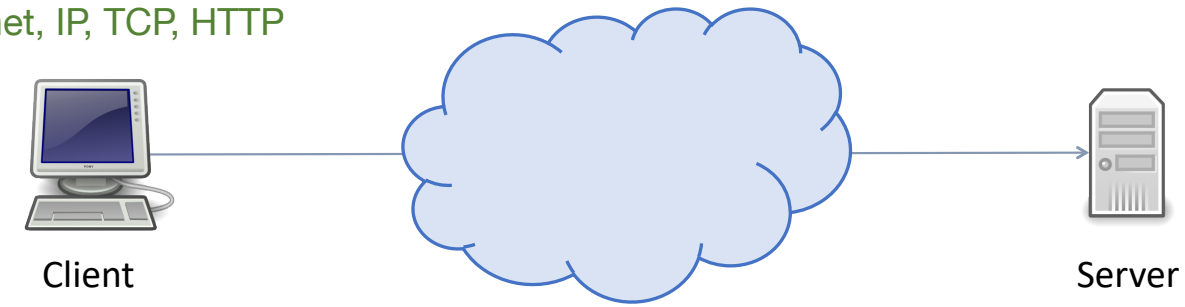

Packet Switching

- Internet provides best-effort delivery of *packets* between hosts
- **Packet:** a structured sequence of bytes
  - Header: metadata used by network
  - Payload: data to be transported
- Packets are *forwarded* by *routers* from sender to destination host
  - Each packet is treated independently

Routers

- Receive outgoing packets from local hosts and attempt to deliver them to destination
- Deliver incoming packets to local hosts

Internet Message Processor

Linksys WRT54G

Cisco CRS-1

Protocol Layering

- A network isn't defined by one protocol, but an entire ecosystem of protocols!
- Networks use a stack of layers
- Lower layers provide services to layers above
  - Don't care what higher layers do
- Higher layers use services of layers below
  - Don't care how lower layers implement services
- Layers define abstraction boundaries
  - At a given layer, all layers above and below are largely opaque

The Internet Protocol Stack

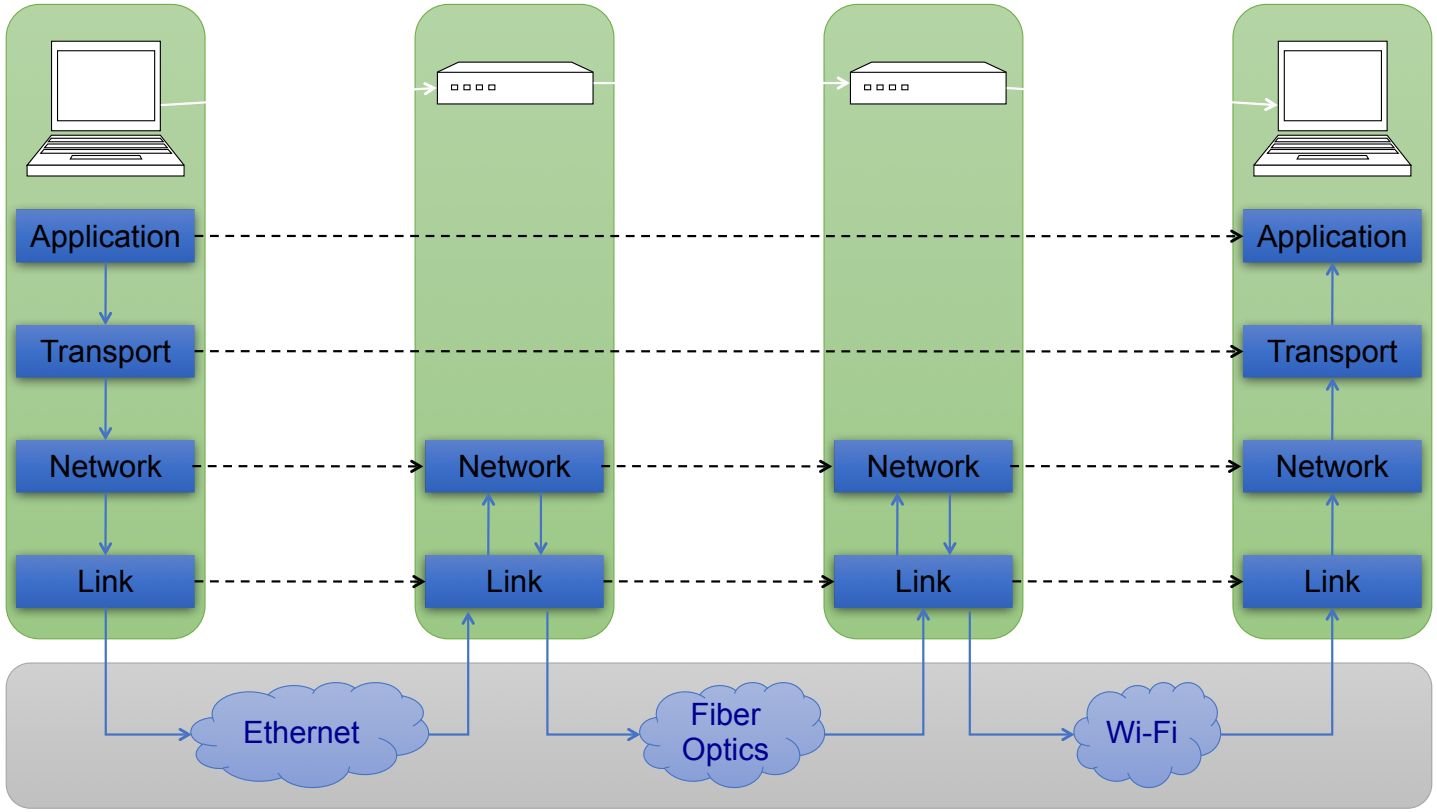

The Internet Protocol Stack
(a Bottom-Up Approach…)

- Physical Layer: Transmits the raw bits of a packet over a physical data link.
- Link Layer: Transmits packet from one host to another host that it is physically connected to.
- Network Layer: Transmits packet from a host in one network to a host in another network (i.e., **internet**working)
- Transport Layer: Transmits packets over a *stateful connection* between two hosts
- Application Layer: Transmits packets from one process to another process

