Karlsruhe Institute of Technology

Resilient Networking

Module 3: Routing Security

Thorsten Strufe - Disclaimer: This module prepared in cooperation with Mathias Fischer, Michael Roßberg, and Günter Schäfer

Winter Term 2020 –KIT / TUD

Competence Center for Applied Security Technology

Course Overview

- Introduction
- Background
  - Graphs and graph theory
  - Crypto basics (Symmetric/Asymmetric/MACs)
- Resilient Routing (Attacks on BGP, SBGP)
- IPsec
- TLS
- DNS Security
- DDoS and Countermeasures
- Resilient Overlay Networks / Darknets
- Intrusion Detection and Response

Module Outline

- How does routing work in the first place?
- Routing protocols
- Threats to routing
- Attacks on routing protocols
- Defense: S-BGP / SIDR
- Pretty good BGP
- Topology-based routing security

Recall: 4+ Layers of TCP/IP

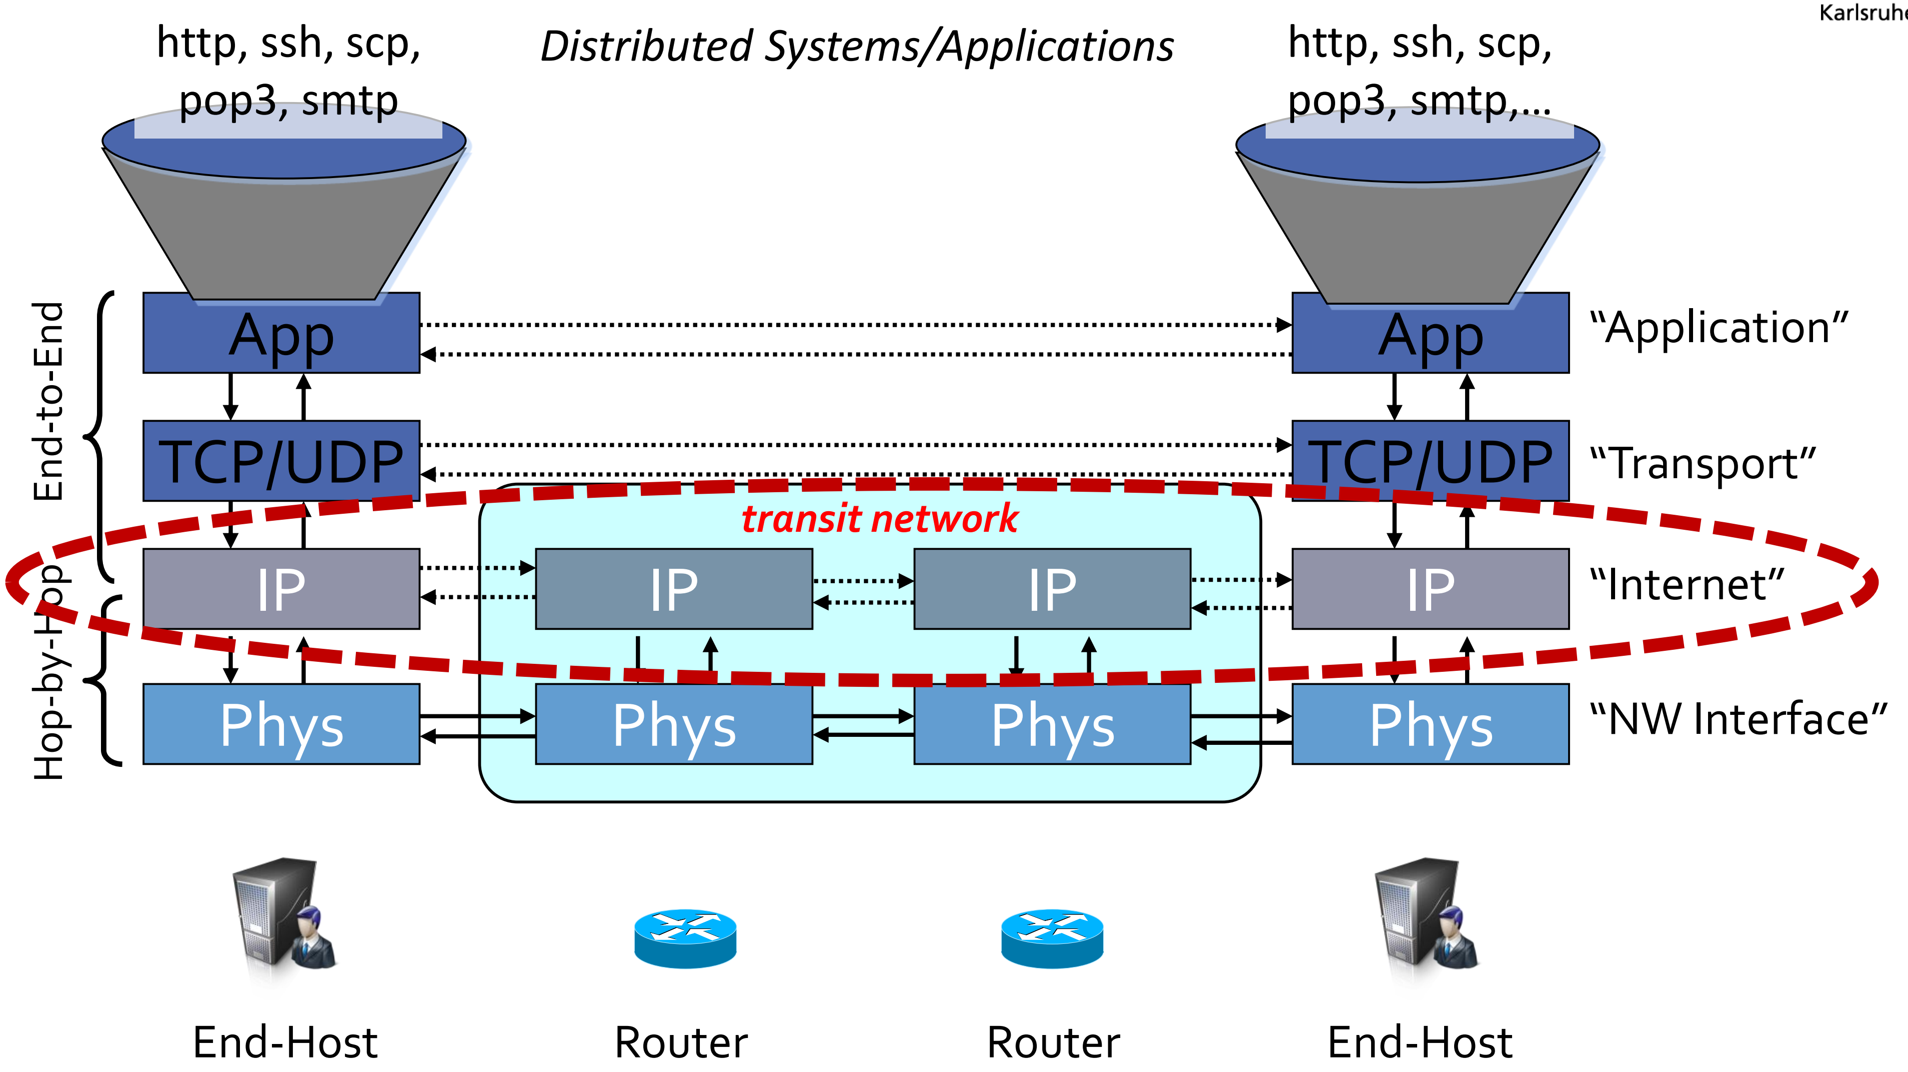

Who to call, where to send my packet?

- Now, h
  cable?
- Here: a
  house..

So where do I send this packet?

Who to call, where to send the packet? ☺

And where do I call???

- N
- H
- E ... aller than H
- C ... on't forget w ... don't a
- S

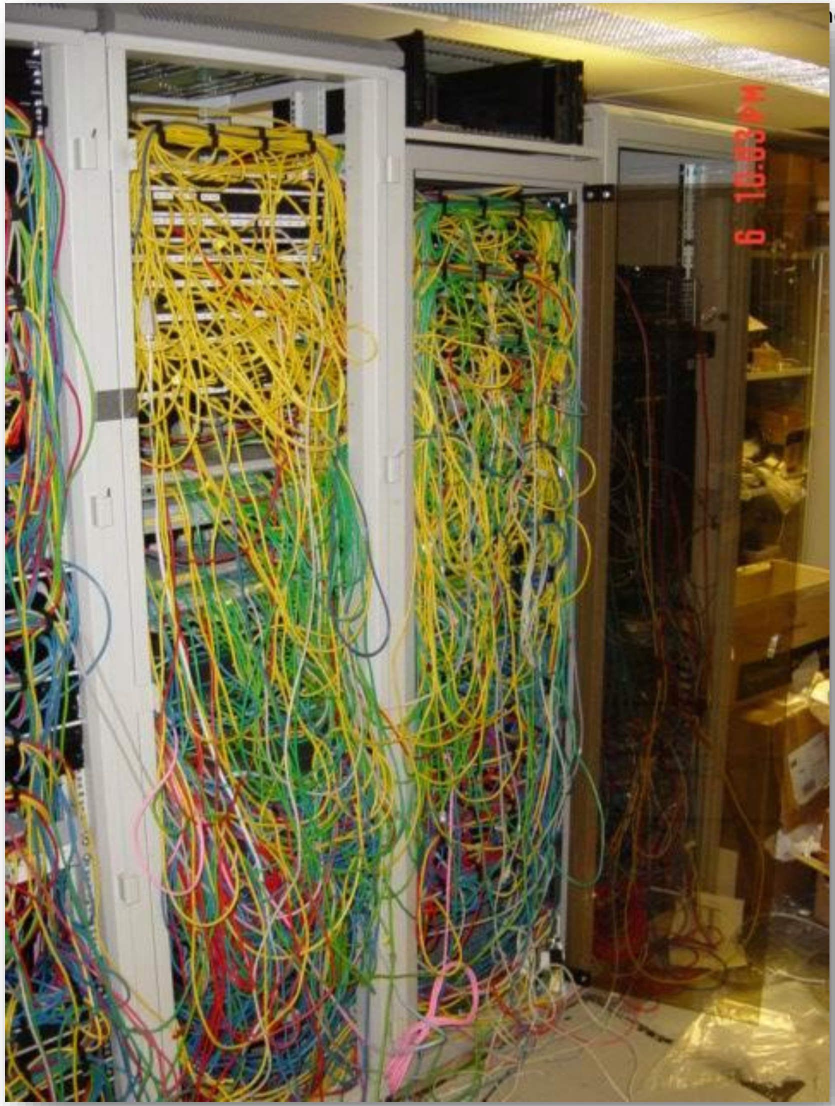

Can it get worse? :-)

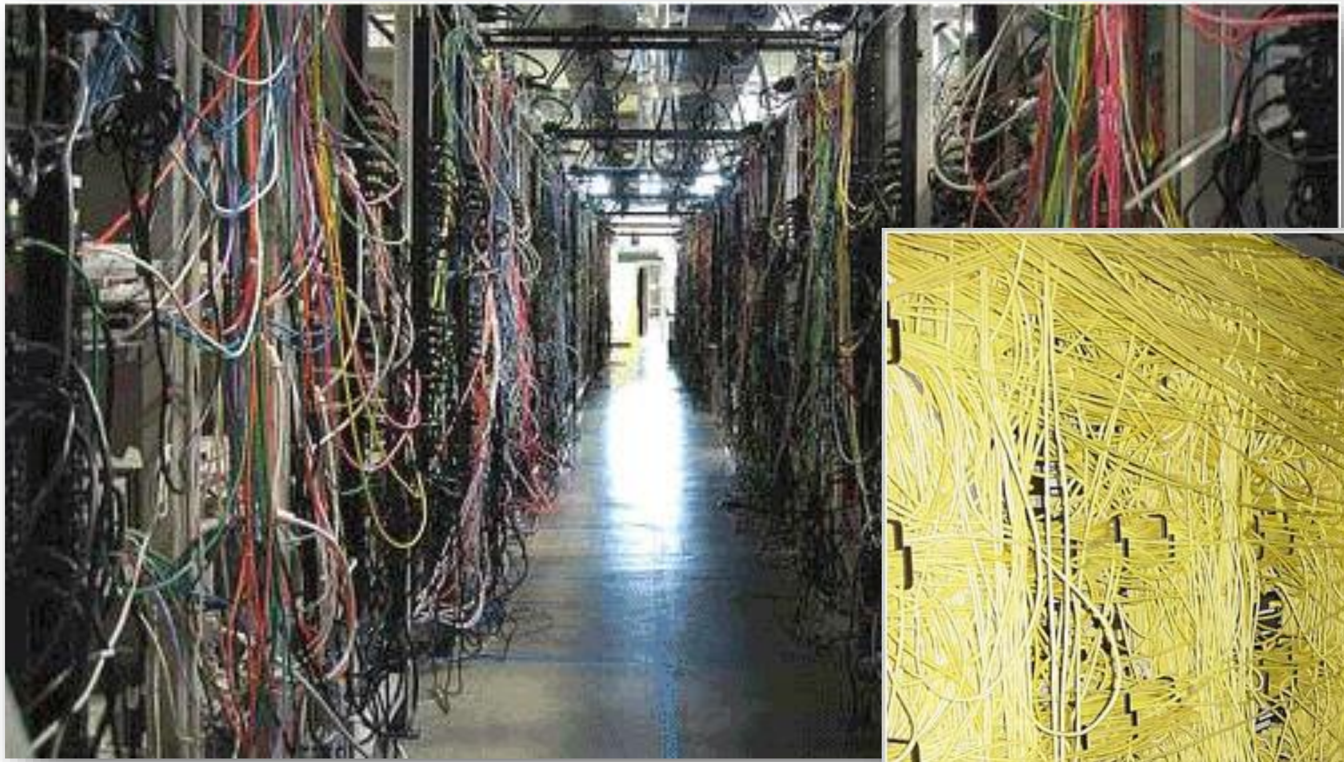

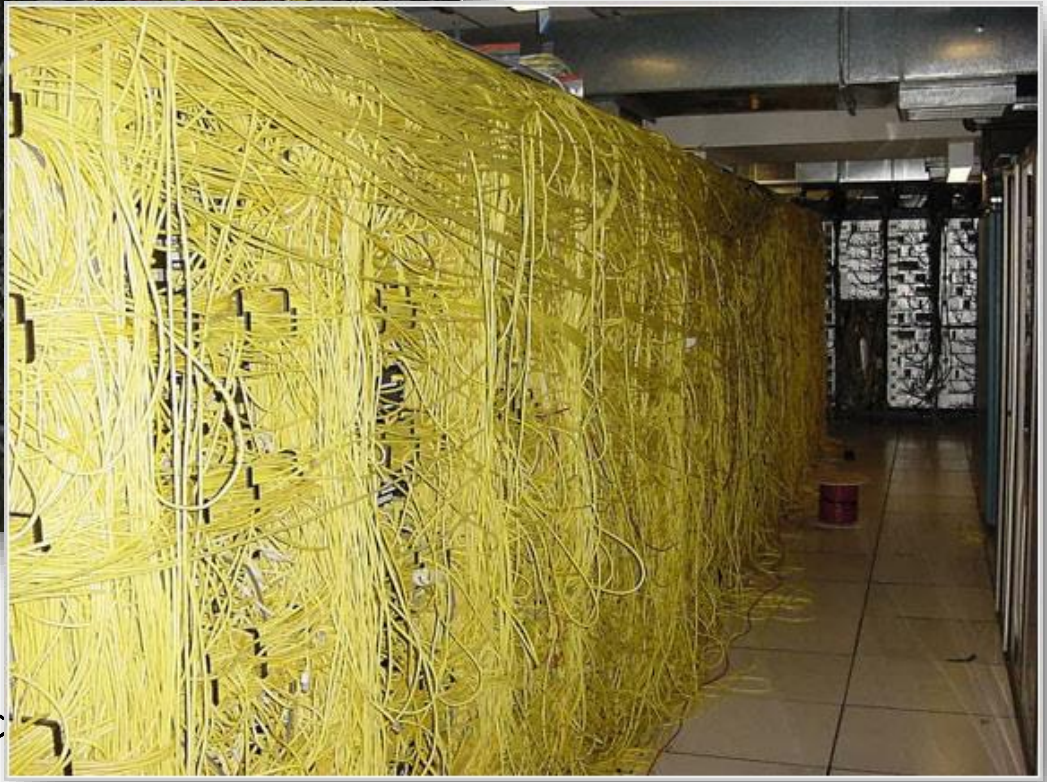

- ...it can!
- But structure helps!?
- These examples are already struc

Network Layer Functions

- Transport packet from sending to receiving hosts
- Network layer protocols in every host and router
- Three important functions:
  - Path determination: route taken by packets from source to destination: Routing algorithms
  - Forwarding: move packets from router's input to appropriate router output
  - Call setup: some network architectures require router call setup along path before data flows (the Internet does not!)

