CSCI 1800 Cybersecurity and International Relations

Internet Naming and Routing

John E. Savage

Brown University

Outline

- The Domain Name System (DNS)
  – Protecting the DNS from attacks
- History of Naming Policy
- Internet routing
  – The Border Gateway Protocol (BGP)
  – Protecting BGP from attacks
- Routing Policy

The Domain Name System

The Domain Name System (DNS)

- DNS is the "telephone directory" for the Internet.
- DNS is a distributed, hierarchical, naming system.
- DNS translates host names into IP addresses.
  - www.example.com translates to the addresses *192.0.32.10* (IPv4) and *2620:0:2d0:200::10* (IPv6).
- Names are hierarchical
  - .com is a top-level domain
  - example.com is a second-level domain of .com
  - aaa.example.com is sub-domain of example.com

Domain Names

- Four types of top-level domain (TLD):
  - Country codes (2 letters, e.g. .ca, .au, .de, .hu, .uk)
  - Sponsored codes (e.g. .coop, .jobs, .post, .gov, .mil, .int)
  - Historical top level (e.g. .com, .net, .edu, .org,)
    - ~1,540 active TLDs, e.g. .IBM, .NYC, .REISE, COOKINGCHANNEL
- Domain names are registered and assigned by domain-name registrars† who are accredited by the Internet Corporation for Assigned Names and Numbers (ICANN).

† See http://www.icann.org/registrar-reports/accredited-list.html

Organization of the DNS

- The DNS resolves names into IP addresses.
- Root name servers hold IP addresses for top-level name servers, e.g. .edu, .uk. and .net.
- Top-level name servers hold IP addresses for sub-domain name servers, e.g. example.com.

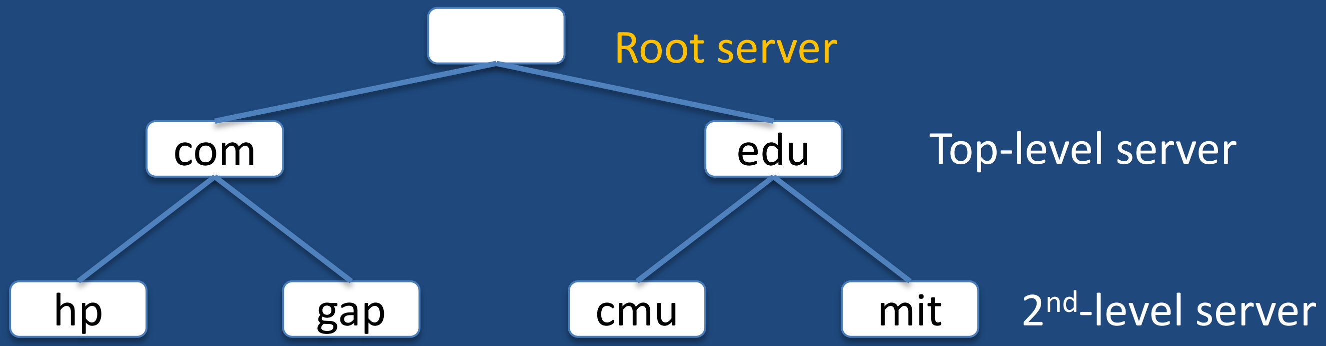

Querying the DNS

- Local caches hold records mapping domain names to IP addresses. If the time to live (TTL) for a domain expires, another lookup is done. TTL about 2 hours
- When local cache is queried for a name that is not in the cache, it is fetched via root server and cache is updated with new mapping.
- Root server is asked for IP address of name server for top-level domain, which is asked for IP address of second-level domain server, etc., until authoritative server is reached, which returns correct IP address.

DNS Cache Poisoning

- Eve tricks DNS cache into mapping a domain name to fake IP addr
  - Users will go to fake IP address until TTL reached

- Steps Eve takes to poison the cache:
  1. Eve sends a request IP address for DNS name not in cache
  2. Cache asks authoritative server **S** for mapping, sending to it a 16-bit ID. The server responds with same ID after delay
  3. Eve guesses 16-bit ID but responds to cache before **S** does with incorrect answer.
  4. If Eve guesses ID correctly, DNS accepts her answer and ignores later input from authoritative server S.
  5. Cache is poisoned with fake IP address for the domain name.