Layering of protocols

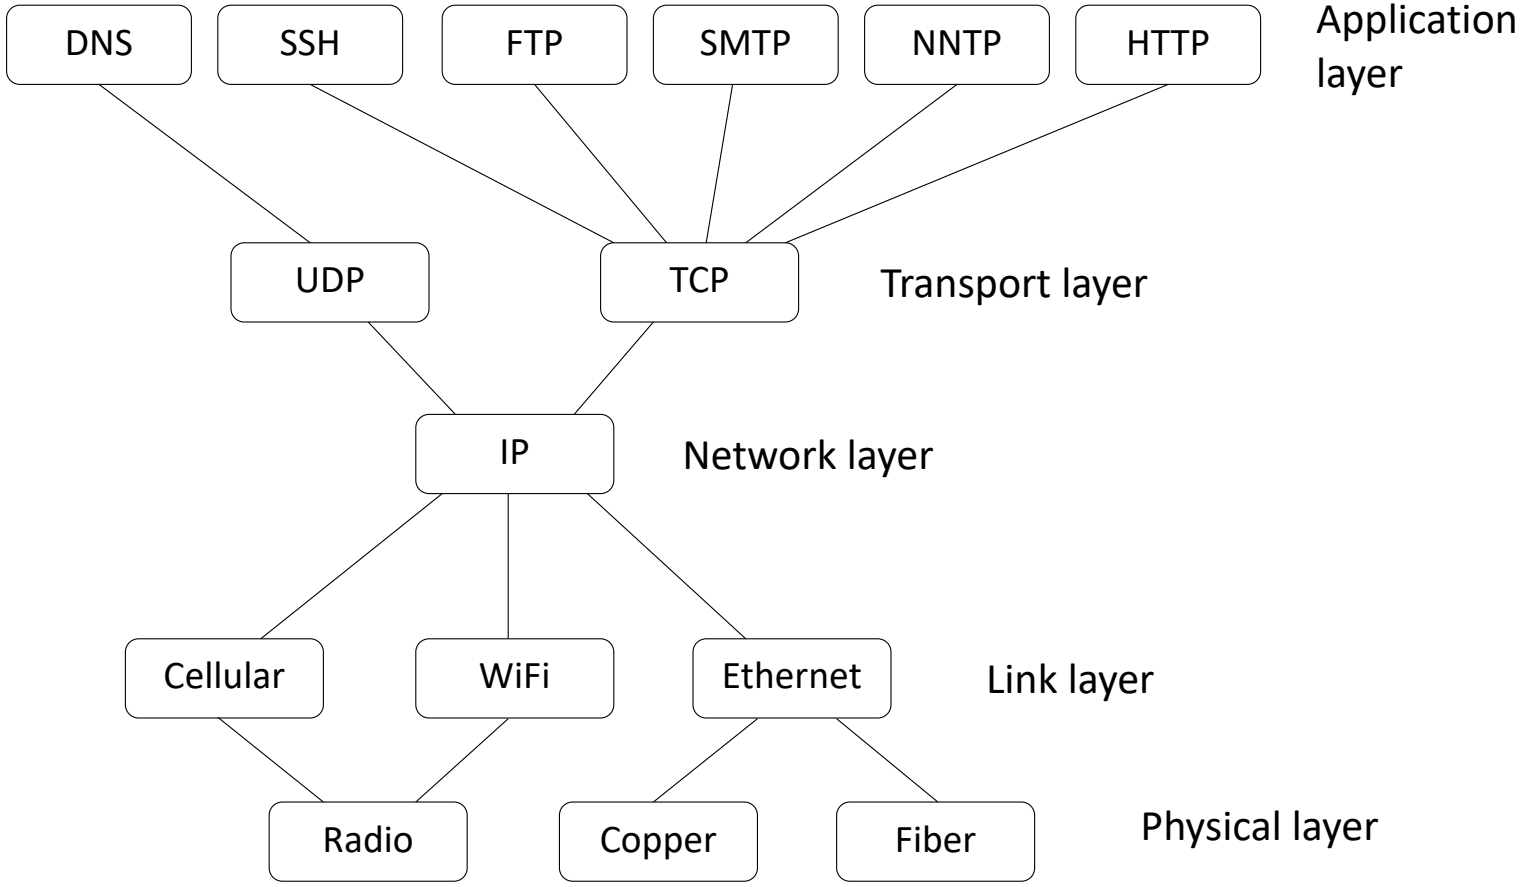

Internet Packet Encapsulation

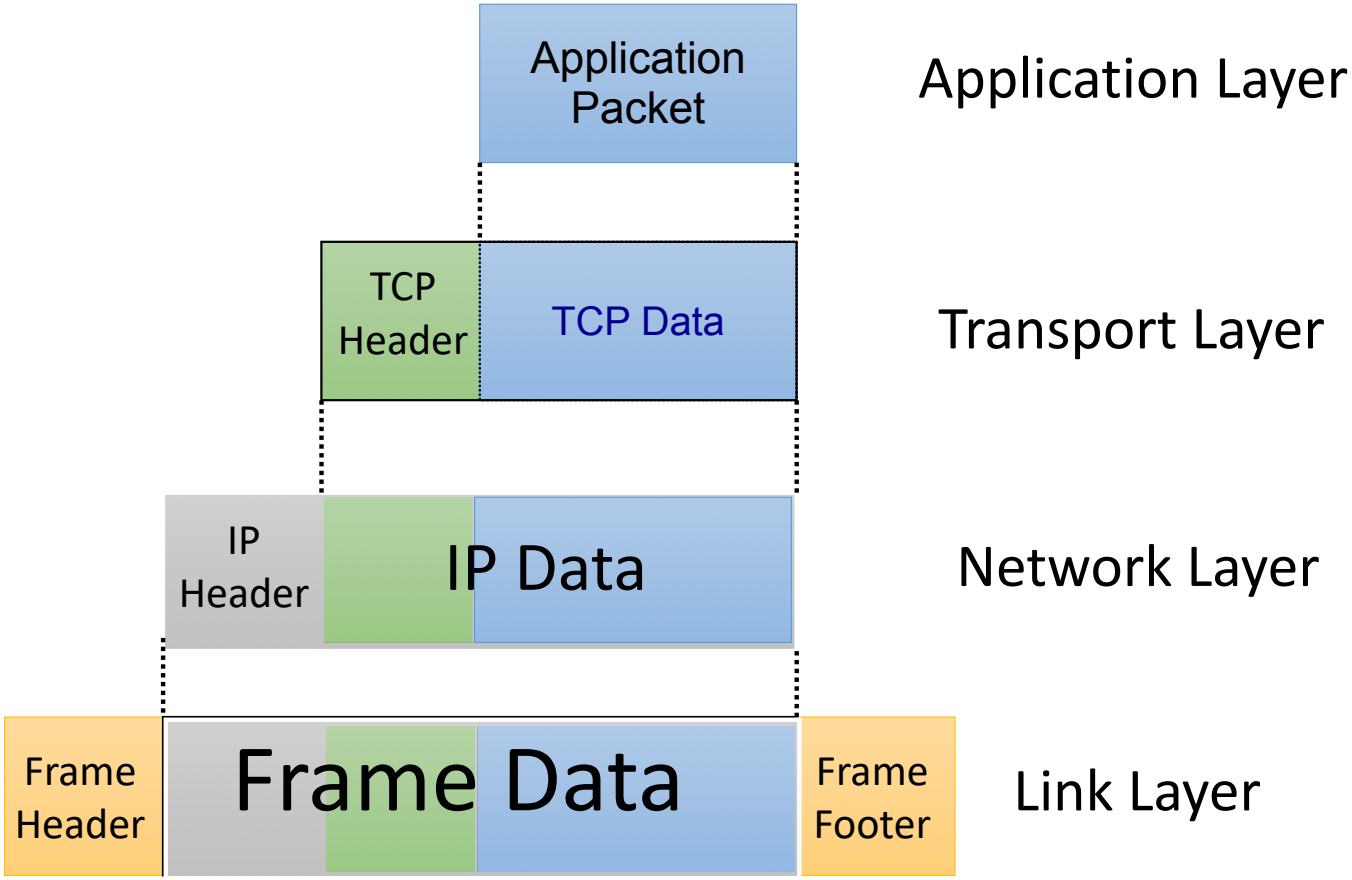

Addressing

- Link layer: MAC address: 48 bits, e.g., 2c:54:91:88:c9:e3
  - Meaningful only on *local network*
  - Usually fixed per device (though MAC address privacy more common)
- IP layer: IP address, 32-bits (v4) or 48-bits (v6)
  - 1.2.3.4 (v4) or 2001:db8:0:0:0:800:200c:7334 (v6)
  - Short forms: 2001:db8**::**800:200c:7334 and 2001:db8**:**