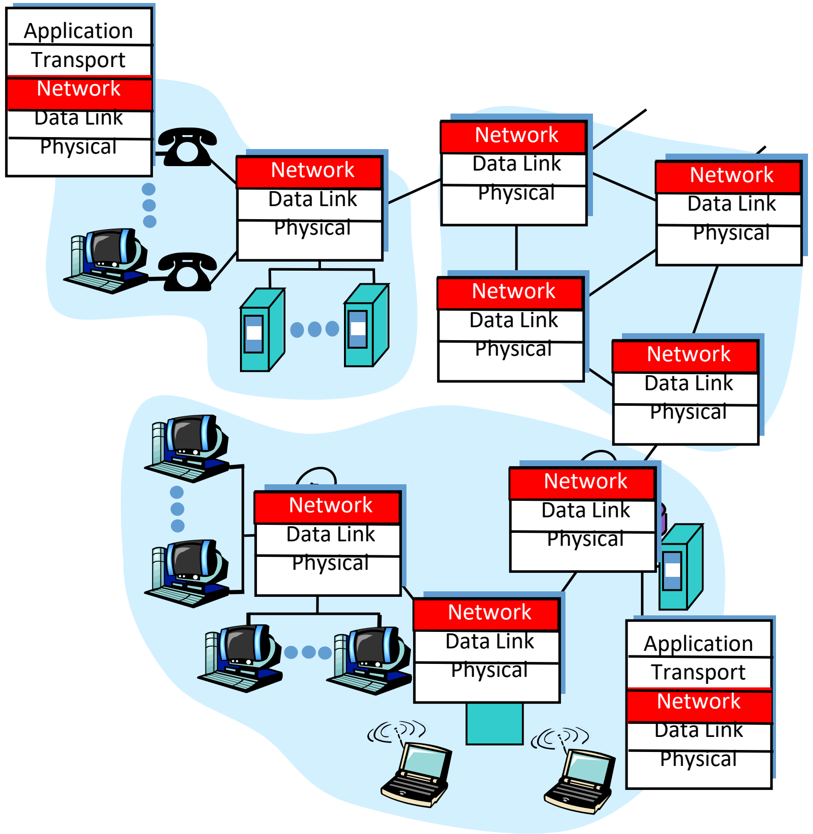

The Internet Network Layer

- No call setup at network layer, no network-level concept of “connection”
- Routers do not process or store state about end-to-end connections
- Packets are typically routed using destination host ID
- Packets between same source-destination pair may take different paths

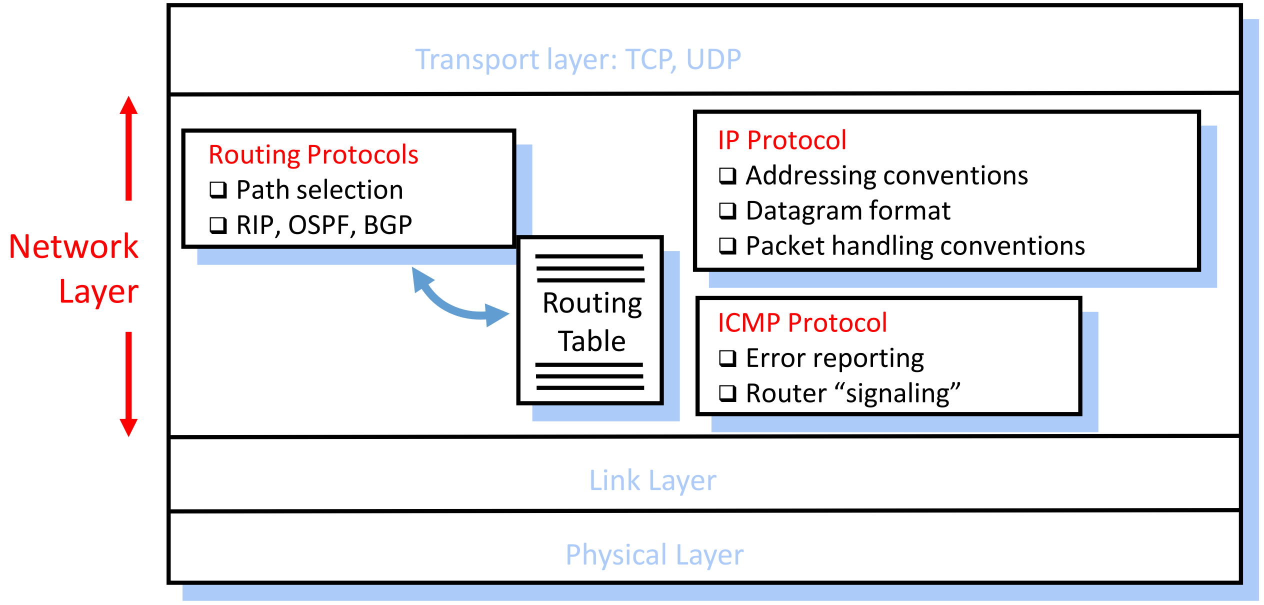

Internet Names and Addresses

	Example	Organization
Host name	www.ietf.org	hierarchical
IP address	132.151.1.35	topological (mostly)
MAC address	8:0:20:72:93:18	flat, permanent

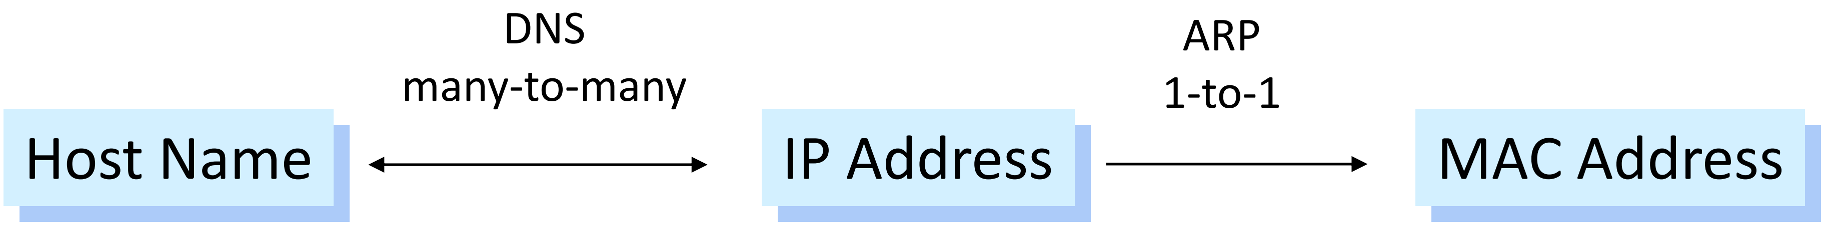

IP Addressing (1)

- IP Address:
  32/128-bit identifier for host or router interface
- Interface:
  Connection between host, router and physical link
  - Router’s typically have multiple interfaces
  - Host may have multiple interfaces
  - IP addresses associated with interface, not host or router

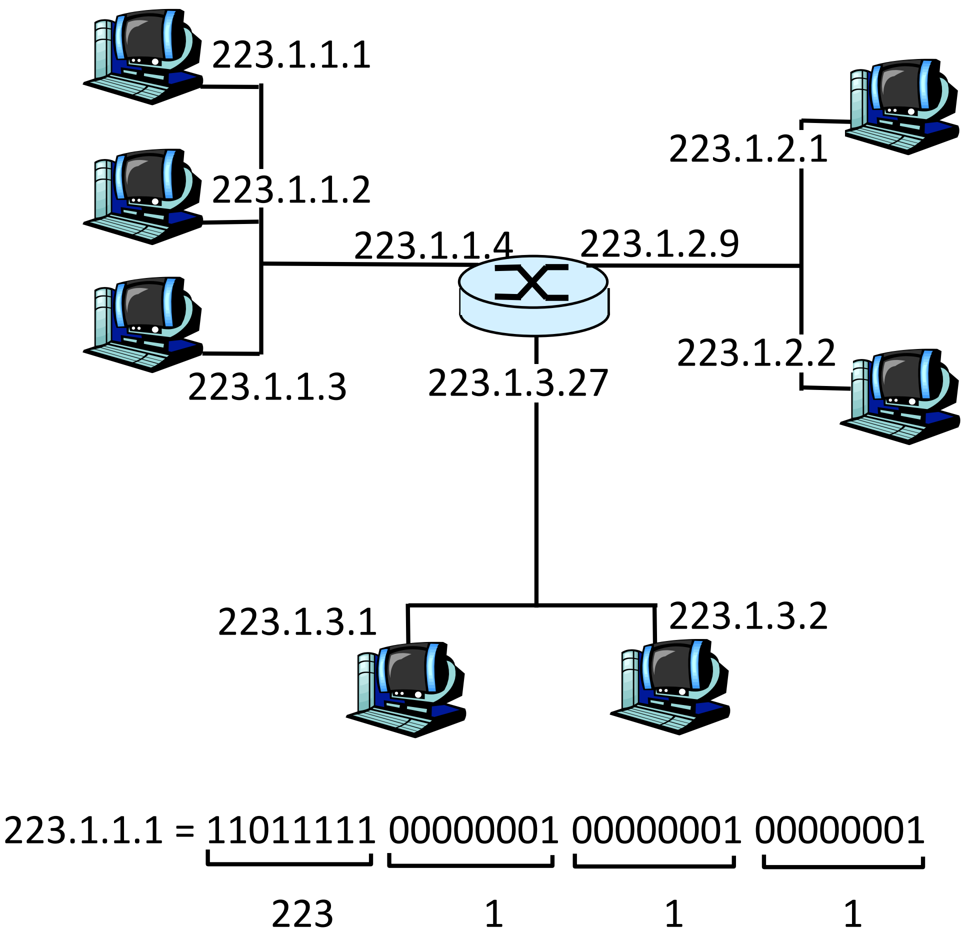

What Happens in the Routers?

- Intermediate Systems (routers) on the Internet have two main functions
- Routing: Determine which route to use
- Forwarding: What happens when a packet arrives?

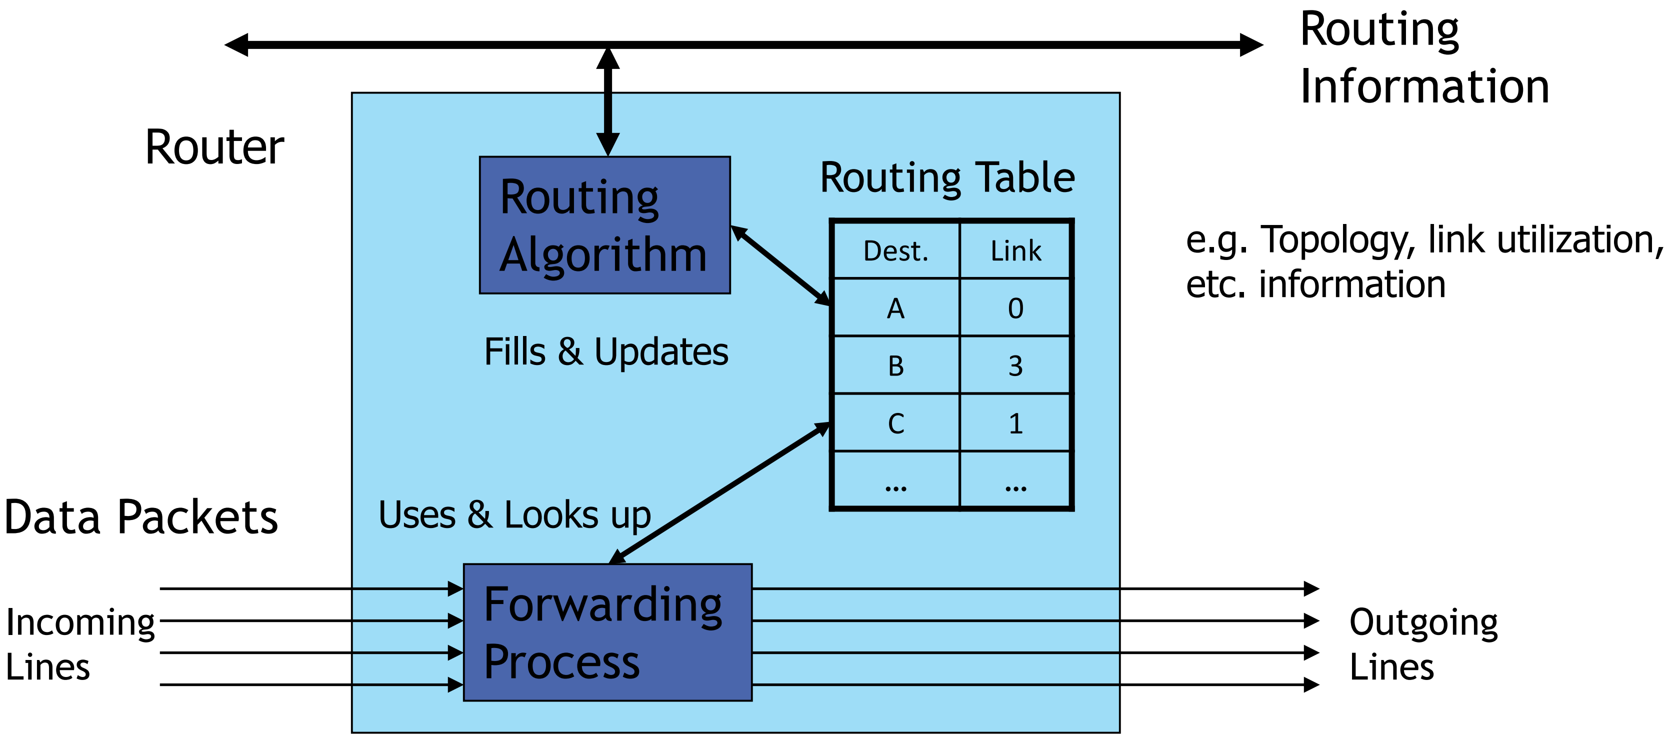

Threats to routing

General Threats to Routing Protocols

- Threats to routing can be characterized according to:
  - Threat source:
    - Subverted link or subverted / rogue router and their possible misbehavior
  - Threat consequence (generic):
    - Disclosure of (routing) information
    - Deception of other routers (e.g. with forged messages)
    - Disruption of normal (router) operation
    - Usurpation (= gaining control over a routers operation, e.g. by “stealing” traffic originally to be routed by that router)
  - Threat consequence zone:
    - Single node / part of a network / whole Internet
  - Threat consequence period:
    - Only during attack / for a certain period of time

Generally Identifiable Routing Threats (1)

- **Disclosing routing information**:
  - Deliberate exposure of routing information: e.g., by a subverted router in order to disclose routing information
  - Eavesdropping on routing exchanges: different attacking technique, also leading to disclosure of routing information
  - Traffic analysis: by eavesdropping on forwarded data traffic, an attacker can gain insight about routing information