Protecting DNS Caches

- Problems in protecting DNS caches:
  – 16-bit IDs on DNS queries are short, too easily guessed
  – It only takes 64K* tries to find correct ID
- How to harden DNS caches:
  – Only allow updates from within local network.
    - If update is from outside local network, don't trust it.
  – Provide port number when querying root zone and require that responses have correct port no. and ID.
  – Number of choices goes from 2^{16} to 2^{32}!

* K = 2^{10} = 1,024

Public Key Cryptography

- Alice and Bob have public and private keys PrvA, PubA and PrvB, PubB
- Bob encrypts a message for Alice using her public key PubA. She decrypts it using her private key PrA.
- Alice sends messages to Bob the same way.
- Using this method, they can communicate in secret.

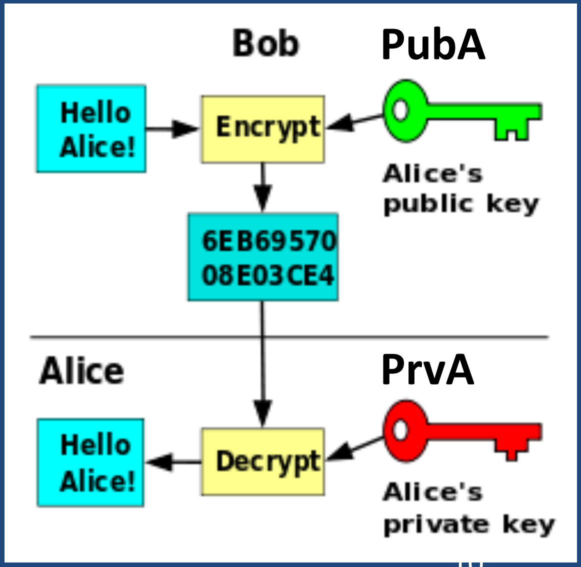

Cryptographic Signing of Messages

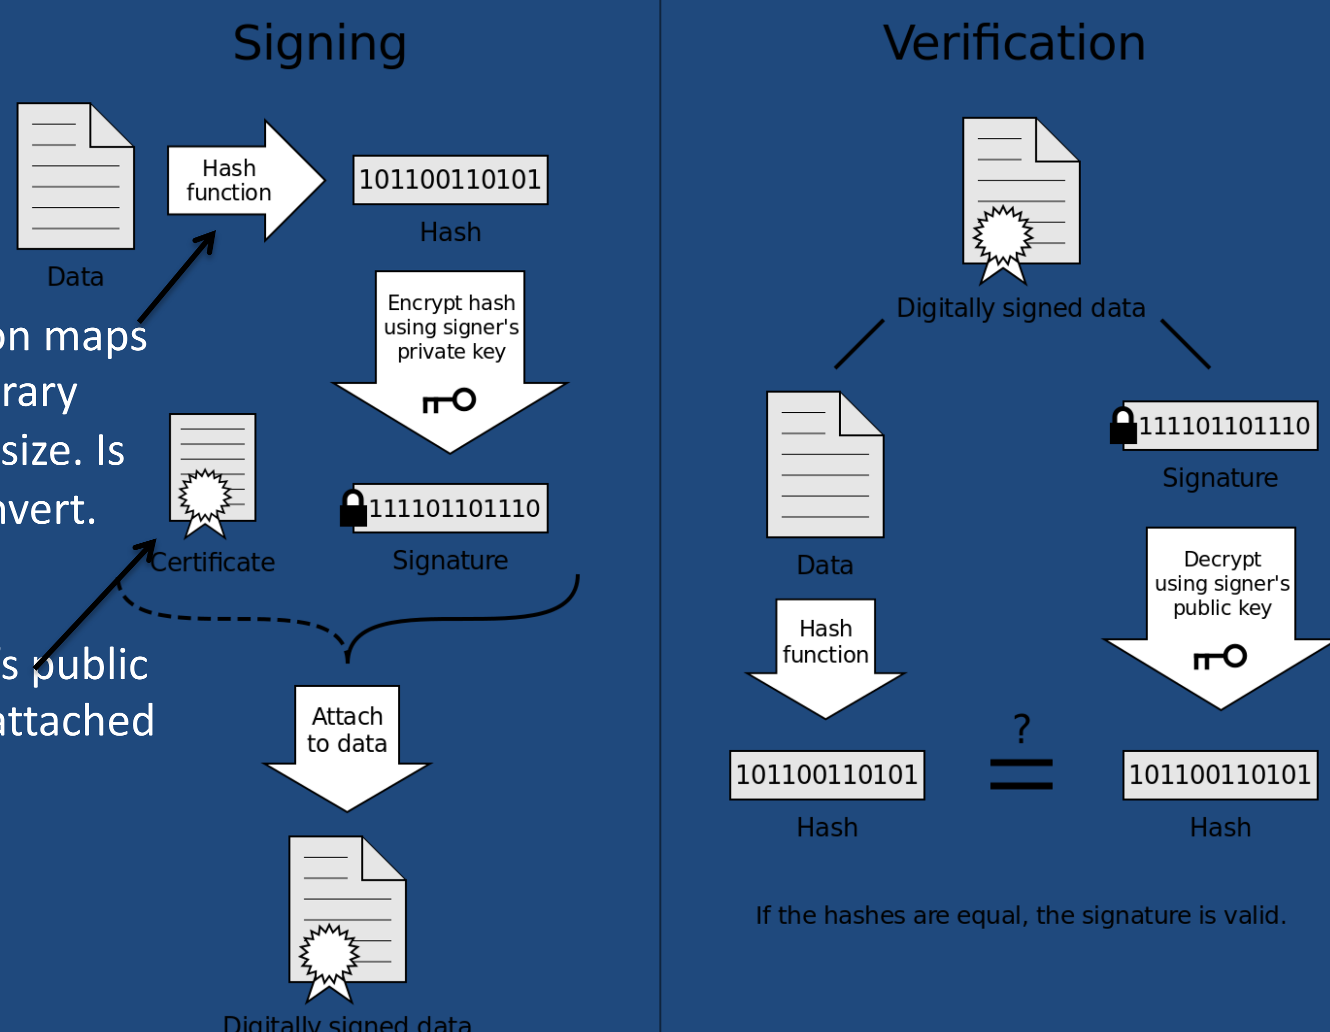

Hash function maps data of arbitrary size to fixed size. Is difficult to invert.

Signer’s public key is attached

DNSSEC: Security Extensions to DNS

- DNS is not secure! DNSSEC provides trust
- Under DNSSEC, DNS replies are cryptographically signed using public key encryption.
  - A message identifying sender is encrypted by sender.
  - Public decryption key is used to verify author.
- Source has authority granted by issuer of keys
- Chain of trust here. Ultimately, must trust root.
- Most TLDs are protected by DNSSEC

History of Naming Policy

Names Matter

- Domain names can be expensive,
  – insurance.com cost $35.6 million in 2010
  – cars.com cost $872 M in 2014
  – Suffixes such as .xxx , .sucks may be controversial.
- Who should have the authority to decide on ownership and assignment of domain names and IP addresses?

Early Days

- In early 1970s naming system consisted of small file called “hosts.txt” placed at each host.
- In 1978 Jon Postel of USC was given no-bid USG contract to run Internet naming & numbering
- By mid 1980s Postel and SRI had created the modern domain name system.
- By 1990s DoD required contract bidding.

Commercialization of Internet

- In May 1990 Government Systems, Inc. wins contract to administer the root (Postel's job) which it hands over to Network Solutions.
- In 1995 Network Solutions wins right to charge for registering domain names.
- Domain names become very popular and Network Solutions earns fabulous profits.
- Engineers disenchanted.