Addressing

- Link layer: MAC address: 48 bits, e.g., 2c:54:91:88:c9:e3
  - Meaningful only on *local network*
  - Usually fixed per device (though MAC address privacy more common)
- IP layer: IP address, 32-bits (v4) or 48-bits (v6)
  - 1.2.3.4 (v4) or 2001:db8:0:0:0:800:200c:7334 (v6)
  - Short forms: 2001:db8**:***[0:0:0]***:**800:200c:7334 and 2001:db8**:***[0:0:0:0:0:0]*

Layering of protocols

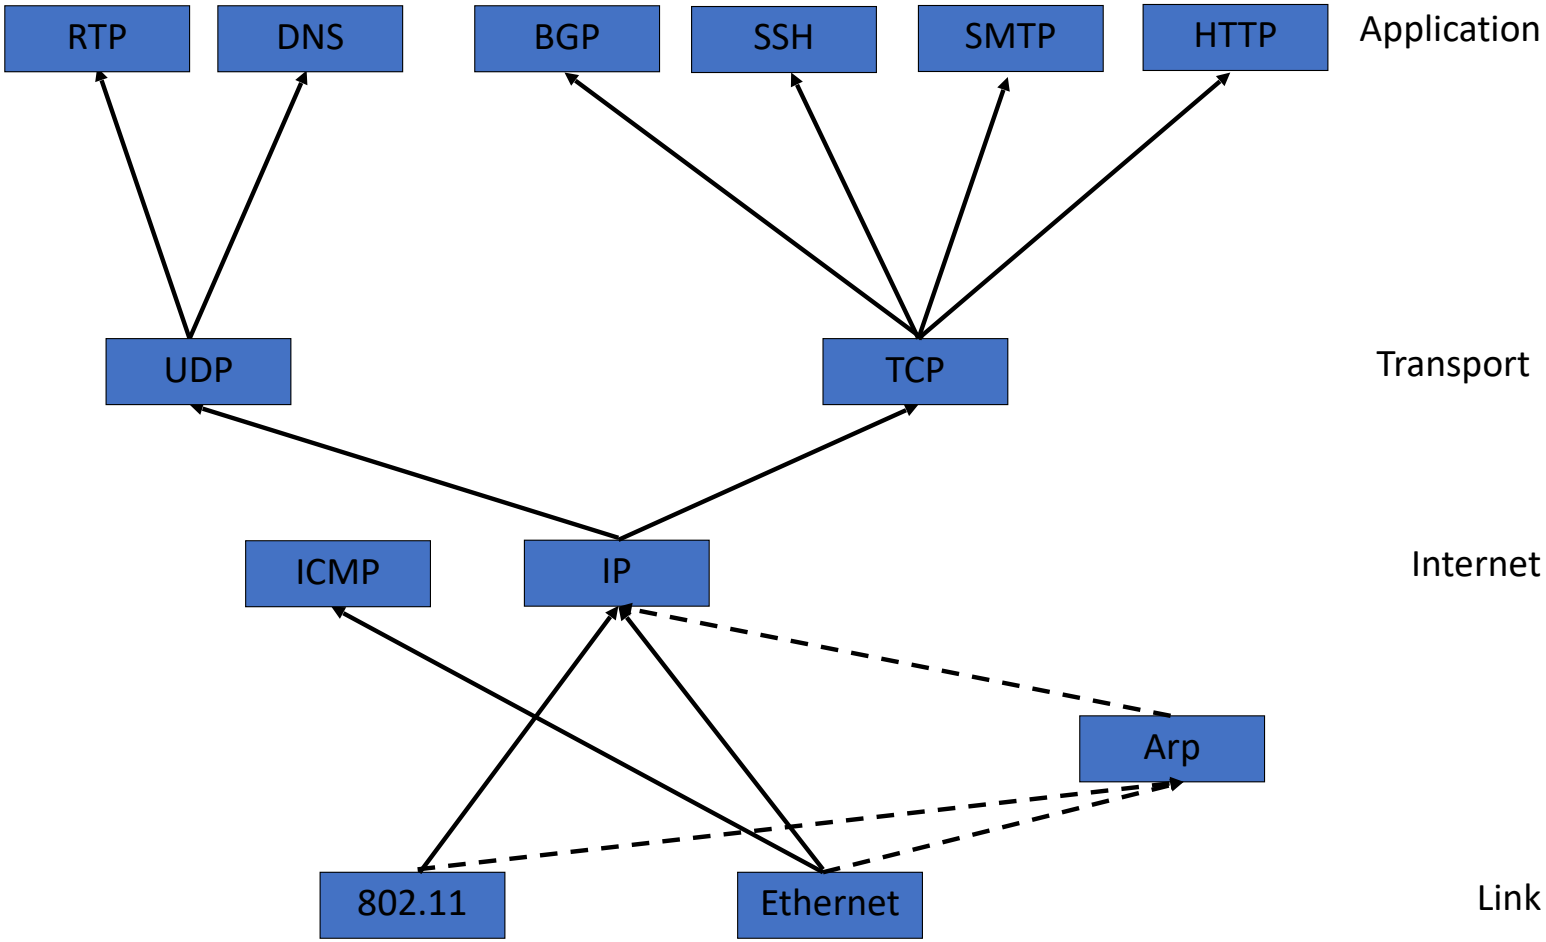

Layering of protocols

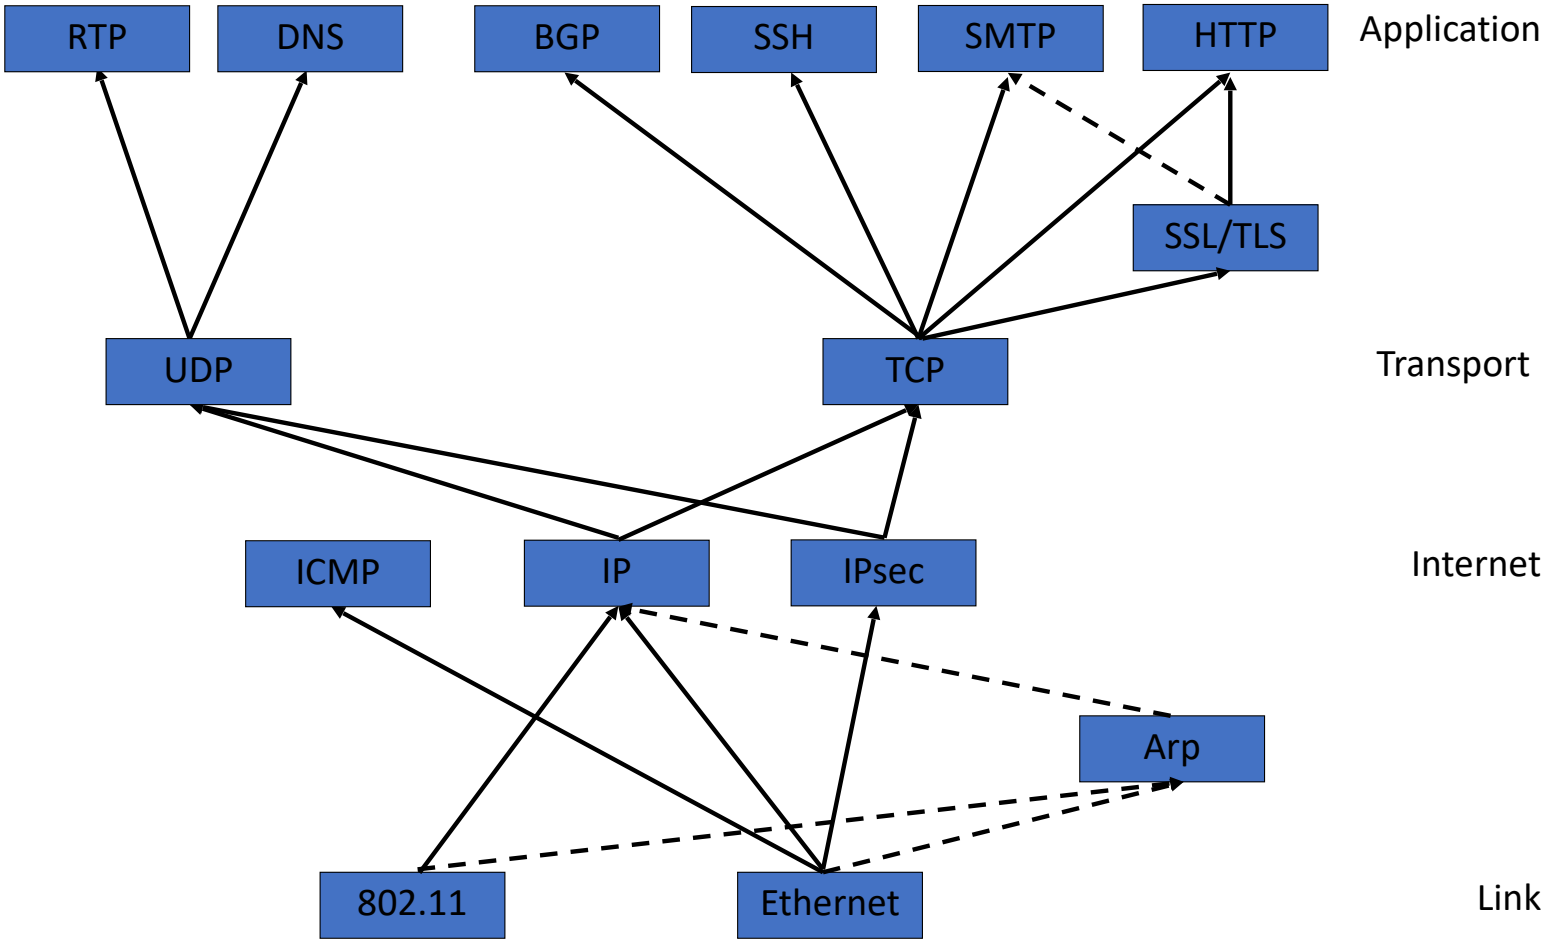

Cryptography Toolbox

- Cryptographic hash functions
- Message Authentication Codes
- Symmetric encryption and decryption
- Asymmetric encryption and decryption
- Digital signatures
- Key exchange

Symmetric

Asymmetric

Basic Idea

- Based on *key agreement*
  - Establish a shared secret key using key agreement (e.g., Diffie-Hellman)
  - Authenticate shared secret (digital signatures scheme)
- Use shared secret for symmetric cryptography
  - Symmetric encryption and decryption
  - MACs

- Based on *key exchange*
  - Generate random session key, encrypt using peer's public key, and send to peer
  - Authenticate encrypted key (digital signature scheme)
- Use shared secret for symmetric cryptography
  - Symmetric encryption and decryption
  - MACs