- **Masquerade**:
  - An entity claims the identity of a router (sometimes also called spoofing)
  - Masquerade is usually performed in order to realize further attacks

- **Interference**:
  - An attacker inhibits the exchange of routing information between routers, e.g. by delaying or deleting routing messages or receipts, breaking synchronization, etc.
  - The consequence is usually (partial) disruption of routing operations

Generally Identifiable Routing Threats (2)

Karlsruhe Institute of Technology

- **Falsification of routing information**:
  - Either by an originator (forging) or a forwarder (modification)

- **Overclaiming**:
  - Announcing better routes / link capacity than available
  - Goals can be to attract traffic to a certain area in order to control the traffic or to mislead the traffic so that it will not be delivered at all or with higher delay
  - Consequences for the network are potential overload of single routers, increase of overall traffic load

- **Underclaiming**:
  - Announcing inferior routes / link capacities than actually exist
  - Potential goals are to keep traffic out of certain areas of the network, e.g. in order to avoid forwarding of traffic at certain routers or to increase attractiveness of alternative routes
  - Potential consequences are that certain destinations become unreachable, and the overall traffic load in the network increases (because packets take inferior routes)

Generally Identifiable Routing Threats (3)

Karlsruhe Institute of Technology

- **Resource exhaustion**:
  - E.g. by an attacker that announces frequent changes in his routing information, or triggers a router to create an excessive amount of state information which can not be handled by other routers
  - Sometimes also referred to as overload
  - Goal is degradation / disruption of routing protocol operation
- **Resource destruction**:
  - Link destruction: either physically ("cutting") or by strong interference
  - Node destruction: e.g., physically or logically by exploiting weaknesses in the router software (OS, routing software)
  - Depending on the network topology, the consequences can be either of local or global scope (single network / part of network unreachable or network partitioning)

Routing Threats Consequences (General)

Karlsruhe Institute of Technology

- **Network congestion**:
- **“Blackhole”**
- **Looping**:
- **Partitioning**:
- **Frequent route changes**:
- **Instability of the routing protocol**:
- **Routing overload**:

Routing Threats Consequences (General)

Karlsruhe Institute of Technology

- **Network congestion**: more traffic is routed through a specific part of the network than would usually be
- "**Blackhole**": packets go into a certain router/region and "disappear"
- **Looping**: traffic is forwarded along a route that loops (this causes both traffic to disappear and congestion)
- **Partitioning**: some portion of the network believes that it is partitioned from the rest of the network when in fact it is not
- **Frequent route changes**: resulting in unnecessary routing processing and message exchanges as well as large variations in forwarding delay
- **Instability of the routing protocol**: convergence towards a global forwarding state is not achieved
- **Routing overload**: routing protocol messages become a significant part of the overall transported traffic the network carries

Routing Threats Consequences (Specific)

- Consequences regarding a specific target host / network:

Routing Threats Consequences (Specific)

- Consequences regarding a specific target host / network:
  - **Delay**: traffic from / to a target host / network is routed along routes that are inferior to the route the traffic would otherwise take
  - **Cut**: some part of the network believes that there is no route to the target host / network when, in fact, there is
  - **Starvation**: the traffic destined for the target host/network is routed to a part of the network that can not deliver it
  - **Eavesdropping**: traffic is routed through some router or network that would normally not “see” this traffic, so that an attacker can eavesdrop on the traffic or at least monitor the traffic pattern
  - **Controlled delivery**: traffic is routed through a router / network so that an attacker can selectively delay, delete or modify packets destined to a target host / network

(Local) Countermeasures: IP Fast Reroute

- Link or node failures result in period of disruption to the delivery of traffic

- IP Fast Reroute as restoration mechanism on network layer until network and its (Intra-AS) routing finally re-converges again
  - Local backup computation without need of informing neighbors about failure

- Variations of IP Fast Reroute mechanisms
  - Equal Cost Multipath (ECMP)
  - Loop-free Alternates
  - U-Turn Alternates
  - Tunnels
  - Not-via Addresses

IP Fast Reroute - ECMP

- Failover to detour shortest paths, so we need more paths…

- Equal Cost Multi-Path (ECMP)
  - Routers can reach destination by multiple preferably link-disjoint paths of same cost
  - Alternate paths can be pre-computed and used in case of failure

$$cost(r_i, d) = cost(r_j, d), i \neq j$$

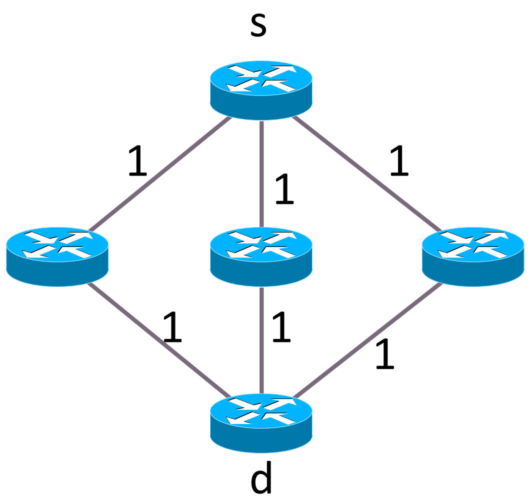

IP Fast Reroute – Loop-free Alternates (LFA)

Karlsruhe Institute of Technology

- Using pre-computed alternate next hop a in event of link failure that disrupts primary next hop p
- While choosing loop-free alternates, micro loops need to be prevented:

$$cost(a,d) < linkcost(a,s) + cost(s,d), \qquad cost(a,d) < cost(s,d)$$

- In case of failure, traffic flow from source s to destination d is not disrupted
- Network converges in background
- Failure of link (r_0, r_2)
  - Primary next hop $p = r_2$ of node r_0 not reachable
  - Rerouting at $\boldsymbol{s = r_0}$ by using $\boldsymbol{r_1}$ as LFA
- Failure of primary next hop $p = r_2$ of node r_0
  - At $s = r_1$ rerouting via r_0 and via $\boldsymbol{r_3}$ possible
  - Rerouting via r_0 prevented as $cost(r_2, d) < cost(r_0, d)$

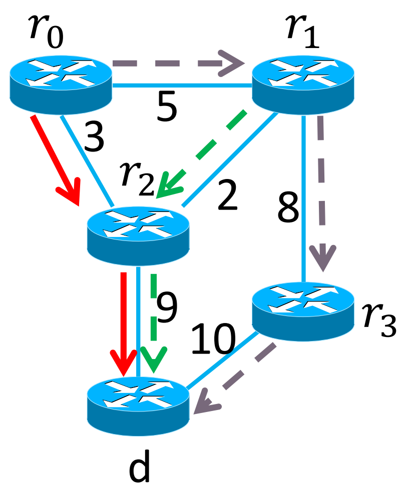

KASTEL

IP Fast Reroute – U-Turn Alternates

- LFA is topology dependent, thus in certain scenarios no LFAs exist
- U-Turn alternate: use neighbor a whose primary next hop to d is again s and that itself has loop free node protecting alternate that does not go through s
- U-Turn selection criteria
  - $cost(a, d) \geq linkcost(a, s) + cost(s, d)$
  - s must always be primary next hop neighbor on all shortest paths from a to d that traverse s
  - Node a has a loop free alternate
- Identifying U-Turn traffic
  - Implicit identification as traffic would be usually forwarded via neighbor
  - Explicit identification via packet marking
- Example: Failure of link (r_0, r_2) results in $s = r_0$ using r_1 as U-Turn alternate

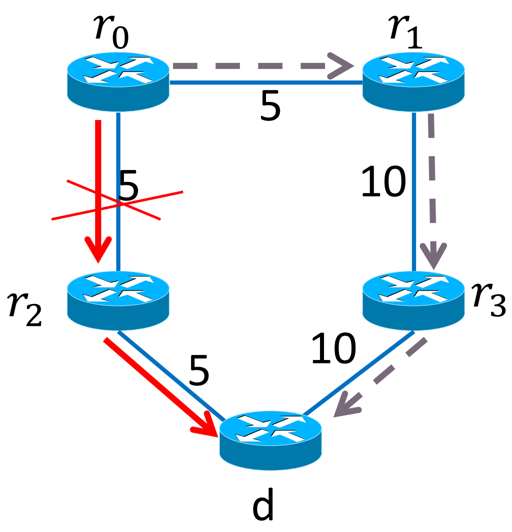

IP Fast Reroute – Tunnels

- When router s detects adjacent failure it uses pre-computed set of repair paths to bypass failed neighbor v or failed link e
- Tunnel is used to carry traffic to a router (tunnel endpoint) at which loop free alternate paths exist via normal forwarding
  - Packets are encapsulated by s and routed towards tunnel endpoint
  - Tunnel endpoint decapsulated packets and forwards them according to shortest path table depending on final destination
- Repair paths
  - Must be created for all neighbors of v plus for v itself to provision for link failures
  - Micro loop prevention necessary, e.g., tunnel endpoint on shortest path to d

IP Fast Reroute – Not-Via Addresses (1)

Karlsruhe Institute of Technology

- Not-Via Addresses are special addresses assigned to each interface
- Once failure is detected, router needs to get packet to destination not via failure
- tunnels traffic towards the Not-via address of the protected component
- Packets addressed to a Not-via address must be delivered to router advertising address, not via component with which address is associated
- Mechanism requires participation of intermediate routers on repair path
    - Routers must be able to tell the link which they must avoid traversing from the semantics of the Not-via address
    - Every router in repair path routes tunneled traffic using shortest paths obtained by running Dijkstra on graph where protected link is excluded

IP Fast Reroute – Not-Via Addresses (2)

- Example
  - Node S routes traffic to D through B and via P
  - When P has failed, S tunnels traffic towards Not-via address Bp of protected component, which is interpreted as shortest path from S to B not going via P

- Drawbacks
  - Each router needs pre-computed backup routes
  - Each router has possibility of being in any shortest path of an initiated repair
  - N-1 Dijkstra calculations required for each of the N routers in the topology