First Attempt at Capturing the Root

- In June 1991 Vint Cerf and others announce formation of Internet Society (ISOC).
  - Goal: Provide Internet governing structure, home, and funding that is independent of USG
  - Milt Mueller: An attempt to self-privatize the Internet.
- In March 1995 Aiken of US Energy Department asks ISOC what authority ISOC is claiming.
- Vint Cerf responds implying that it is preferable that Internet be run by ISOC, not USG

Role of ISOC

- ISOC writes "Generic Top-Level Domain Memorandum of Understanding" (gTLD-MoU), which looks like international legal document, designed to give Internet policy to ISOC.
- International Telecommunications Union agreed to recognize it and be repository for gTLD-MoU.
  - Formal signing ceremony on May 1, 1997
  - Group of ISPs release tentative Internet Constitution

United States Reacts

- Ira Magaziner (‘69), USG Internet policy czar, responds
  – Commercialization of Internet will be boon to US
  – To foster growth, Internet must not be regulated
  – It must be predictable and secure
  – Only the US has ultimate authority over Internet’s deep structure including naming and routing
  – USG needed to ensure Internet growth and independence
- Issue comes to head with ISOC at 12/1997 DC meeting at which Magaziner states USG case forcefully.
- 1/28/1998 Postel protests by seizing control of root but relents when Magaziner issues legal threat to USC.

ICANN Created in 1988

- Internet Corporation for Assigned Names and Numbers (ICANN), non-profit organization, is created in 1998 to oversee Internet-related tasks
  - ICANN coordinates
    - Domain name system (DNS)
    - IP addresses, allocation of addresses to Internet registrars*
    - Management of root servers and top-level domains
    - Numbers assigned to protocols and autonomous systems
  - Ensures Internet stability and security
  - Consults broadly with users, technologists, govs.

* See http://www.icann.org/registrar-reports/accredited-list.html

Major Internet Governance Event

- On 3/14/14 USG announced "its intent to transition key Internet domain name functions to the global multi-stakeholder community"* if the following goals are met:
  - "Support and enhance the multi-stakeholder model,
  - Maintain the security, stability, and resiliency of Internet DNS,
  - Meet the needs and expectations of the global customers and partners of the IANA services; and
  - Maintain the openness of the Internet."
- **No transition if the role of USG is replaced by another government or an intergovernmental organization.**

* NTIA Press Release, http://www.ntia.doc.gov/press-release/2014/ntia-announces-intent-transition-key-internet-domain-name-functions

2016 US Supervision of ICANN Ends

- After substantial revision of its bylaws, ICANN allowed to operate without USG supervision.
- However, ICANN and its new subsidiary, PTI (an acronym for post-transition IANA), are US corporations subject to US law.
- These changes are in a special set of ICANN bylaws that cannot be changed without difficulty.

Internet Routing

Autonomous System (AS)

- Each AS is a separately managed network.
- An AS is connected to a few other ASes.
- ASes decide the routes that packets will follow.