Client Hello includes a nonce *(random number generated by client)*

Illustrated TLS connection with explanations:
`https://tls.ulfheim.net/`

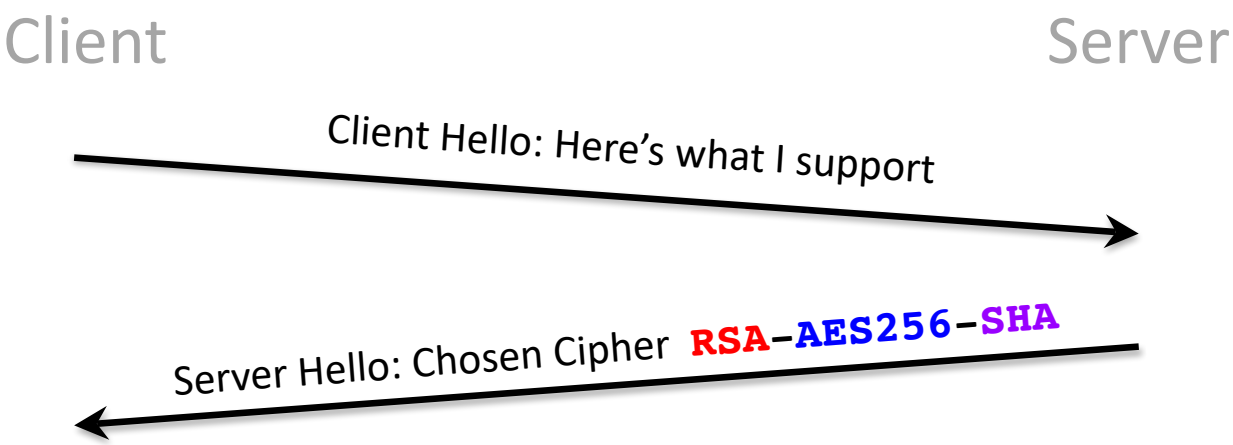

RSA-AES256-SHA

Key Exchange

Data Transfer Cipher

Message Digest / Authentication Code

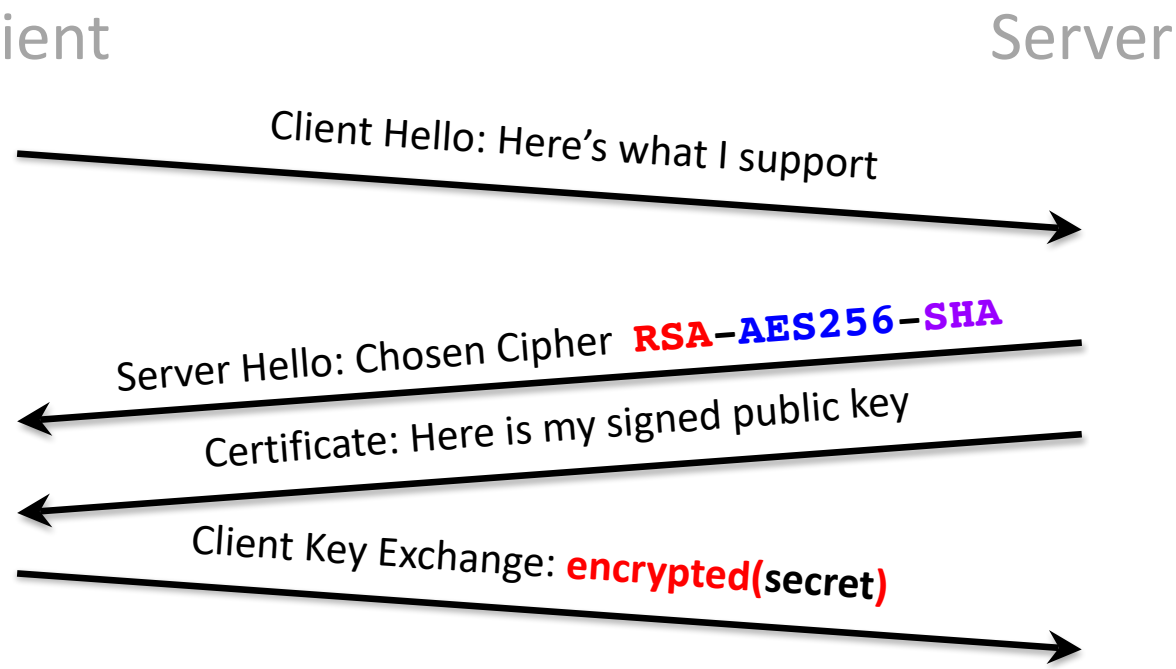

Encrypted using Server's public key *(the public key from the Certificate)*
This means that *only the server* can decrypt the secret!

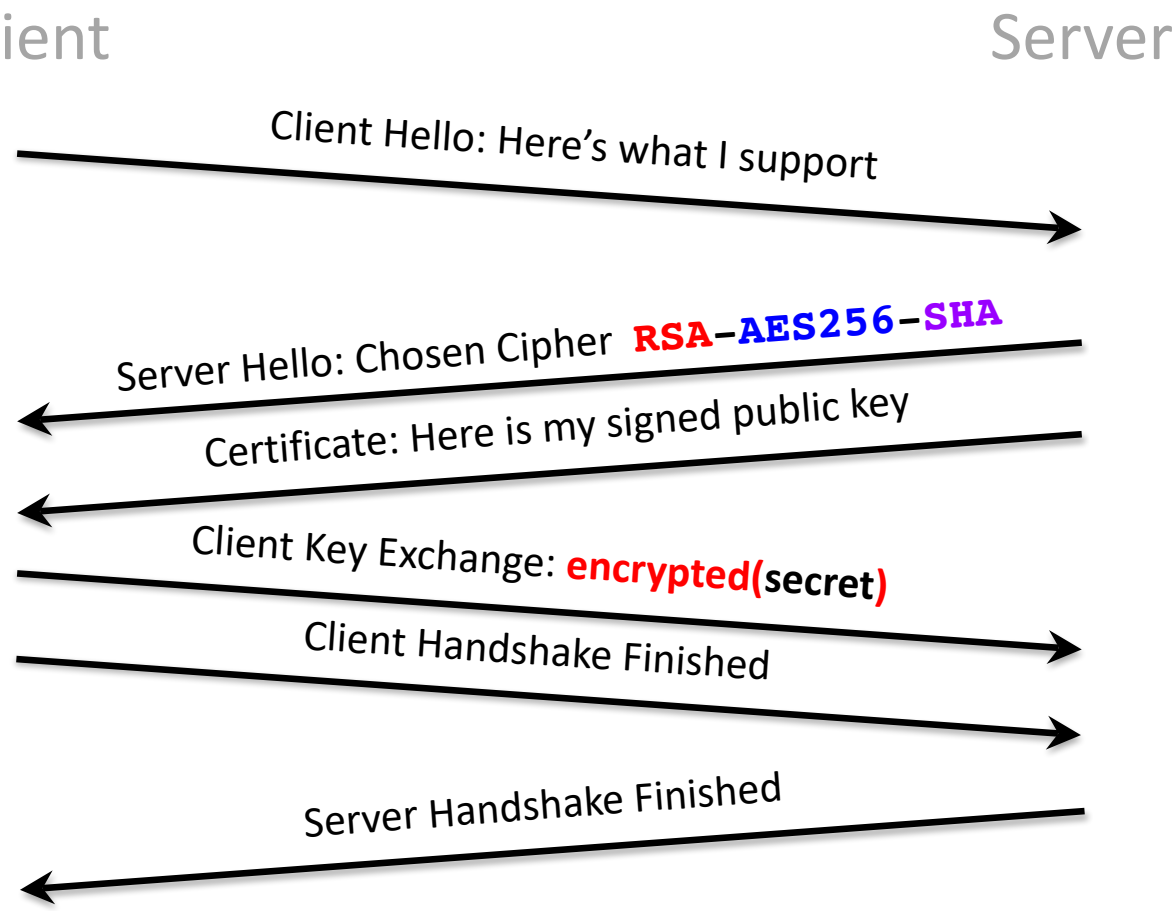

The Handshake Finished message contains a MAC of the handshake so far

Why?

Session Keys

- Shared secret used to derive session keys
  - Symmetric key for data encryption and decryption
  - Symmetric key for MACs
- Session keys must be hard to guess
  - Pick shared secret uniformly at random
- All communication after handshake is encrypted and MAC'd using negotiated suite

Session Resumption

- Key exchange is expensive
  - Several round trips
  - Asymmetric cryptography
- Session resumption
  - Server creates a **ticket** that encapsulates session state (keys, auth)
  - Client uses ticket to **resume** a session

DNS

- People refer to names as adv-sec-sp25.nikita.phd or facebook.com
- Need to translate this to IP addresses => Domain Name System
- DNS includes:
  - Registrars
  - Authoritative servers
  - Caching resolvers
  - Clients

Domain Name System

- Application-layer protocols (and people) usually refer to Internet host by *host name*
- Host names organized into hierarchy

www.illinois.edu

Domain Name System

- Application-layer protocols (and people) usually refer to Internet host by *host name*
- Host names organized into hierarchy

www.illinois.edu

top-level
domain (TLD)

Domain Name System

- Application-layer protocols (and people) usually refer to Internet host by *host name*
- Host names organized into hierarchy

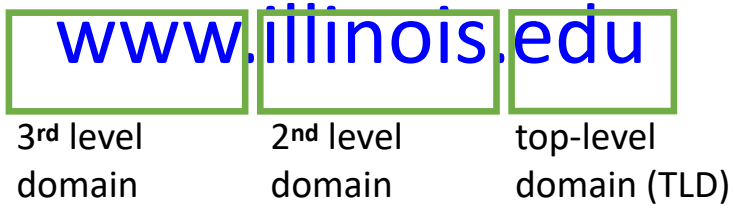

DNS Hierarchy

- Each level allocates names to next level
- TLDs allocated by ICANN
  - ccTLD: country-based TLDs, two letters (e.g. `.us`)
  - gTLD: arbitrary names, 3+ letters (e.g. `.com`)
- TLD operated by different registries
- Registrars are agents that register domains for a person or organization in a particular TLD
- Organizations have control over its subdomains
  - e.g. UIUC decides what domains have suffix illinois.edu