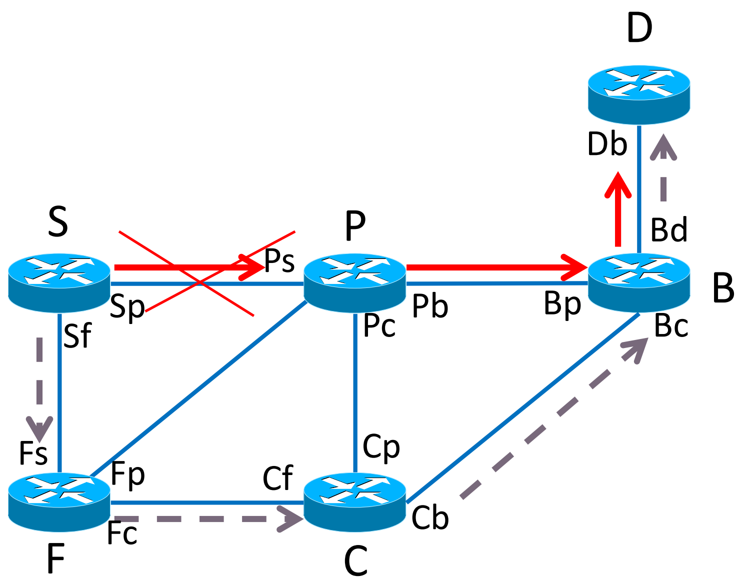

So far (within AS), so good…

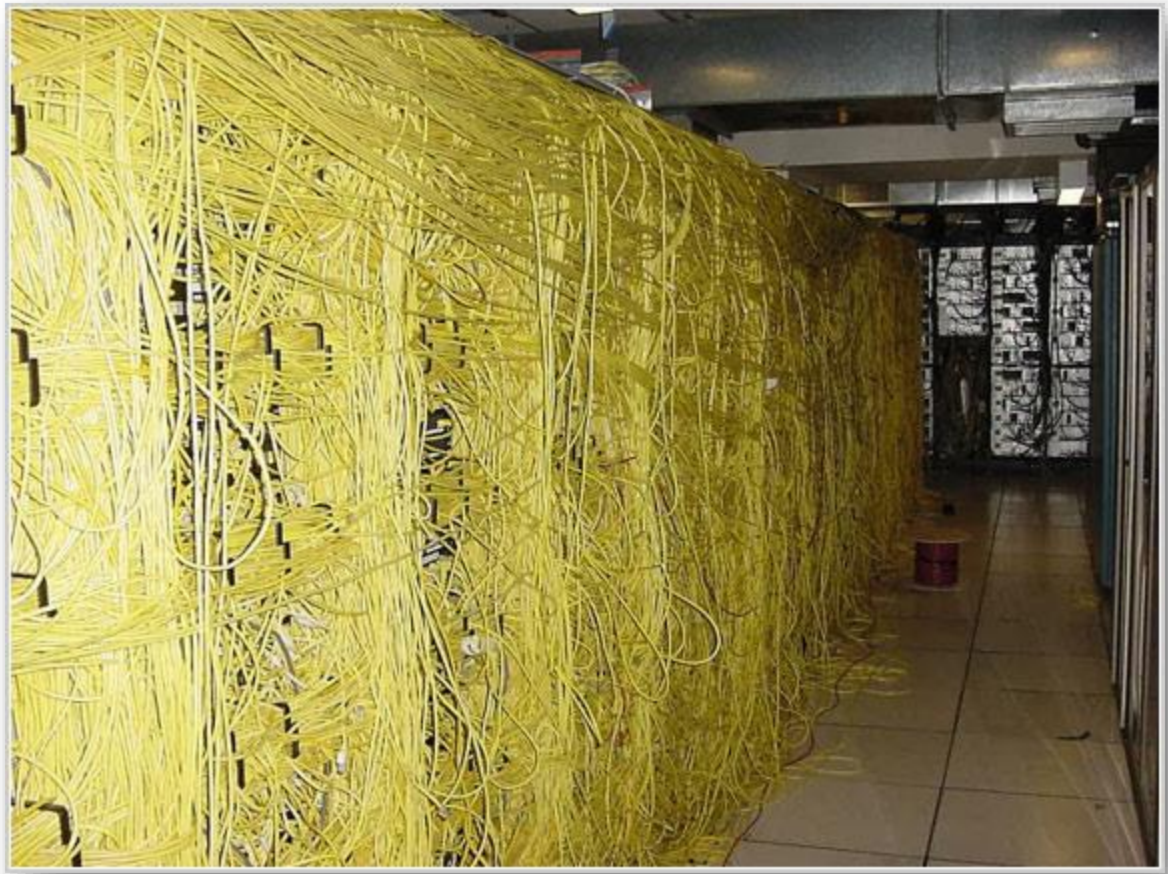

The Growth of The Net, Users…

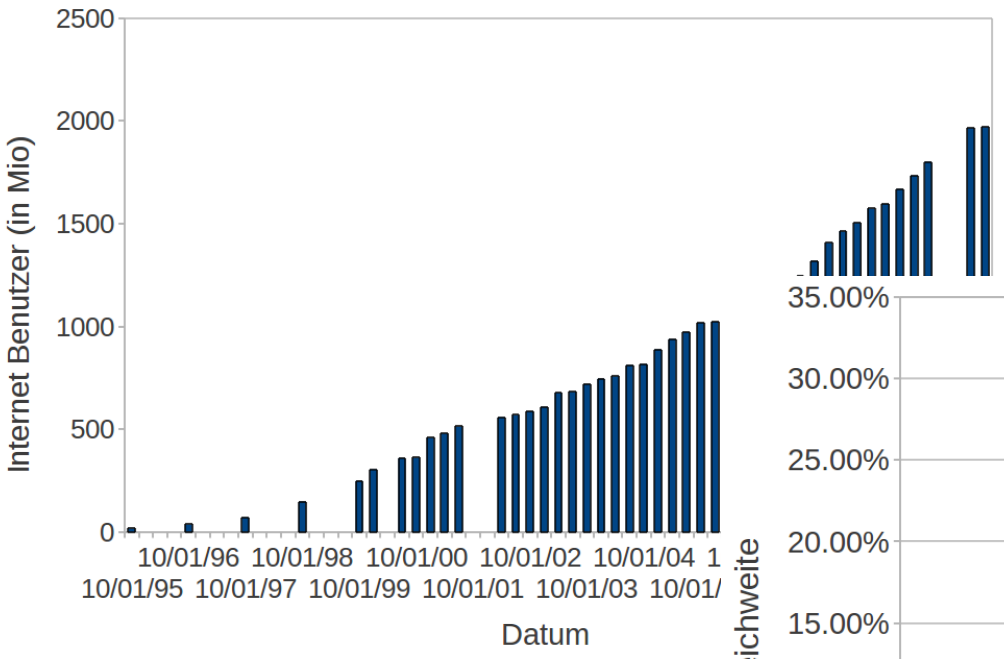

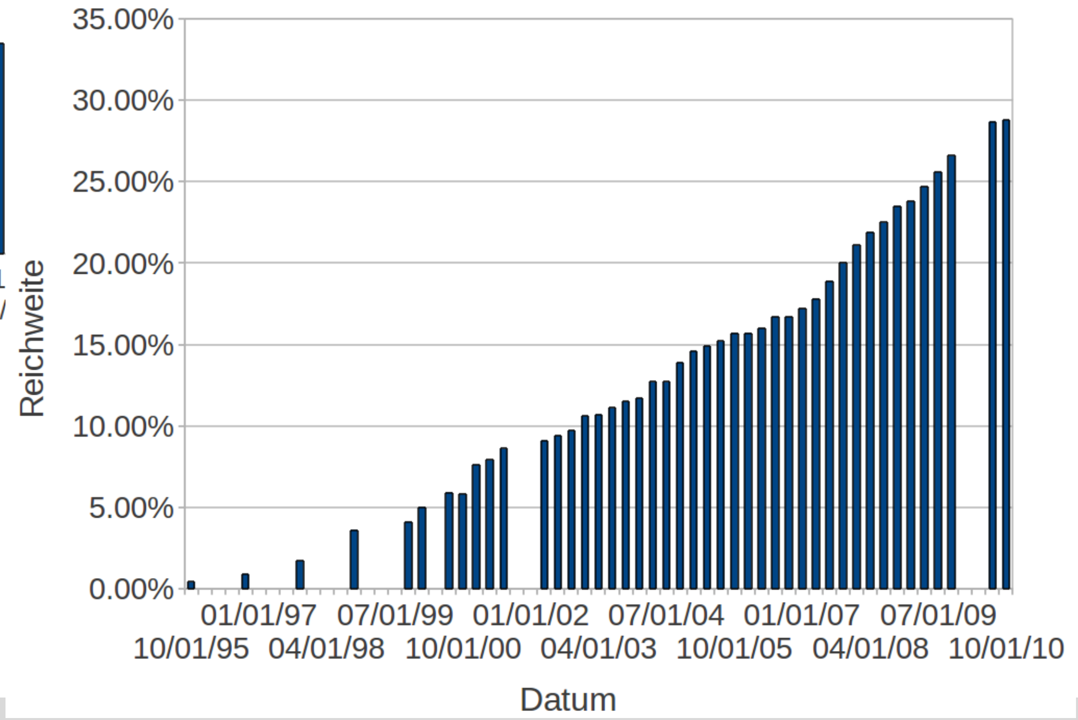

Growth: Example of Germany

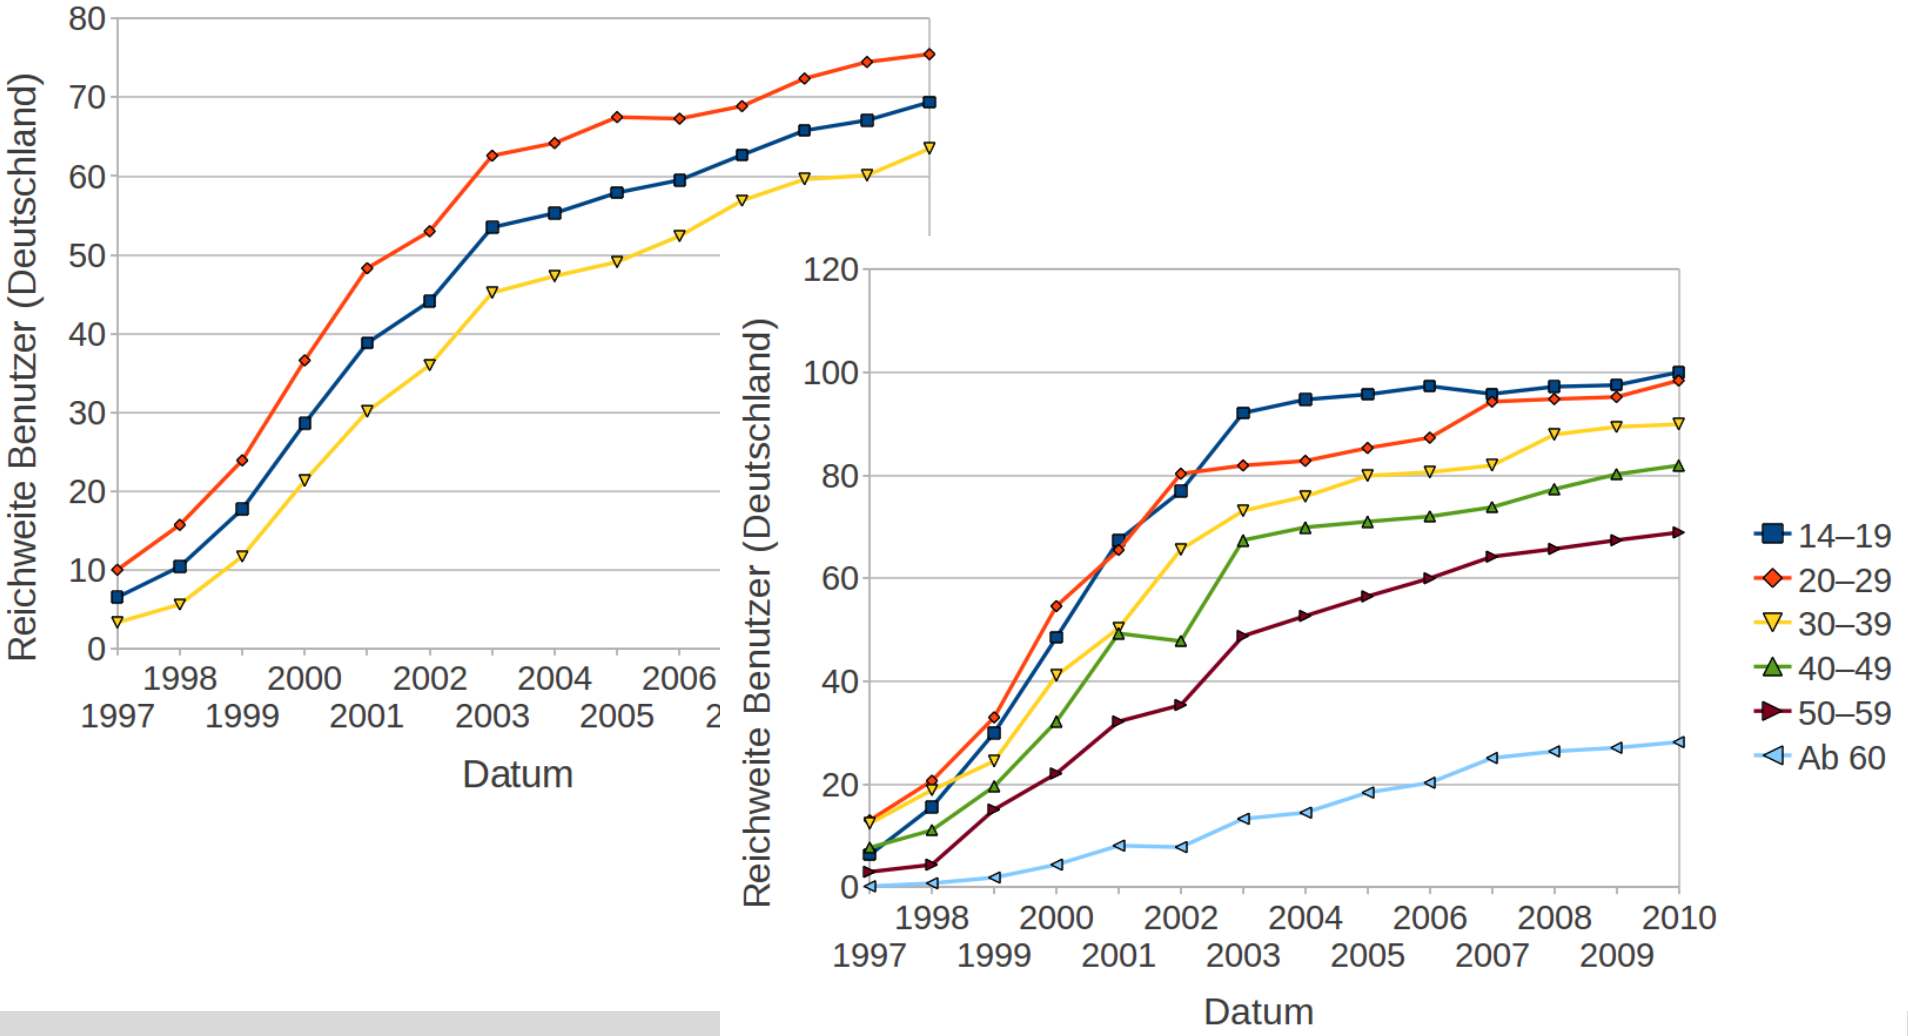

Growth: Devices, Worldwide

Karlsruhe Institute of Technology

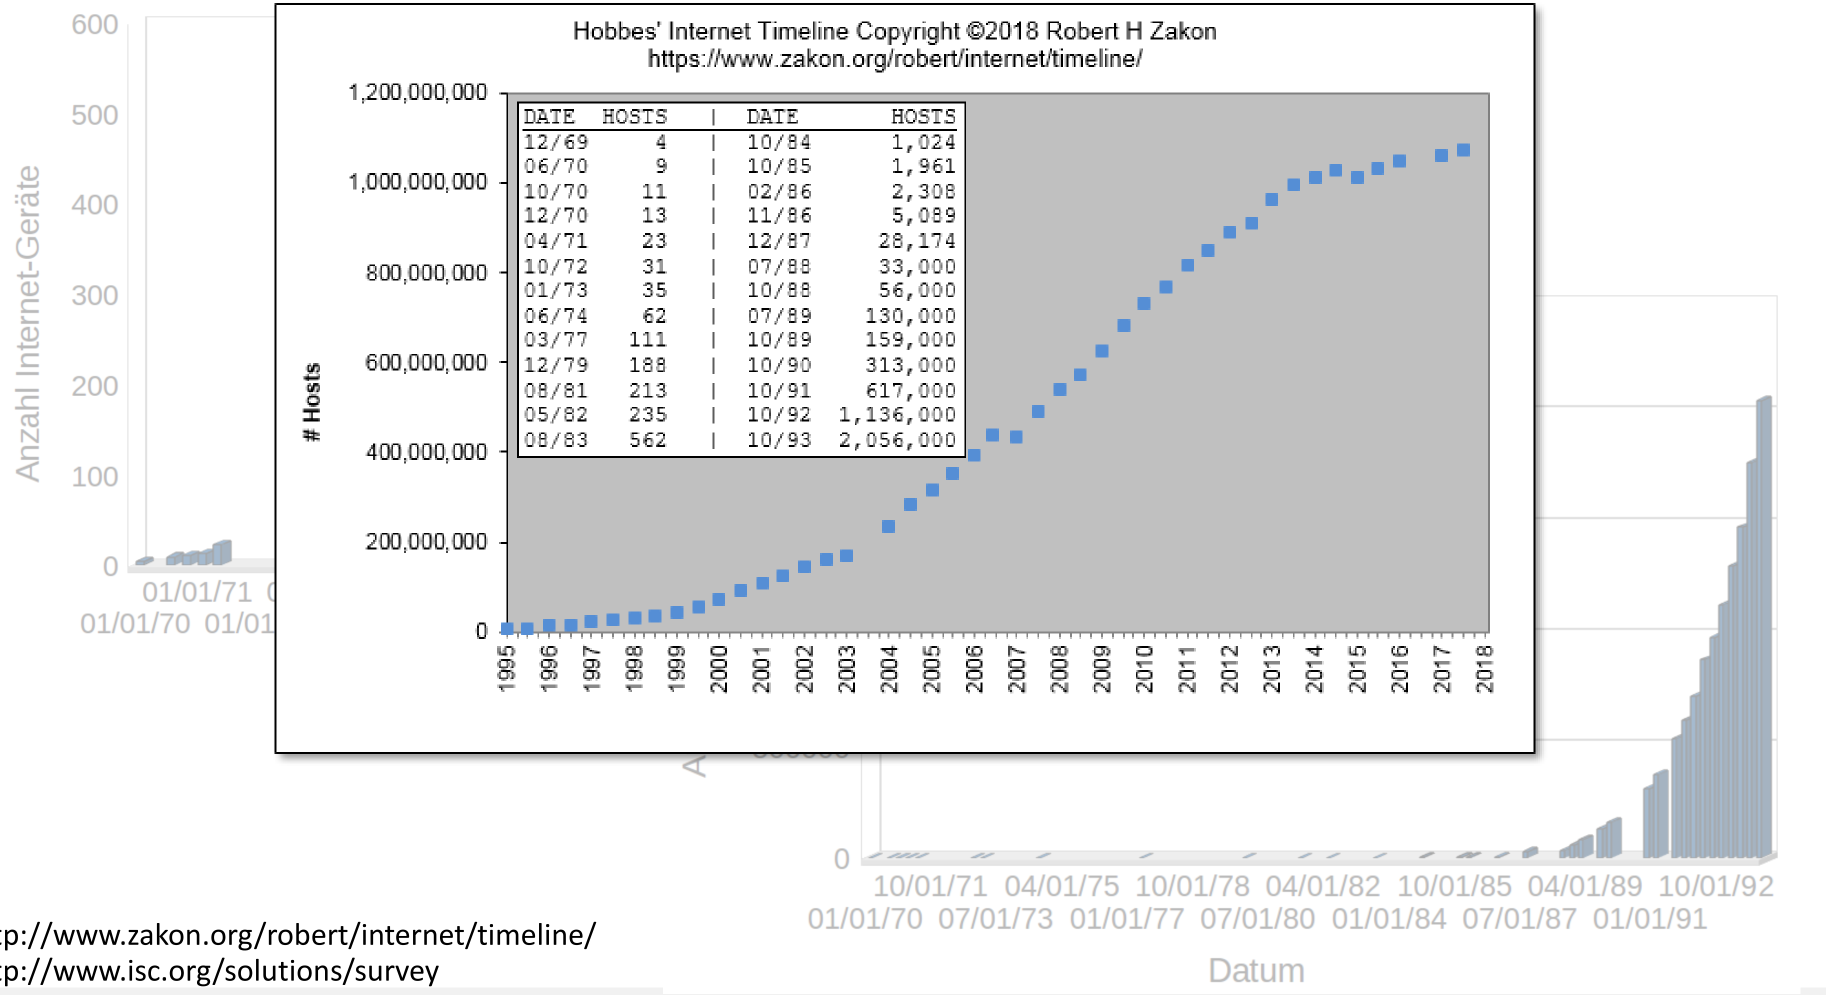

DATE	HOSTS	DATE	HOSTS
12/69	4	10/84	1,024
06/70	9	10/85	1,961
10/70	11	02/86	2,308
12/70	13	11/86	5,089
04/71	23	12/87	28,174
10/72	31	07/88	33,000
01/73	35	10/88	56,000
06/74	62	07/89	130,000
03/77	111	10/89	159,000
12/79	188	10/90	313,000
08/81	213	10/91	617,000
05/82	235	10/92	1,136,000
08/83	562	10/93	2,056,000

http://www.zakon.org/robert/internet/timeline/
http://www.isc.org/solutions/survey