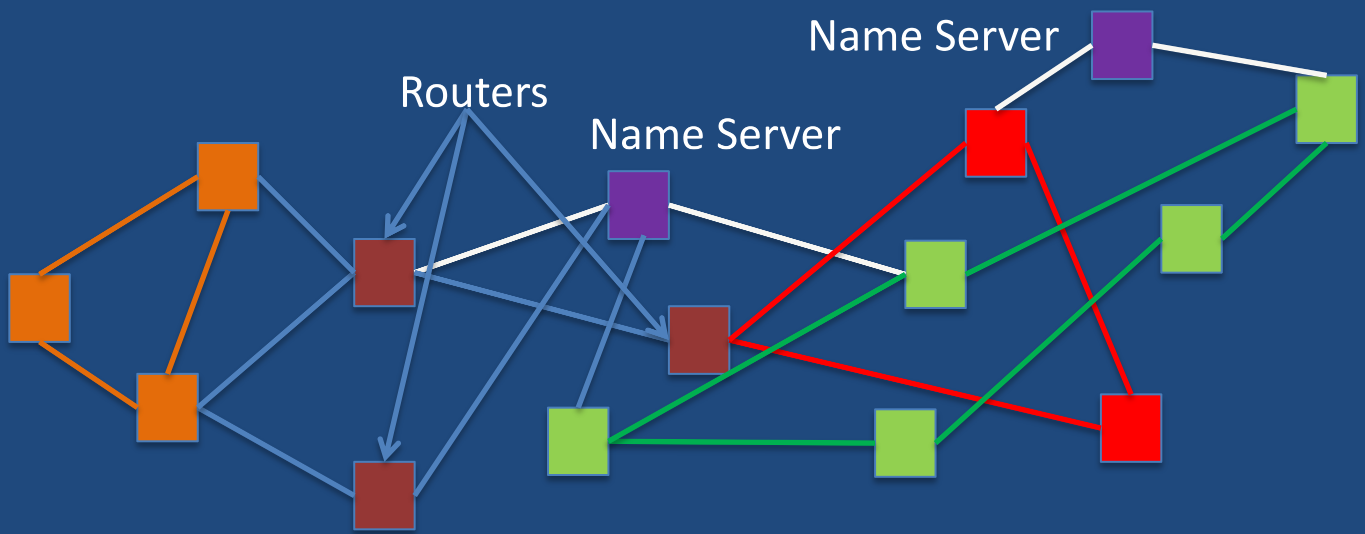

Three ASes, three routers, and two domain name servers (DNS)