$ dig adv-sec-sp25.nikita.phd

```
; <<>> DiG 9.10.6 <<>> adv-sec-sp25.nikita.phd
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 18393
;; flags: qr rd ra ad; QUERY: 1, ANSWER: 2, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1232
;; QUESTION SECTION:
;adv-sec-sp25.nikita.phd.       IN      A

;; ANSWER SECTION:
adv-sec-sp25.nikita.phd. 300    IN      A       104.21.11.141
adv-sec-sp25.nikita.phd. 300    IN      A       172.67.166.39
```

DNS Server Roles

- **Authoritative server:** provides authoritative information for a set of domains
  - Does not handle queries about other domains
- **Recursive resolver:** provides recursive resolution of a domain to return requested record to client
  - Handles queries about all domains
- Same protocol for both types of servers
  - Distinction is in intended purpose only

DNS Name Resolution

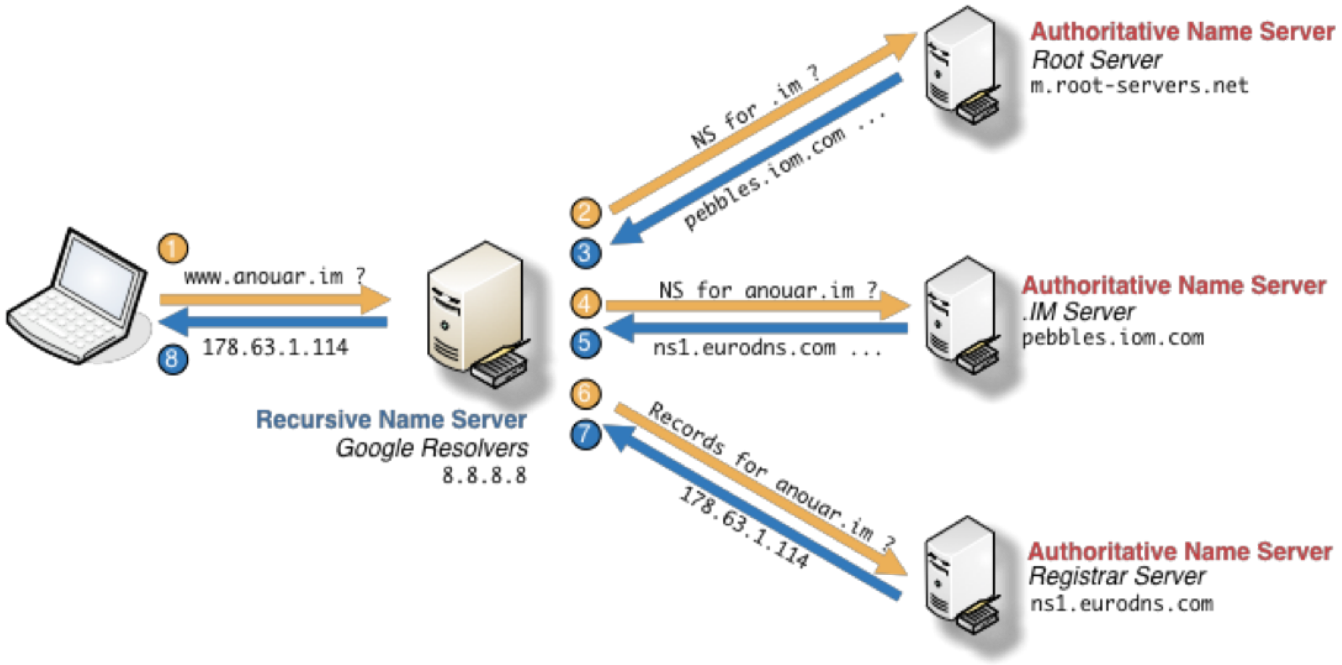

source: http://anouar.adlani.com/2011/12/useful-dig-command-to-troubleshot-your-domains.html

Caching and Additional Records

- Recursive resolvers (and clients) cache DNS entries
  - TTL is used to indicate how long caching is allowed
- DNS responses often include additional records that help resolution
  - Fertile ground for attack!

```
; <<>> DiG 9.10.6 <<>> @a.gtld-servers.net courses.grainger.illinois.edu
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 49286
;; flags: qr rd; QUERY: 1, ANSWER: 0, AUTHORITY: 3, ADDITIONAL: 7

;; QUESTION SECTION:
;courses.grainger.illinois.edu. IN      A

;; AUTHORITY SECTION:
illinois.edu.           172800  IN      NS      dns1.illinois.edu.
illinois.edu.           172800  IN      NS      dns2.illinois.edu.
illinois.edu.           172800  IN      NS      dns3.illinois.edu.

;; ADDITIONAL SECTION:
dns1.illinois.edu.      172800  IN      A       130.126.2.100
dns1.illinois.edu.      172800  IN      AAAA    2620:0:e00:b::53
dns2.illinois.edu.      172800  IN      A       130.126.2.120
dns2.illinois.edu.      172800  IN      AAAA    2620:0:e00:c::53
dns3.illinois.edu.      172800  IN      AAAA    2600:1f16:8b2:2e53::53
```

```
; <<>> DiG 9.10.6 <<>> @a.gtld-servers.net courses.grainger.illinois.edu
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 49286
;; flags: qr rd; QUERY: 1, ANSWER: 0, AUTHORITY: 3, ADDITIONAL: 7

;; QUESTION SECTION:
;courses.grainger.illinois.edu. IN      A

;; AUTHORITY SECTION:
illinois.edu.           172800  IN      NS      dns1.illinois.edu.
illinois.edu.           172800  IN      NS      dns2.illinois.edu.
illinois.edu.           172800  IN      NS      dns3.illinois.edu.

;; ADDITIONAL SECTION:
dns1.illinois.edu.      172800  IN      A       130.126.2.100
dns1.illinois.edu.      172800  IN      AAAA    2620:0:e00:b::53
dns2.illinois.edu.      172800  IN      A       130.126.2.120
dns2.illinois.edu.      172800  IN      AAAA    2620:0:e00:c::53
dns3.illinois.edu.      172800  IN      AAAA    2600:1f16:8b2:2e53::53
paypal.com.             172800  IN      A       130.126.2.103
```

Bailiwick Rules

- **Bailiwick rule:** defines what response records a recursive resolver will accept
- **Bailiwick** (general def.): the area of authority of a legal officer, e.g., a set of territories
  - Synonym – Jurisdiction!
- **Bailiwick** (DNS def.)**:** set of domains about which a server is has direct or indirect authority to speak
  - Bailiwick determined by the initiator of query
- Answer should be relevant *(in response to request)*
- Answer should be in bailiwick