Hierarchical Routing

- Usual routing lectures are an idealization:
  - All routers are assumed to be identical
  - Network is assumed to be “flat”
  - ==> Dijkstra.
- … Real world, however, looks different

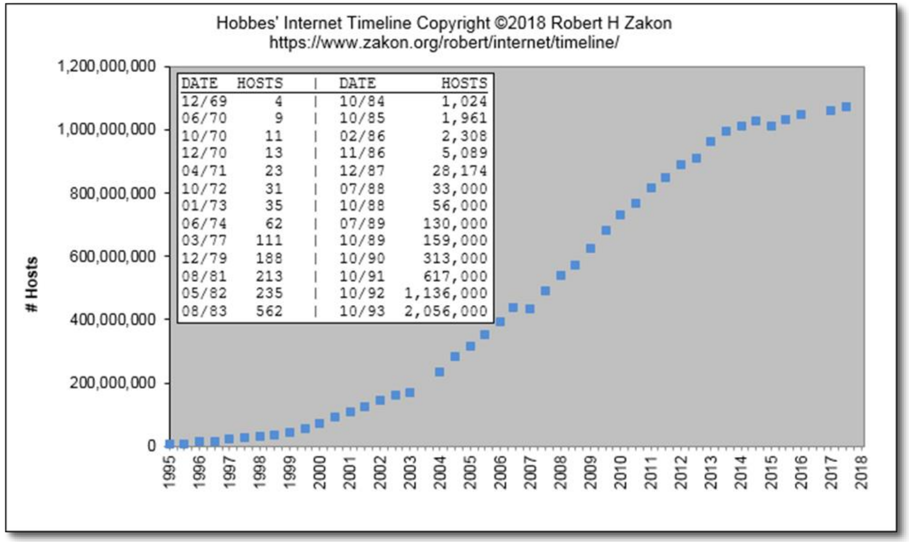

DATE	HOSTS	DATE	HOSTS
12/69	4	10/84	1,024
06/70	9	10/85	1,961
10/70	11	02/86	2,308
12/70	13	11/86	5,089
04/71	23	12/87	28,174
10/72	31	07/88	33,000
01/73	35	10/88	56,000
06/74	62	07/89	130,000
03/77	111	10/89	159,000
12/79	188	10/90	313,000
08/81	213	10/91	617,000
05/82	235	10/92	1,136,000
08/83	562	10/93	2,056,000

Scale (>1 billion destinations!):

- Can’t store all destinations in routing tables
- Routing table exchange would overload links

Administrative autonomy:

- Internet = network of networks
- Each network admin may want to control routing in its own network

[1]https://www.cia.gov/library/publications/the-world-factbook/rankorder/2184rank.html

Routing on the Internet

- The Global Internet consists of Autonomous Systems (AS) interconnected with each other:
  - Stub AS: small corporation (only one link to the Internet)
  - Multihomed AS: large corporation (multiple links, but no transit)
  - Transit AS: provider
- On the Internet today we see many autonomous systems (~70350)
  - Have different sizes
  - Exchange services with each other as equals or as provider/customer
  - Have different relations to each other
- Every AS has a unique number
- Every AS must know a route to every network

http://www.cidr-report.org/as2.0/

Routing on the Internet

- The Global Int… …S) interconnected with each othe…
  - Stub AS: sm…
  - Multihome…
  - Transit AS: …
- On the Interne… …ns (>70350)
  - Have differ…
  - Exchange s… …r
  - Have differ…
- Every AS has a …
- Every AS must know a route to every network

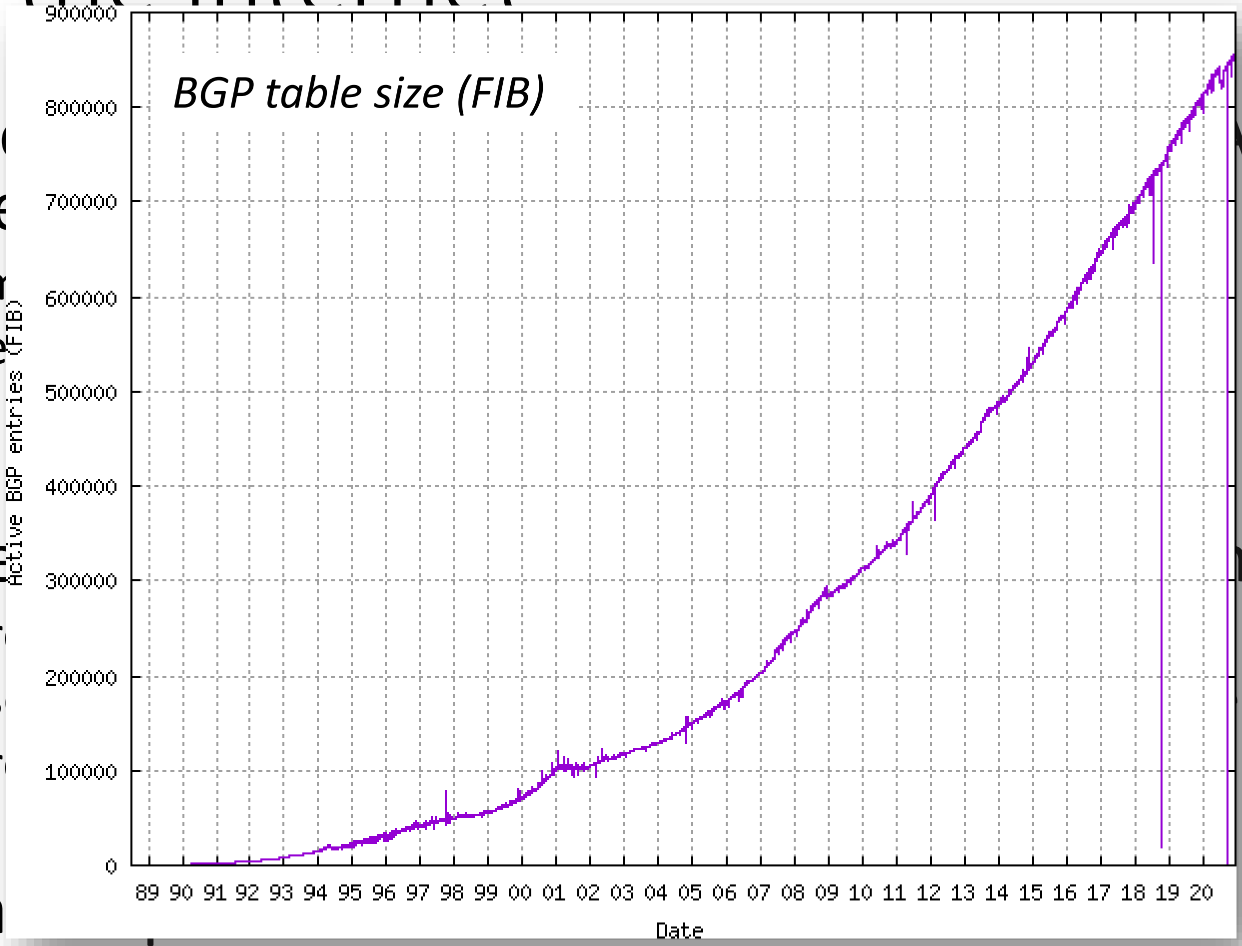

http://www.cidr-report.org/as2.0/

Routing on the Internet

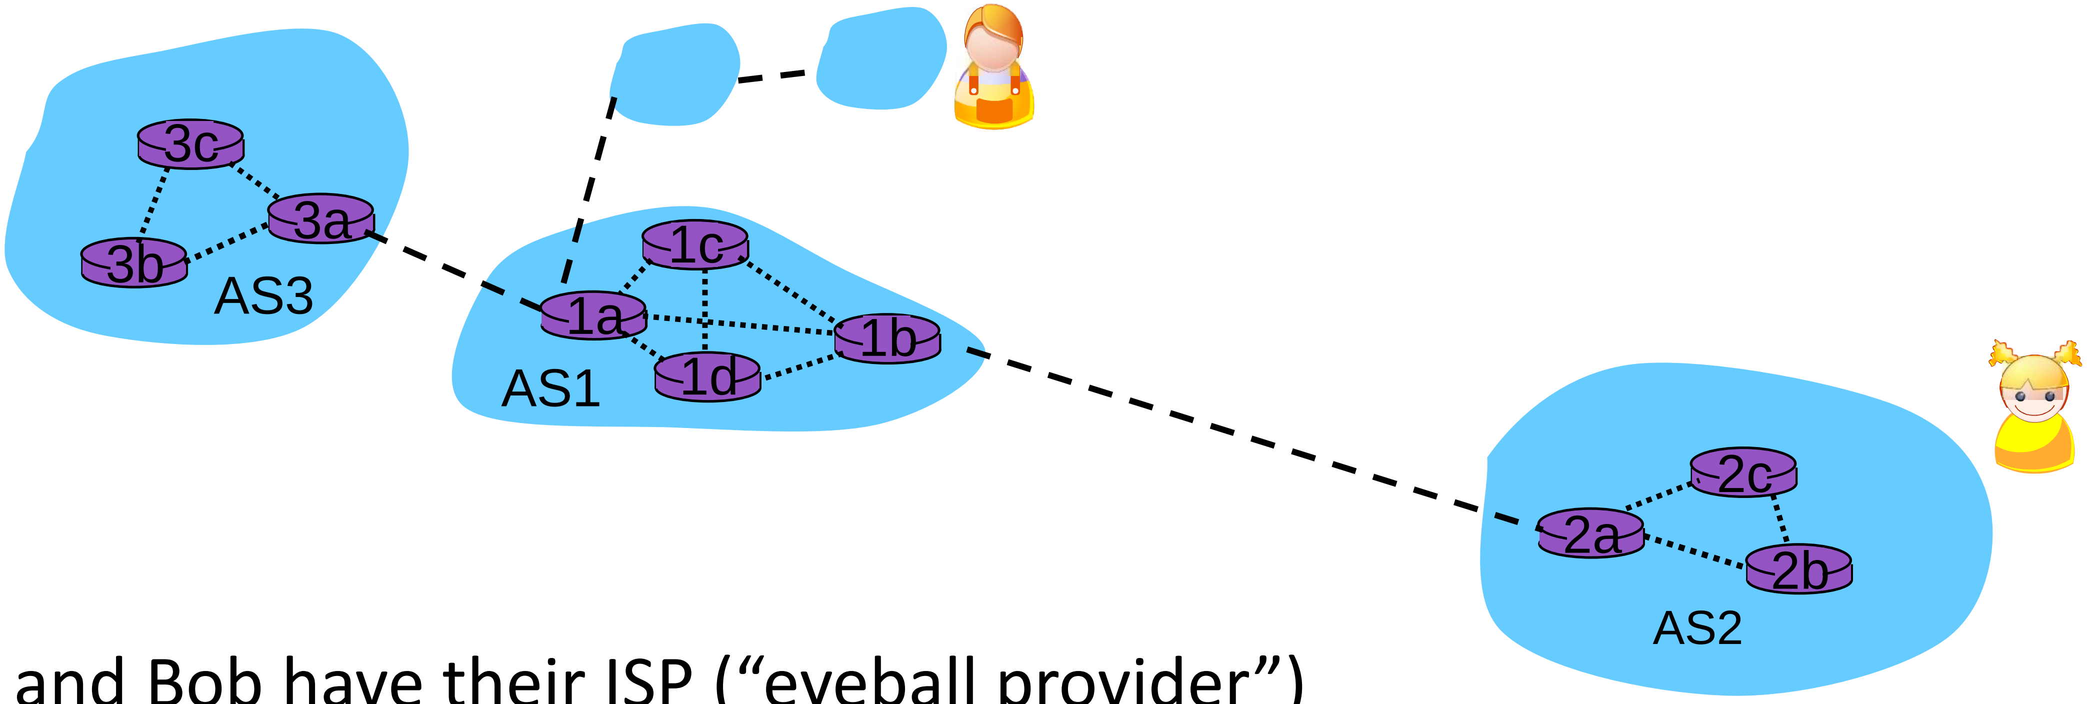

- Alice and Bob have their ISP (“eyeball provider”)
- Their ISPs are peering, or connected via transit providers
- How do their access routers know of each other?

Hierarchical Routing: Autonomous Systems

- Autonomous systems (AS)
  aggregate routers into regions,
- Routers in same AS run same routing protocol
  - "Intra-AS" routing protocol
  - Routers in different AS can run different intra-AS routing protocol

Gateway Routers

- Special routers in AS
- Run intra-AS routing protocol with all other routers in AS
- Also responsible for routing to destinations outside AS
  - Run *inter-AS routing* protocol with other gateway routers

Intra/Inter-AS routing

Karlsruhe Institute of Technology

- **Two-level routing:**
  - Intra-AS: administrator is responsible for choice
    - Intermediate System to Intermediate System (IS-IS): Link State
    - Open Shortest Path First (OSPF): Link State
    - Routing Information Protocol (RIP): Distance Vector
    - Interior Gateway Routing Protocol (IGRP): Distance Vector (Cisco proprietary)
  - Inter-AS: unique standard
    - Border Gateway Protocol (BGP): Path Vector (sort of distance vector, but with path information for loop avoidance)

Inter-AS and Intra-AS Routing

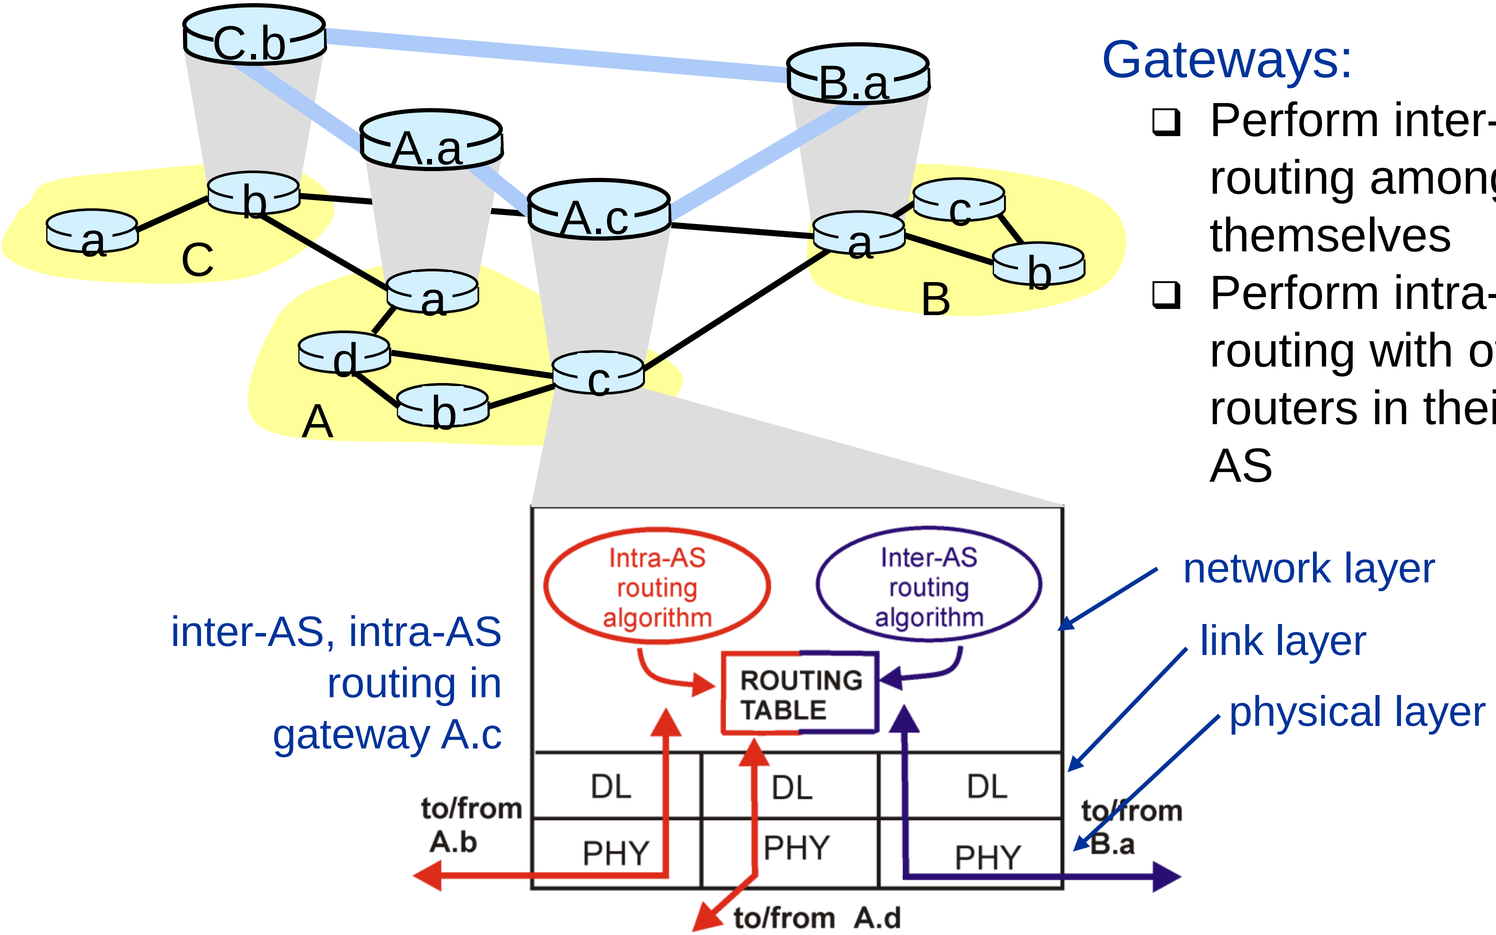

Gateways:

- Perform inter-AS routing amongst themselves
- Perform intra-AS routing with other routers in their AS

Why Different Inter- and Intra-Domain Routing?