Intra-Network Routing

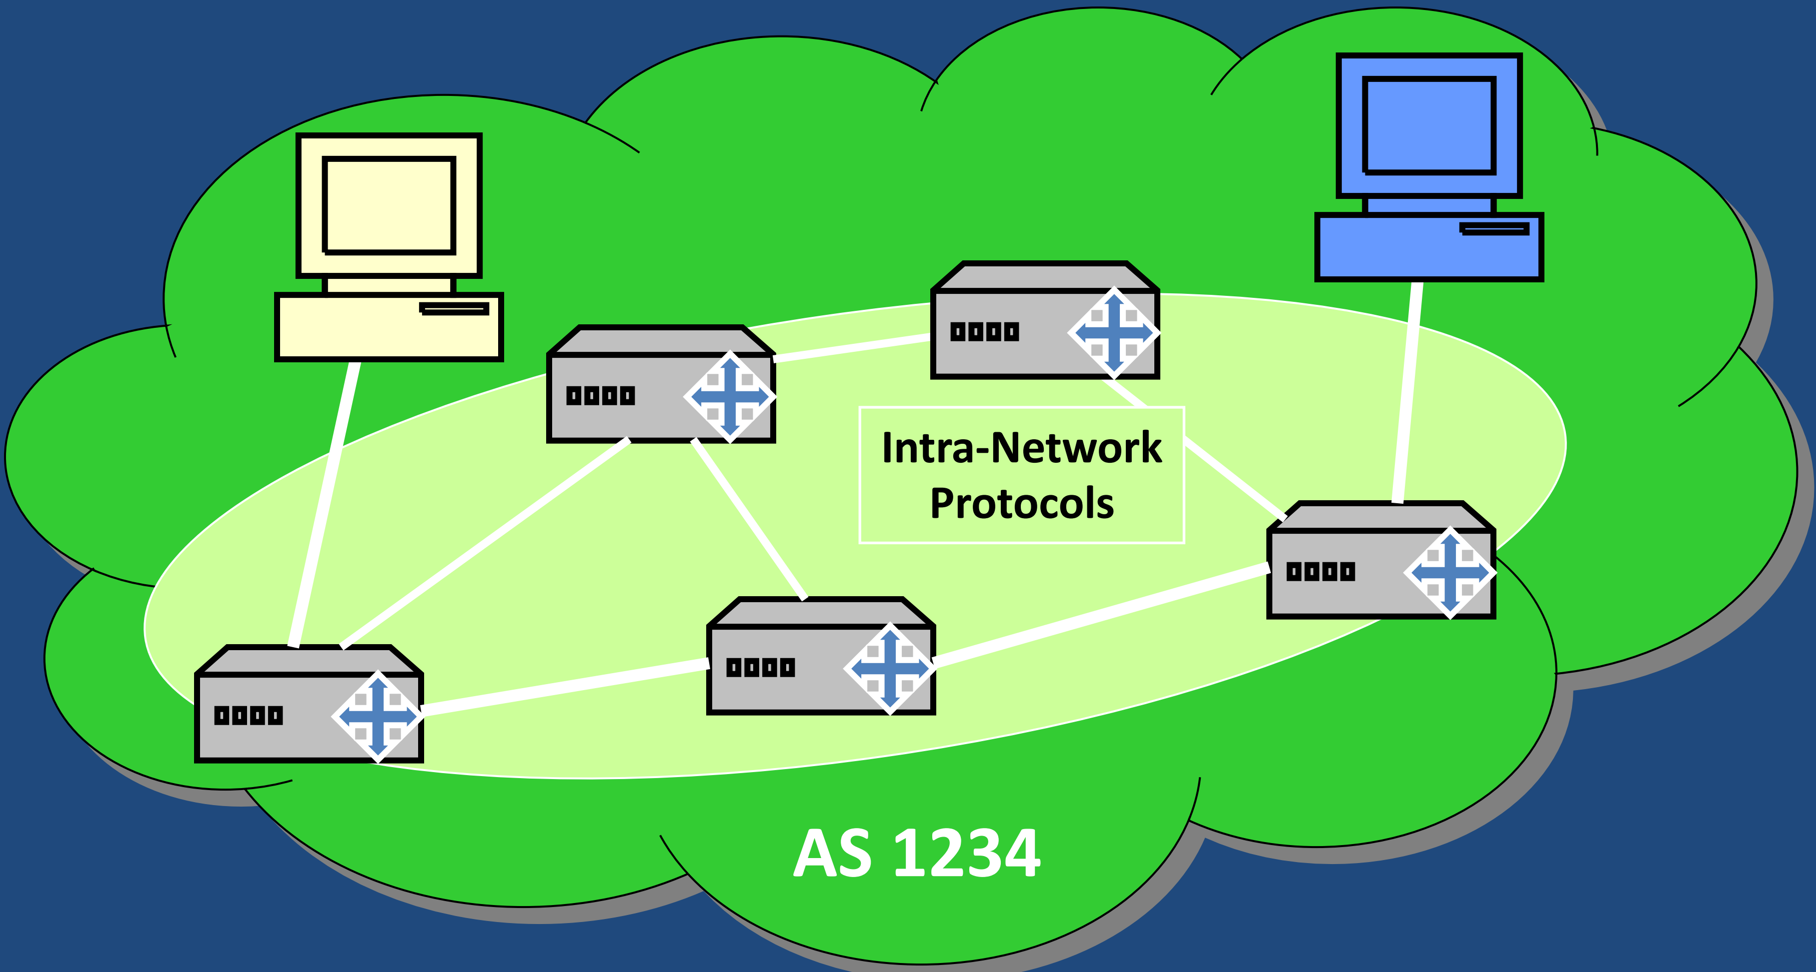

Inter-Network Routing via BGP