Bailiwick Checking Rule from BIND

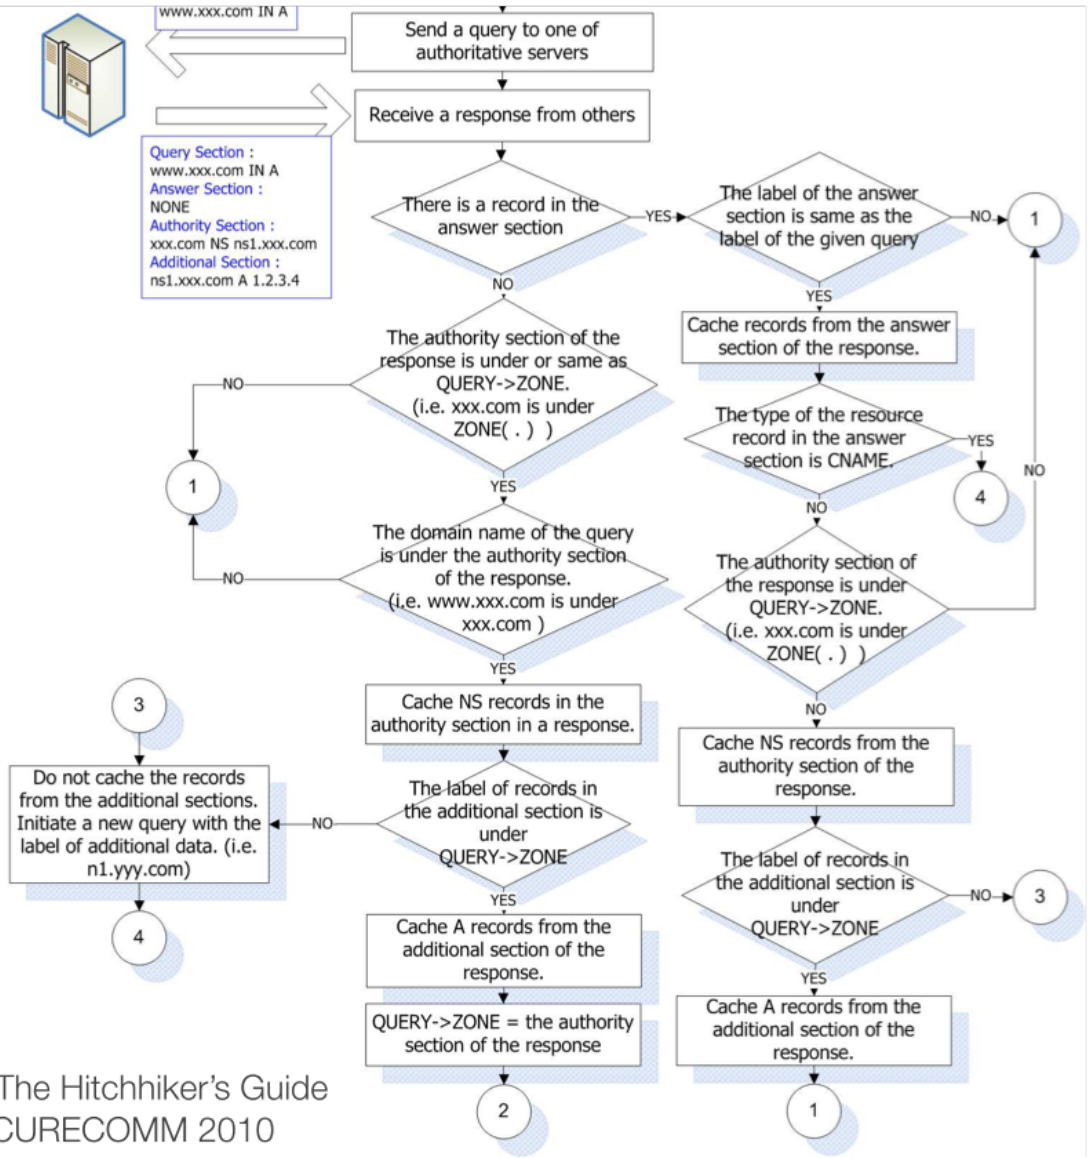

source: Son and Shmatikov, “The Hitchhiker’s Guide to DNS Cache Poisoning” SECURECOMM 2010

BIND Bailiwick Rule (roughly)

- Authorities must be for queried domain
  - `dns1.illinois.edu` accepted as authority for illinois.edu only when initiating query was for subdomain of `illinois.edu`
- Additional records must be *in bailiwick* for query
  - A record for `dns1.illinois.edu` accepted because `edu` server has indirect authority over `dns1.illinois.edu`

Routing

- Routers need to decide how to deliver packets to destination
- Internet divided into *autonomous systems*
  - Intra-AS routing is managed entirely by the AS
  - Inter-AS routing is managed using Border Gateway Protocol
  - BGP goals: reachability, performance, policy, ~~security~~

BGP brief overview

- Autonomous systems advertise prefixes
  - 130.126/16: 130.126.0.0 – 130.126.255.255
  - 130.126.17.128/26: 130.126.17.128 – 130.126.17.192
  - When two prefixes overlap, more specific one is preferred
- BGP used to talk to neighboring ASes:
  - Export routes: promise to deliver to prefix
  - Import routes: use neighbor to deliver to prefix
- Advertisements include AS PATH: list of ASes used to reach prefix