- Scale:
  - Hierarchical routing saves table size, reduced update traffic

- Policy:
  - Inter-AS: admin wants control over how its traffic routed, who routes through its net.
  - Intra-AS: single admin, so no policy decisions needed

- Performance:
  - Intra-AS: can focus on performance
  - Inter-AS: policy may dominate over performance

Internet Inter-AS Routing: Border Gateway Protocol (1)

- Border Gateway Protocol (BGP) is the current de facto standard

- BGP is a path vector protocol:
  - Similar to distance vector protocol, but path info allows to avoid loops
  - Each border gateway broadcast to neighbors (peers) entire path (i.e, sequence of 16-bit AS identifiers) to destination
  - E.g., Gateway X may (or may not!) send its path to dest. Z:
- Path (X,Z) = X,Y1,Y2,Y3,…,Z

- Suppose gateway X send its path to peer gateway W:
  - W may or may not select path offered by X
  - Cost, policy (don't route via competitors AS), loop prevention reasons.
  - If W selects path advertised by X, then:
    - Path (W,Z) = w, Path (X,Z)
  - Note: X can control incoming traffic by controlling its route advertisements to peers:
    - e.g., don't want to route traffic to Z $\Rightarrow$ don't advertise any routes to Z

Internet Inter-AS Routing: Border Gateway Protocol (2)

- BGP messages are exchanged using TCP:
  - Simplifies BGP (no own error control / timeouts needed)
  - Routes from a peer are kept until withdrawn or TCP connection to that peer breaks $\Rightarrow$ allows for incremental updates

- BGP messages (non exhaustive list):
  - `OPEN`: opens TCP connection to peer and authenticates sender
  - `UPDATE`: advertises new path (or withdraws old)
    - Network Layer Reachability Information (NLRI): a length and a prefix per UPDATE, may contain several AS paths (route aggregation)
  - `KEEPALIVE`: keeps connection alive in absence of UPDATES; acknowledges OPEN request
  - `NOTIFICATION`: reports errors in previous msg; also used to close connection

Hierarchical Addressing: Route Aggregation

- Hierarchical addressing allows efficient advertisement of routing information:

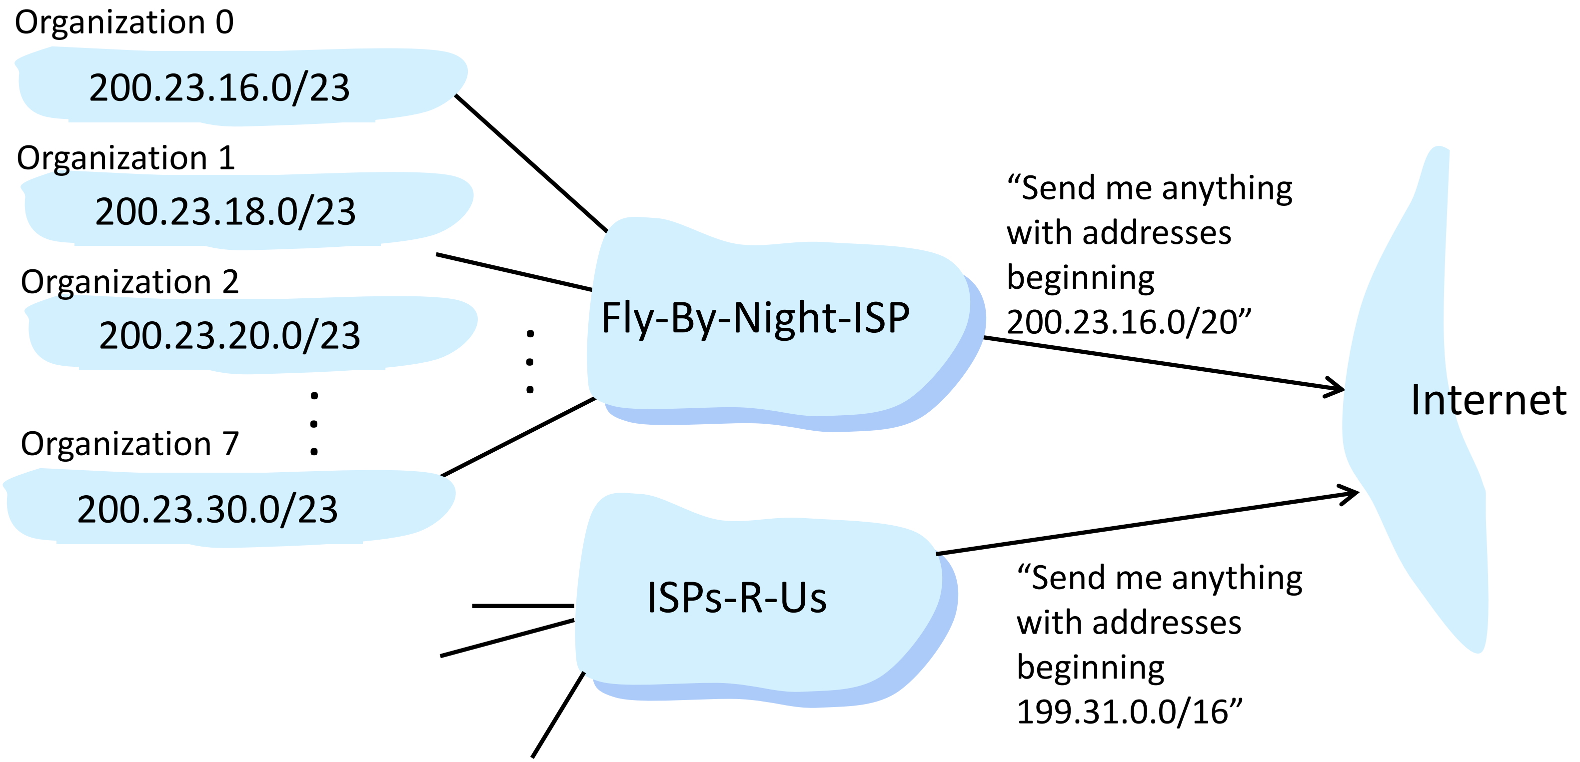

Hierarchical Addressing: More Specific Routes

- ISPs-R-Us has a more specific route to Organization 1:

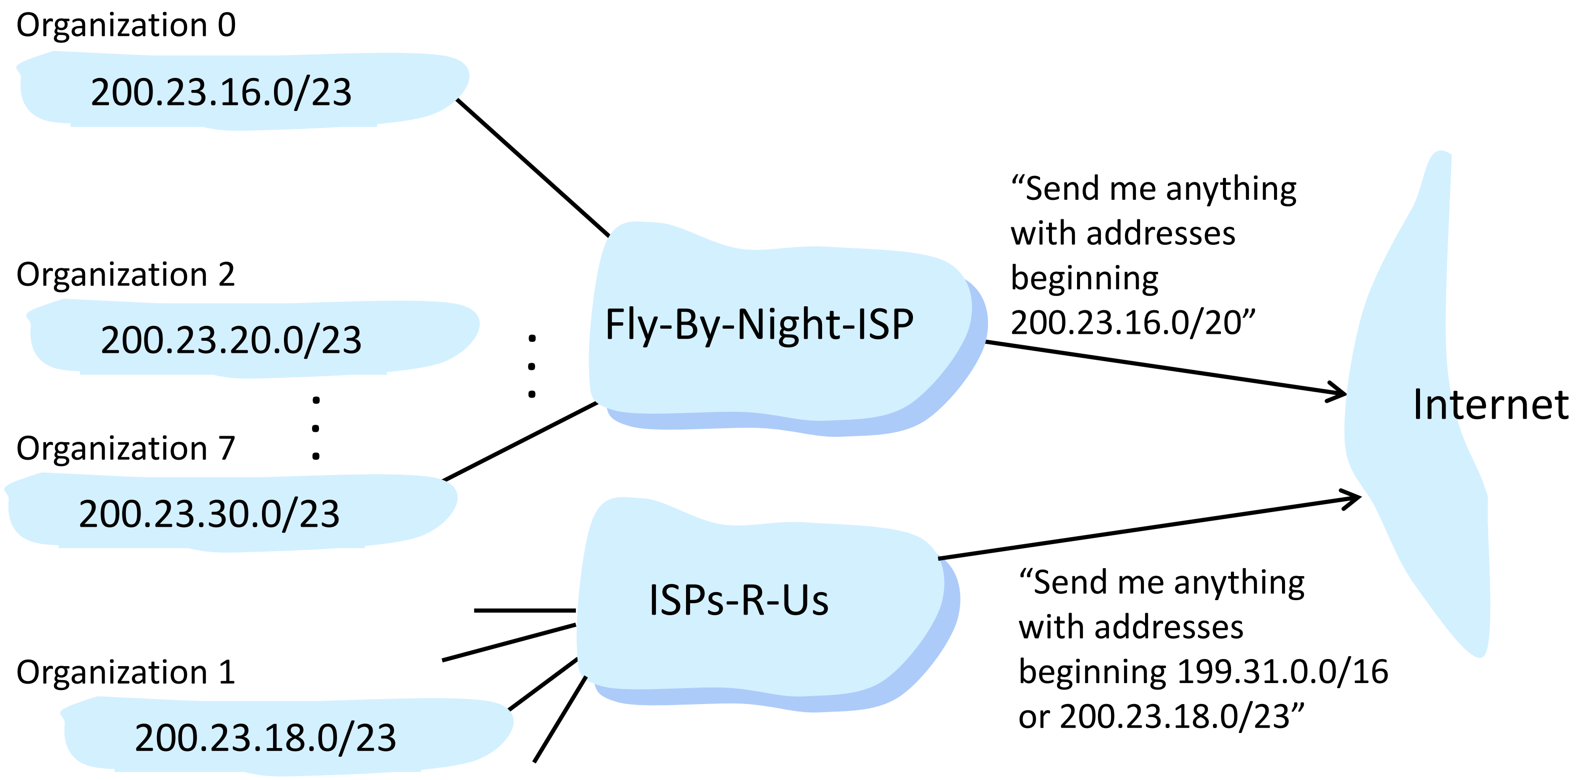

Inter-AS Routing Threats in the Internet

- Inter-AS routing threats mainly concern BGP operation

- Attack Scenarios:
  - Disabling parts of the Internet by disrupting Internet routing tables
  - Forcing multi-homed AS to use alternate paths to / from an outside AS instead of the preferred path
  - Disabling a single- or multi-homed AS
  - Creating traffic “blackholes”

- Such attack scenarios can e.g. be realized by:
  - announcing to “host” IP addresses ranges for which no ownership exists
  - inserting unauthorized “prefixes” into routing table (= announcing paths for networks for which no authorization to route exists)
  - modifying or forging routing messages during transmission
  - resource destruction

Examples from the real world…

- Pakistan Telekom “vs.” youtube.com (Feb 24th 2008)
  - PT (AS 17557) wanted to block traffic to youtube
  - Mistakenly advertised routes to 208.65.153.0/24 (AS 36561, youtube)
  - PCCW Global (upstream provider, AS 3491) forwarded announcement…

http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study

An example from the real world…

Examples from the real world…

- Pakistan Telekom “vs.” youtube.com
  - PT (AS 17557) wanted to block traffic to youtu
  - Mistakenly advertised routes to 208.65.153.0/
  - PCCW Global (upstream provider, AS 3491) for
- “Fat-finger error” or not, malicious intent e
  - Long term study by Symantec/QCRI (incl. Span
  - >2.5k IP blocks hijacked (~99% short-lived)
  - 64 abused for spamming:

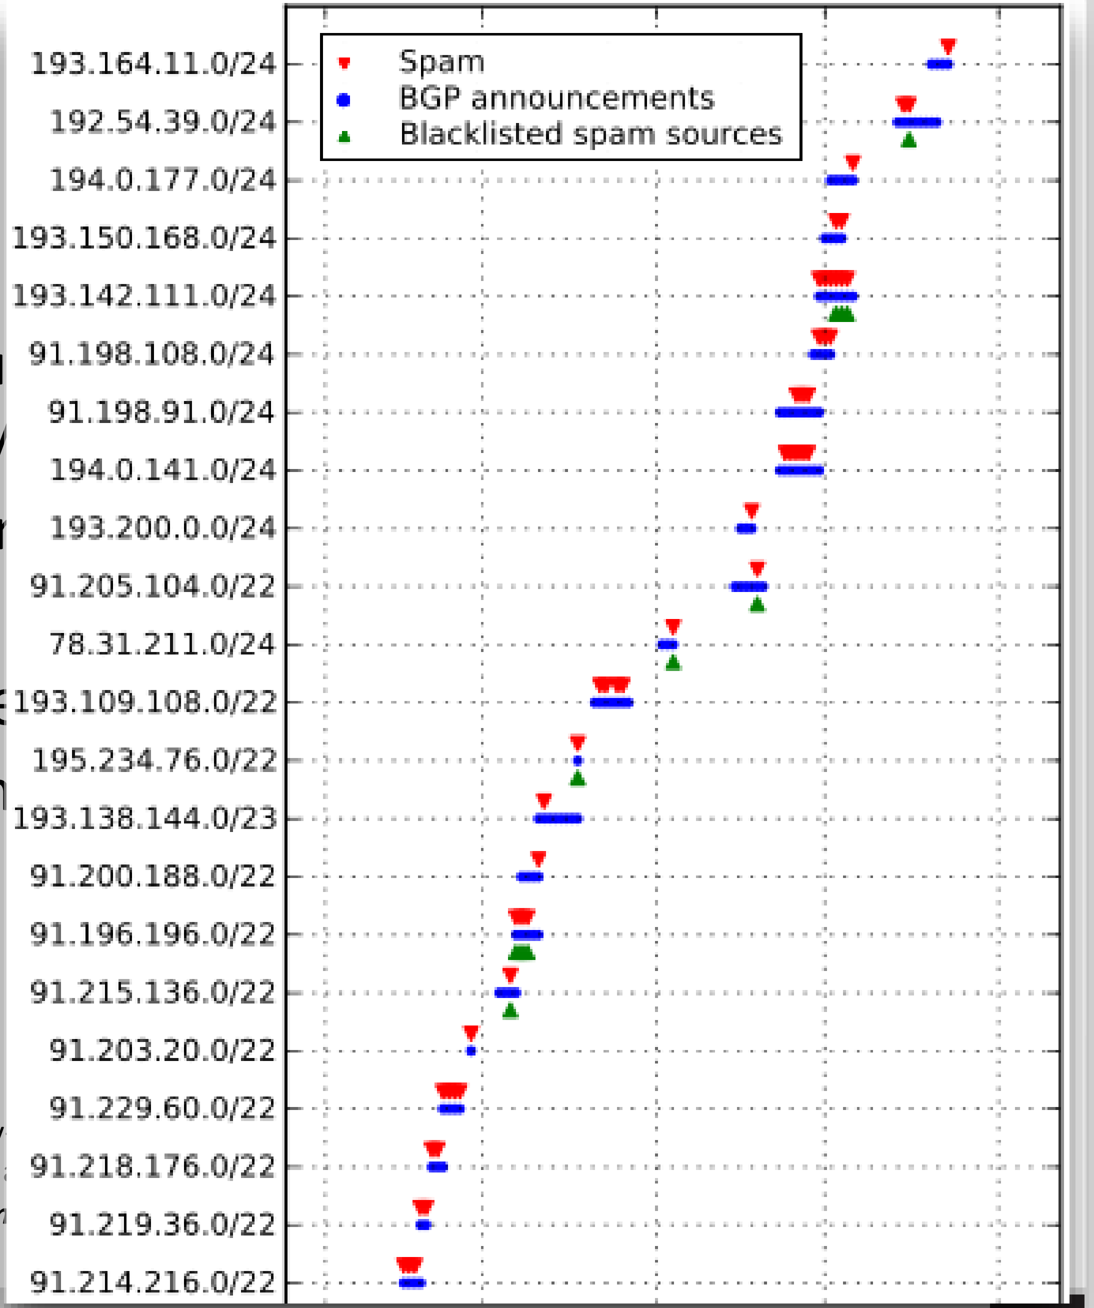

http://www.ripe.net/internet-coordination/news/industry
Vervier et al.: “Mind Your Blocks: On the Stealthiness of M
https://www.nanog.org/sites/default/files/monday_gen

Hijacking for Censorship [1]

- Local elections in Turkey, March 2014
- Goal: control „the Internet“ <- control DNS
- Start: hard NULL route
- Set up DNS servers and drop specific requests
- Google public DNS servers at 8.8.8.8 (and 8.8.4.4)
- …hijack 8.8.8.8/32

[1] c.f: https://www.nanog.org/sites/default/files/monday_general_bgp_toonk_63.18.pdf

Hijacking for Censorship [1]

Technology

- Local elections in Turkey, March 20
- Goal: control „the Internet“ <- cont

Youtube.com lookup at Google's 8.8.8.8 DNS server 8.8.8.8 from Turk Telekom

```
;; ANSWER SECTION (1 record)
youtube.com.     86064    IN       A     195.175.254.2
                                         ^^^^^^^^^^^^^^^^
                                  Not a real Youtube IP address
```

Youtube.com lookup at Google's 8.8.8.8 DNS server from The Netherlands

```
;; ANSWER SECTION:
youtube.com.        299    IN A   74.125.136.93
youtube.com.        299    IN A   74.125.136.91
youtube.com.        299    IN A   74.125.136.136
youtube.com.        299    IN A   74.125.136.190
                                ^^^^^^^^^^^^^^^^^
                          Normal Youtube IP addresses
```

- f
- .8 (and 8.8.4.4)
- w.nanog.org/sites/default/files/monday_general_bgp_toonk_63.18.pdf

Examples – Traffic Diversion/MitM

- Early 2013 a Belorussian provider attracted traffic from GlobalOneBel over an uplink to Moscow

- Attacked networks changed daily but continued for a month
- Attracted traffic was forwarded to an unaffected uplink to Frankfurt

- **Difficult to detect**: hosts in US

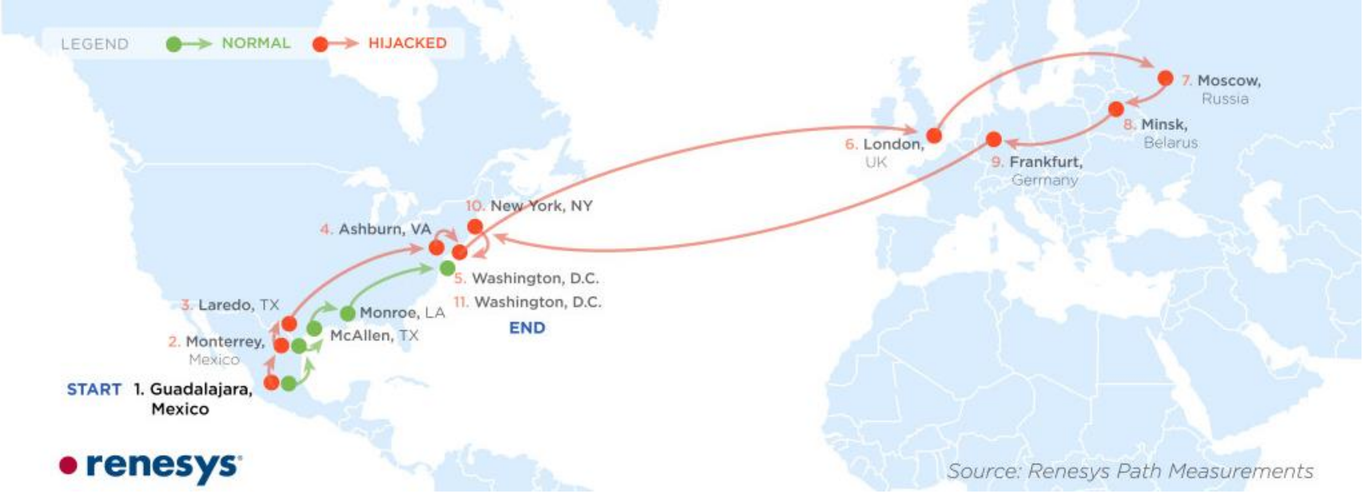

Details: http://www.renesys.com/2013/11/mitm-internet-hijacking

Where do Attacks take Place?

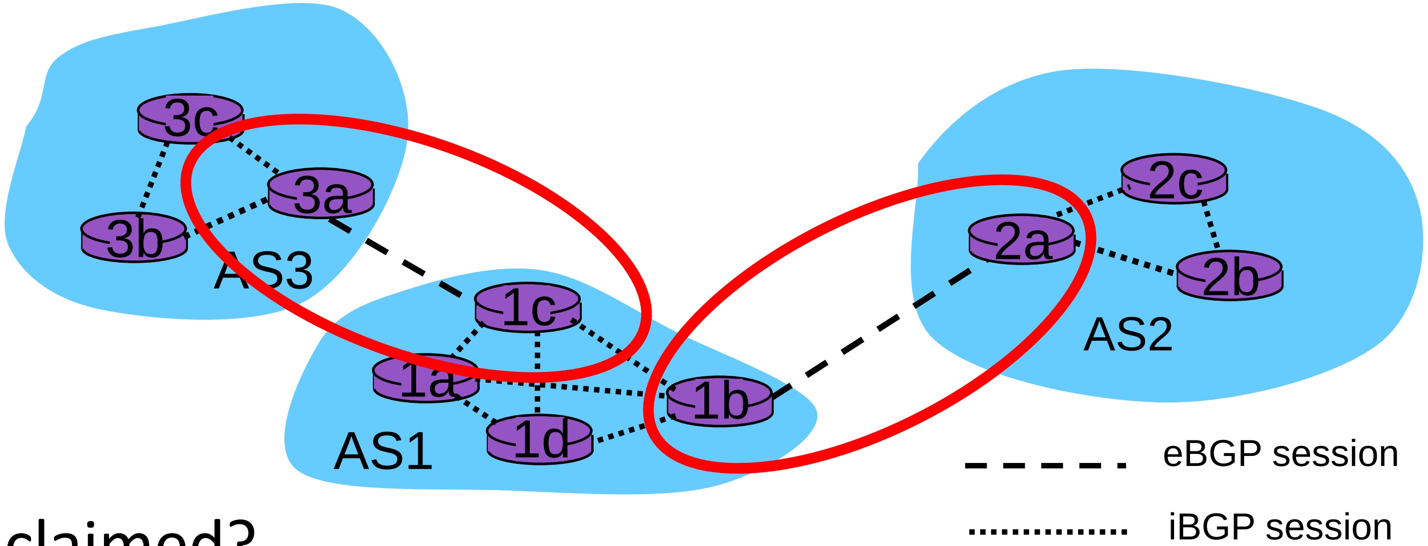

- What is (mis)claimed?
  - Identity
  - Reachability
  - Address ownership

Where are the Attackers?

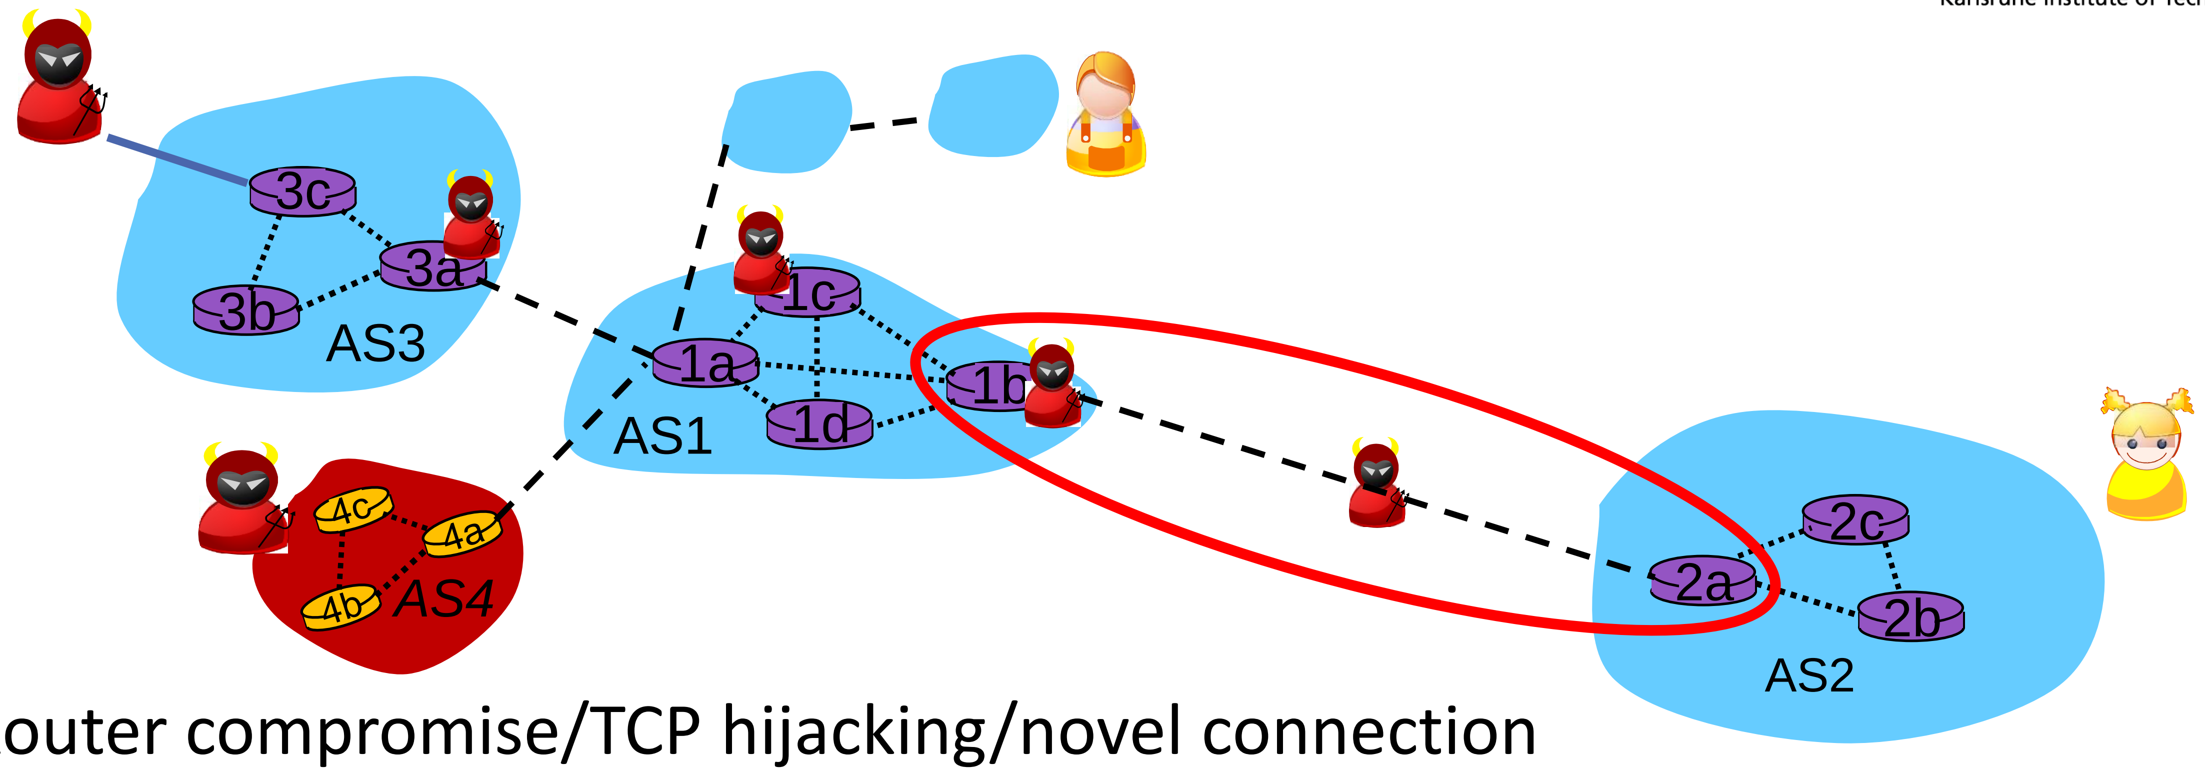

- Router compromise/TCP hijacking/novel connection
- On-path/Off-path
- BGP speakers/end-users?

Securing BGP Operation: Verifying Peer Messages (1)

- Assuming external adversaries:

- Force routers to accept only protocol messages from directly connected peers (if direct links exist):
  - Referred to as BGP TTL Security Hack (BTSH)
  - Idea – directly connected peer routers:
    - send routing messages with IP TTL field set to 255, and
    - accept only routing messages with IP TTL field ≥ 254
  - Messages from attackers which can only reach a target router over multiple hops will be discarded by router

  - ***Why can this mechanism not be implemented as follows?***
    - send routing messages with IP TTL set to 1, and
    - let routers in between automatically discard routing messages after one hop (so routing messages from attacker will not reach the target)

Securing BGP Operation: Verifying Peer Messages (2)

- More general approach:
  - Generalized TTL Security Mechanism
  - Standardized for IPv4 and IPv6 in RFC 5082
  - Routers set TTL=255, but may be multiple hops away
  - Packets are accepted depending on the distance, e.g., with TTL=253 when the router is two hops away
  - More configuration overhead, may be less secure than BTSH
- Better: authenticate routing messages between peers
  - What do you need to authenticate?
  - Signature (Hash) & PKI/CA (IPSec)

Securing BGP Operation: Verifying Peer Messages (3)

- TCP MD5 Signature Option (RFC 2385):
    - Goal: protect BGP exchanges between peers from spoofed TCP segments (attacker who eavesdrop/“guess” correct sequence number)
    - Sender computes an MD5 hash value over each TCP segment and a secret shared with its peer entity
    - The hash value is transported in an option field
    - All options in TCP PDU must not exceed 40 bytes: use 16 Byte long MD5 hash values

    (plus two bytes for TCP option information; type and length)

    - Problem: MD5 is not state of the art, no automatic key negotiation / update procedure defined → deployment difficulties

    (+ known vulnerabilities of manual key mgmt.)