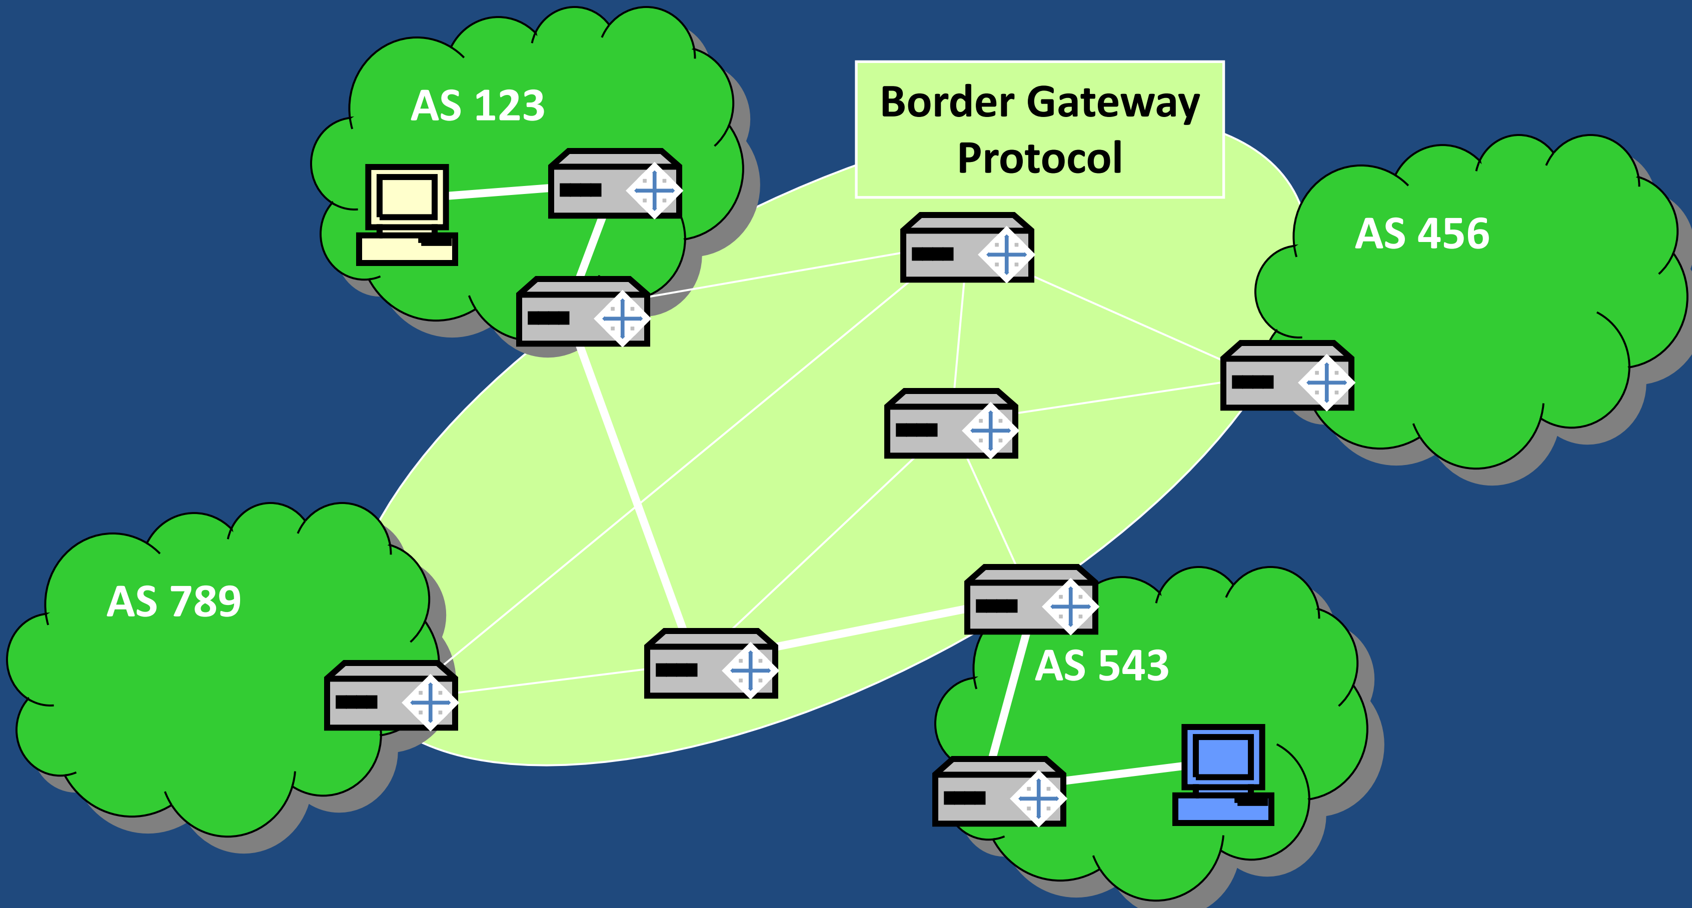

Border Gateway Protocol† (BGP)