ASPATH Attribute

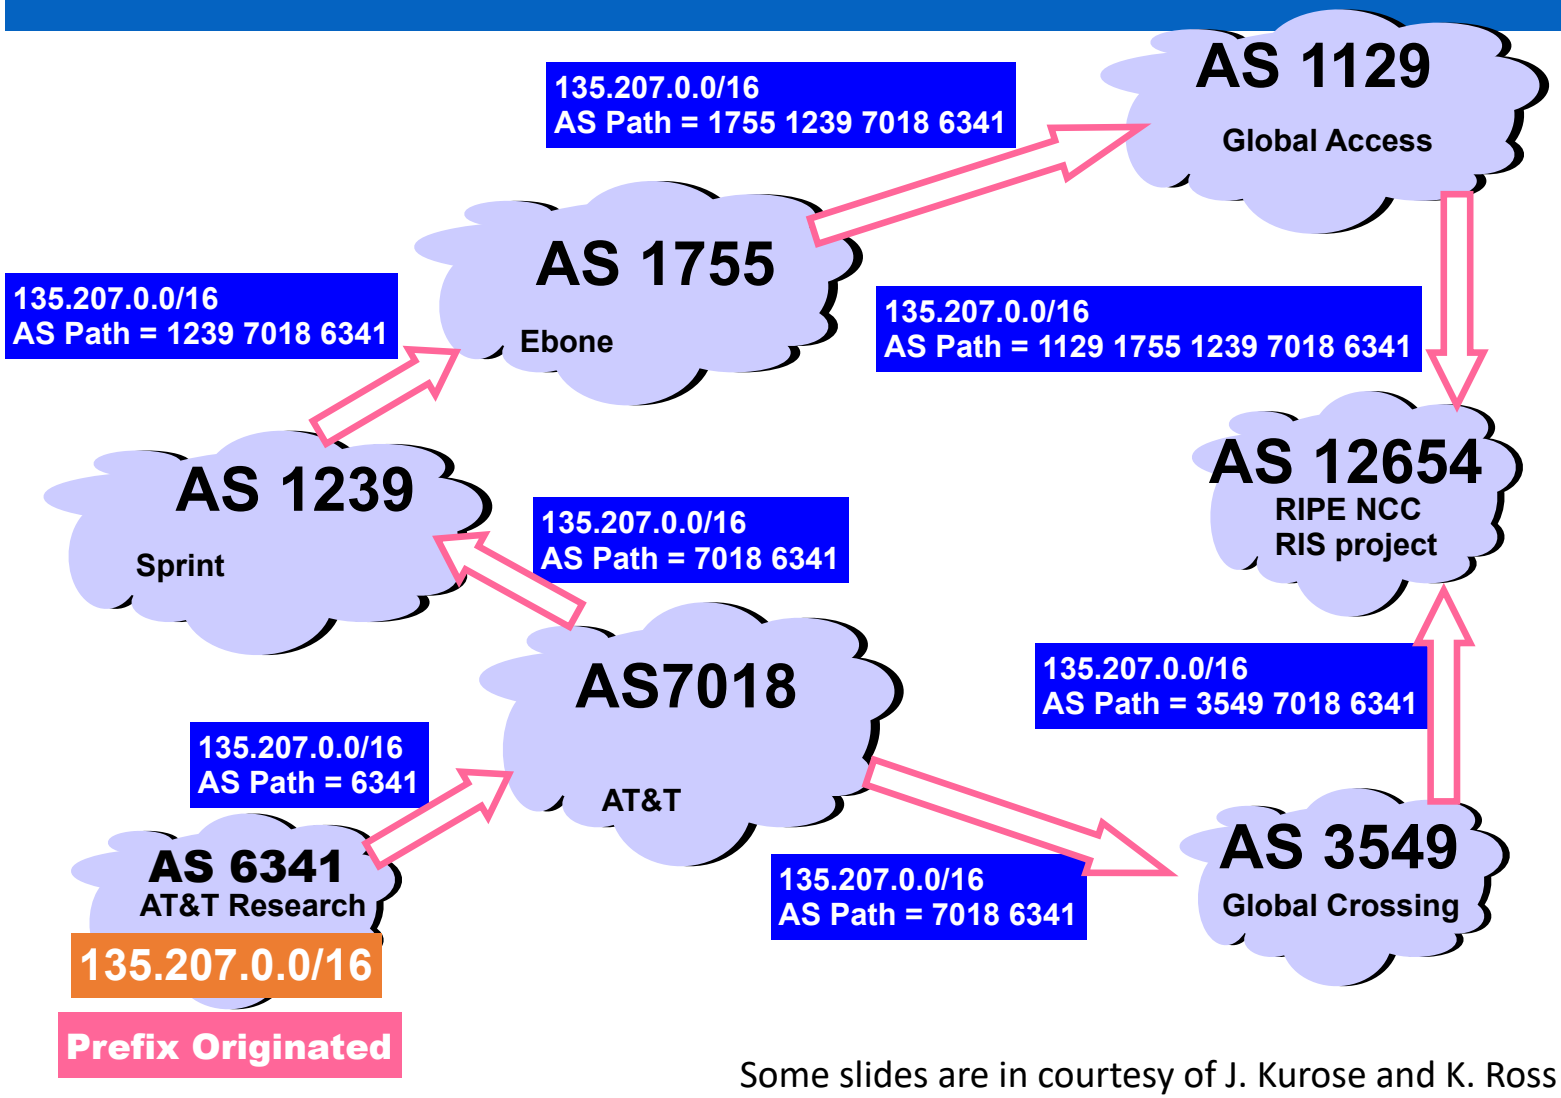

Some slides are in courtesy of J. Kurose and K. Ross

BGP Export and Import Rules

- Export policy: am I willing to carry traffic from neighbor to this destination?
  - Is this going to make me money?
- Import policy: which neighbor will I use to reach destination?
  - What will make me most money?
  - What will be most efficient? (AS path)

BGP Hijacking

BGP attacks hijack Telegram traffic in Iran

With so many users in Iran, it's unsurprising that potentially state-sponsored groups would want an access point into the banned app.

Pakistan's Accidental YouTube Re-Routing Exposes Trust Flaw in Net

A Pakistan ISP that was ordered to censor YouTube accidentally managed to take down the video site arou the world for several hours Sunday. The Pakistani government ordered ISPs to censor YouTube to prevent Pakistanis from seeing a trailer to an anti-Islamic film by Dutch politician Geert Wilders. YouTube has since removed the clip for violating its terms of service, but a screenshot [...]

Popular Destinations rerouted to Russia

Posted by Andree Toonk - December 12, 2017 - Hijack - No Comments

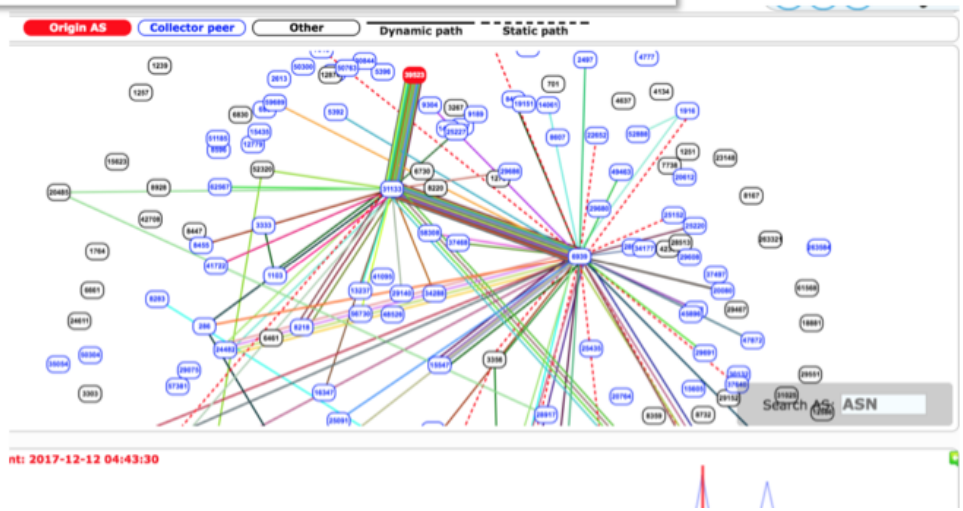

BGP Hijacking

- BGP can be subverted by advertising a route that:
  - Is more specific than existing routes (= preferred)
  - Is more attractive than existing routes
- Both happen regularly due to misconfiguration and attacks
- Defenses
  - Sanity checks of imported routes (e.g., no /32, no /0)
  - Monitoring of large-scale changes
  - *Route origin verification*

Route Origin Verification

- Route Origin Authorization: Prove that AS is allowed to originate prefix
  - Prefix authorities run RPKI to cryptographically authorize prefix origination
- Ideally: route only imported if ROA check passes
  - In practice, deployment spotty
- ROV does not prevent someone from advertising a shorter path than exists
  - Full path verification much more complex