- AS announces prefix of IP addresses reachable via it
  - E.g. Prefix 129.6.5.7/16 denotes set of 32-bit addresses with first 16 bits fixed, i.e. [129.6.0.0, …, 129.6.255.255].
- An announcement shows destination set & path:

  <129.6.5.7/16 reachable via [AS42,AS3,AS701,AS49]>
- AS sends its announcements, and those it receives, to its neighbors.
- AS router uses announcements to create routing tables to choose a neighbor to receive a packet.

†See http://www.washingtonpost.com/sf/business/2015/05/31/net-of-insecurity-part-2/

Some Types of BGP Announcements

- Offer to carry traffic to a set of destinations. An AS announces paths to neighbors.
- Withdrawal of offers.
- Changes in paths for a set of destinations.
- New path attributes.

Some Router Actions

- Checks paths for loops
  - A packet has a TTL that is decremented when it passes a router. It is discarded when its TTL reaches 0.
- Impose policy constraints.
  - E.g. Packets starting in Canada must travel in Canada.
- Withdraw a destination when told to do so.
- Propagate announcements to peers
- Compute/update best paths to destinations.

BGP Threats and Risks

- **Routers are too trusting** – attackers may issue announcements that result in
  - Eavesdroping, delay, and/or disruption of traffic.
  - Redirection of traffic to malicious endpoint.
  - Hijacking (temporarily take over) address space to launch spam, run attacks, etc.
  - Denying service – make an entire network disappear

Some Major BGP Hijacks

- Feb 24, 2008 – For about two hours connection to YouTube was lost around the world due to action by Pakistan Telecom
- April 8, 2010 – For 20 mins. routes to 32,000+ networks were sent to China Telecom, taking Facebook, Twitter, etc. offline.
- November 7, 2016 – Twitter went dark for about 30 minutes
- These and many other examples illustrate fragility of BGP.
- Forbes (4/9/10) called BGP announcements cybernukes.

Spamming

- Spammers – biggest abusers of announcements
  - BGP used to “advertise” a route for a block of addresses that were allocated but unassigned.
  - Large amount of spam is sourced from bogus block
  - BGP then used to withdraw the route to the block
  - Spamming source completely disappears.
  - Untraceable, can’t be audited, not prosecutable.

Routing Policy

Some Router Priorities

- Note that a router may have many announcements for a given prefix
- Most-specific-prefix-first – This always preferred
  - Router prefers 129.6.5.7/32 over 129.6.5.7/16
  - That is, for an IP address in both prefixes, choose announcement with most specific prefix
- Shortest-path-first
  - Given multiple announcements for a prefix, choose the shorter path

A Tragedy of the Commons

- BGP routing space is simultaneously
  - *Everyone's problem*, because it impacts the stability and viability of the entire Internet, and
  - *No one's problem*, in that no single entity manages this common resource
- Who's responsible for reliability of the network?
  - End customers?
  - Service providers?
  - Somebody else?

Making BGP More Robust

- Many proposals to make BGP more robust.
- Latest: Resource PKI (RPKI), cryptographically signed BGP announcements.
- Would increase level of trust but introduces many new issues:
  - Trust anchor can shut down networks.
  - Not widely used.

ARTEMIS -**Neutralising BGP Hijacking Within a Minute***

- An AS can protect itself from BGP hijacking.
- Experiments show that an AS can neutralize a hijack within a minute.
- Approach:
  - Monitor – Receive data from public BGP monitors
  - Detect – Compare announcements with own prefixes
  - Mitigate – Replace hijacked prefix with more specific ones

* https://labs.ripe.net/Members/vasileios_kotronis/artemis-neutralising-bgp-hijacking-within-a-minute

Review

- The Domain Name System (DNS)
  - Protecting the DNS from attacks
- History of Naming Policy
- Internet routing
  - The Border Gateway Protocol (BGP)
  - Protecting BGP from attacks
- Routing Policy