Securing BGP Operation: Verifying Peer Messages (4)

- TCP Authentication Option (RFC 5925):
  - Successor to TCP MD5 Signature with different cryptographic algorithms
  - Better replay protection (even when TCP seq. numbers roll over)
  - Not (yet) widely deployed
  - Still no automatic key negotiation / update procedure defined
- Deployment of IPsec between peers:
  - Provides authentication and replay protection for IP packets
  - Allows for additional confidentiality
  - Leverages key management protocol that may use certificates and private keys
  - Potential problem: Low convergence speed when a router has many peers, e.g. > 1000, as key exchanges may take seconds per neighbor
- Sometimes routers may still be contacted from outside WITHOUT any authentication!

Taking one step back

- …but what is really the core of the problem?

- Attacks are intentionally broken assumptions ->
- What does the adversary lie (uhm, or: err) about?

- Sender and/or content (BGP operations/parameters) of the message:

  - Identity
  - IP address (AS) ownership
  - Reachability information

- How could we solve this (using crypto)?
- Certification of resource ownership and path…

- However, fixed attestations render solutions inflexible…

Problems Beyond Simple Peer-to-Peer BGP Security

Karlsruhe Institute of Technology

- Address space “ownership” verification:
  - Who has been assigned an IP address range and has thus the right to announce this range / delegate the announcement of this range?
- Autonomous System (AS) authentication:
  - To whom has a claimed AS-number actually been assigned?
- Router authentication and authorization (relative to an AS):
  - Are the entities pretending to belong to an autonomous system authentic?
- Route and address advertisement authorization:
  - Who is allowed to announce specific address ranges / routes
- Route withdrawal authorization:
  - Who is allowed to withdraw a route?

→ Need for further security measures, approaches are S-BGP/SIDR

Core Idea of S-BGP/BGPsec

Karlsruhe Institute of Technology

- The validity of a BGP UPDATE is based on four primary criteria:

1. The ***router*** that sent the UPDATE was authorized to act ***on behalf of the AS*** it claims to represent; that is, the AS at the front of the AS path.
2. The ***first AS*** in the AS path was ***authorized***, by the owner of the set of prefixes that are represented in the UPDATE, ***to advertise*** those ***prefixes***.
3. The ***AS*** from which the UPDATE emanates was ***authorized by the preceding AS*** in the AS path (in the UPDATE message) to advertise the prefixes in the UPDATE.
4. If the UPDATE ***withdraws*** one or more routes (specified by the prefixes for the routes), then the ***sender*** must have ***advertised*** each route ***prior to withdrawing*** it.

S-BGP/BGPsec: Combining Standards

Karlsruhe Institute of Technology

- **Address Attestations**:
  - Authorization of subject (by issuer) to advertise specified prefixes/address blocks
- Validation of BGP UPDATEs:
  - New path attribute, using certificates and attestations, to prove authorizations
- Distribution of security specific data:
  - Certificates, certificate revocation lists (CRLs)
- Public Key Infrastructures (PKIs):
  - Secure identification of BGP speakers and of owners of AS's and of address blocks
- IPsec:
  - Provides authentication and integrity of peer-to-peer communication with support for automated key management

https://www.isoc.org/isoc/conferences/ndss/99/proceedings/slides/lynn.pdf

Internet Address Space Ownership

Karlsruhe Institute of Technology

- Internet address space managed hierarchically with the Internet Assigned Numbers Authority (IANA) as root authority for assigning address ranges

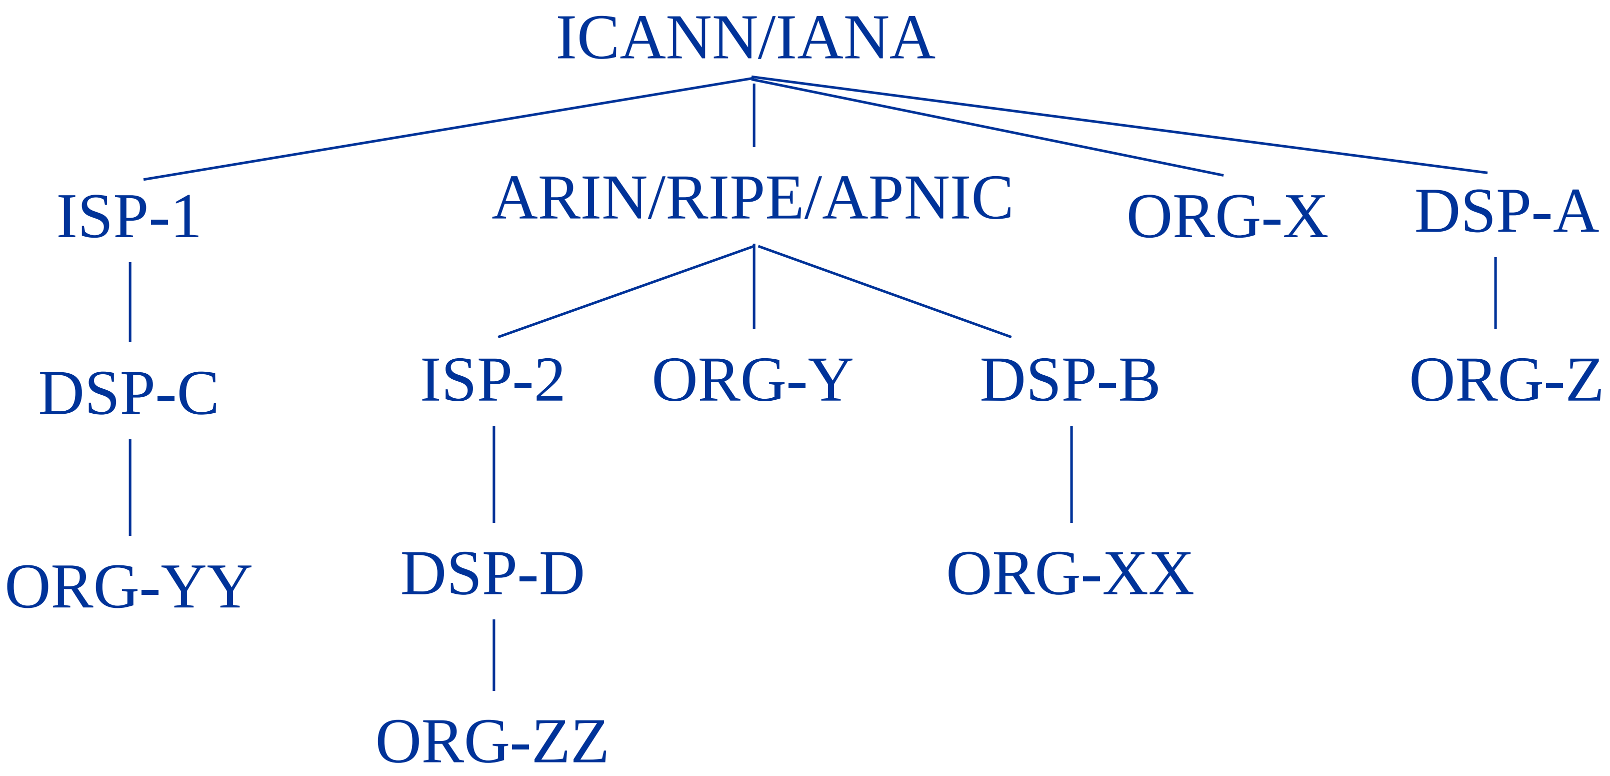

KASTEL

S-BGP: Certificates and Address Space Attestations

- ***ICANN*** issues certificates for ***address space ownership to regional authorities*** and to entities that have direct address allocations (from IANA)
- Each certificate contains extension specifying the ***address space being delegated***, so that certificate validation is address-constrained
- Holders of address space certificates can create an ***address attestation***, authorizing an AS (or a router) to advertise the specified address space

S-BGP: Address Certificates

	Issuer	Subject	Extensions
Root Certificate	IANA	IANA	all addr
Registry Certificate	IANA	Registry	addr blocks
ISP/DSP Certificate	Registry (or IANA)	ISP/DSP	addr blocks
Subscriber Certificate	ISP/DSP (or Registry, IANA)	Subscriber	addr blocks

S-BGP:
AS Ownership/Router Identification

- ICANN issues certificates for AS ownership to:
    - ISPs, DSPs, and organizations that run BGP
- AS operators issue certificates to:
    - Routers as AS representatives

S-BGP: AS and Router Certificates

	Issuer	Subject	Extensions
Root Certificate	IANA	IANA	all ASes
Registry Certificate	IANA	Registry	ASes
AS Owner Certificate	Registry (or IANA)	ISP/DSP or Subscriber	ASes
AS Certificate	ISP/DSP or Subscriber	AS	
Router Certificate	ISP/DSP or Subscriber	Router*	AS, RtrId

* the subject name could be a fully-qualified DNS name

S-BGP:
AS # Allocation and Router PKI Example

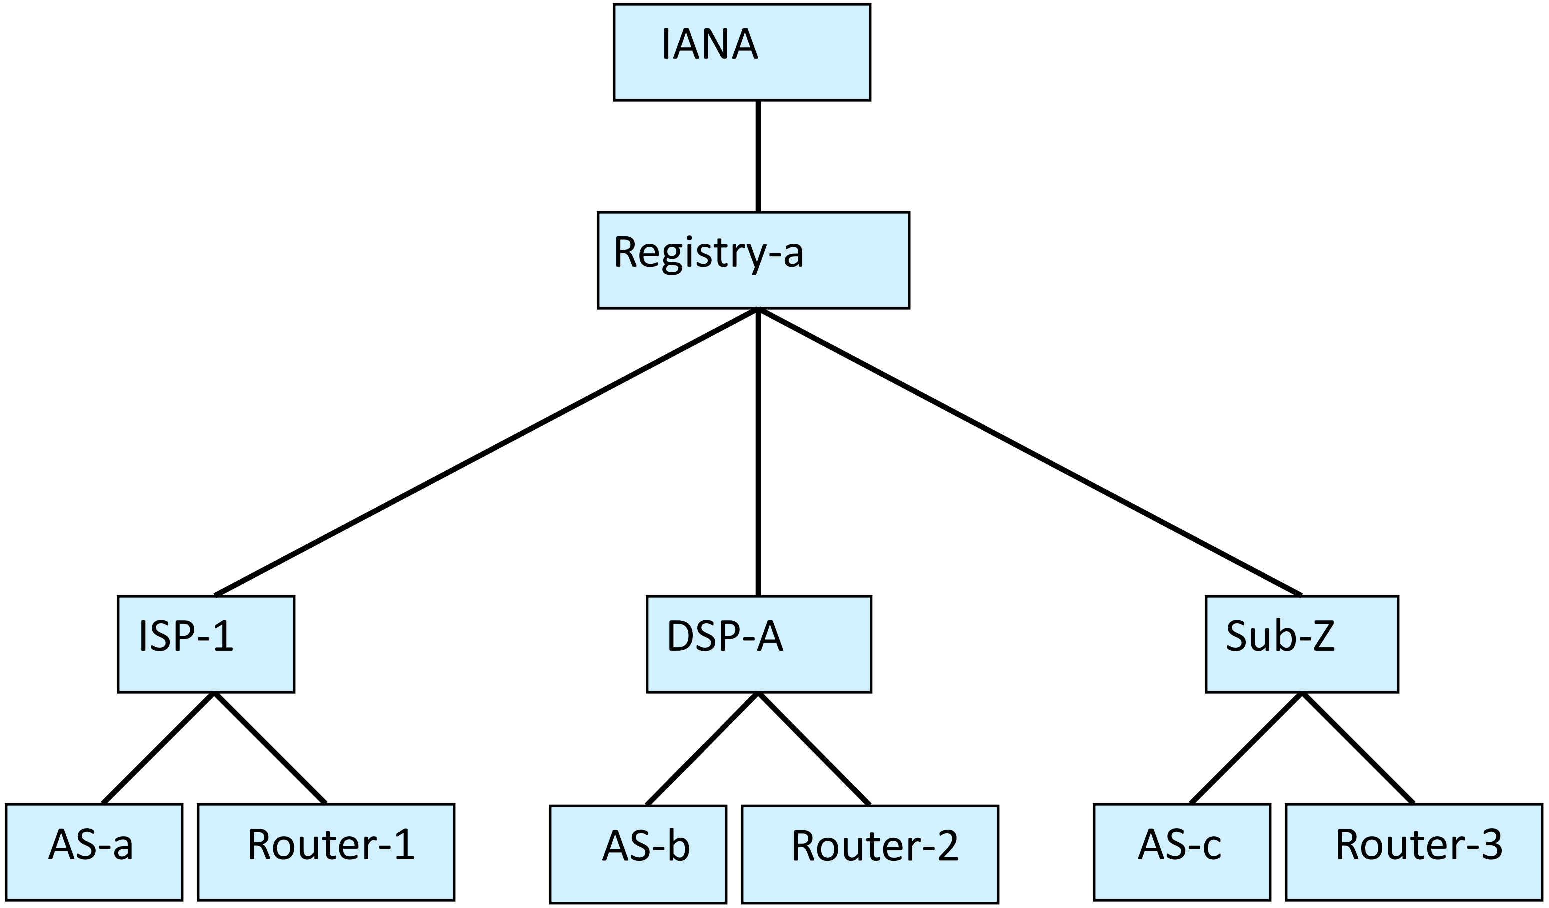

S-BGP: Overview of Attestations

- Holders of AS (or router) certificates generate route attestations that:
  - authorize the advertisement of a route by a specified next hop AS
  - are used to express a secure route as a sequence of AS hops
- Address Attestations:
  - Used to validate that a destination address block is being originated by an authorized AS
- Route Attestations:
  - Used to validate that an AS is authorized to advertise an AS Path
- Each UPDATE includes optional transitive path attribute ATTEST with:
  - one or more Address Attestations, and
  - a set of Route Attestations

S-BGP: Address Attestations

Karlsruhe Institute of Technology

- Address Attestations include identification of:
    - address blocks,
    - their owner's certificate,
    - AS authorized to originate (advertise) the address blocks, and
    - expiration date/time

- Indicate that the AS listed in the attestation is authorized by the owner to originate/advertise the address blocks in an UPDATE

- Digitally signed by owner of the address blocks, traceable up to the IANA via a certification path

- Used to protect BGP from erroneous UPDATEs (authenticated but misbehaving or misconfigured BGP speakers)

S-BGP: Route Attestations

Karlsruhe Institute of Technology

- Include identification of:
    - AS's or BGP speaker's certificate issued by the AS owner,
    - the address blocks and the AS Path (ASes) in the UPDATE,
    - the AS number of the receiving (next) neighbor, and
    - expiration date/time
- Indicate that the BGP speaker or its AS authorizes the receiver's AS to use the AS Path & NLRI in the UPDATE
- Digitally signed by owner of the BGP speaker (or its AS) distributing the UPDATE, traceable to IANA ...
- Used to protect BGP from erroneous UPDATEs (authenticated but misbehaving or misconfigured BGP speakers)

S-BGP/BGPsec: Securing UPDATEs

- Securing an UPDATE:
  - A secure UPDATE consists of an UPDATE message with a new, optional, transitive path attribute for route authorization
  - This attribute consists of a signed sequence of route attestations, nominally terminating in an address space attestation
  - This attribute is structured to support both route aggregation and AS sets
  - Validation of the attribute verifies that the route was authorized by each AS along the path and by the ultimate address space owner

S-BGP/BGPsec: Propagation of an S-BGP UPDATE

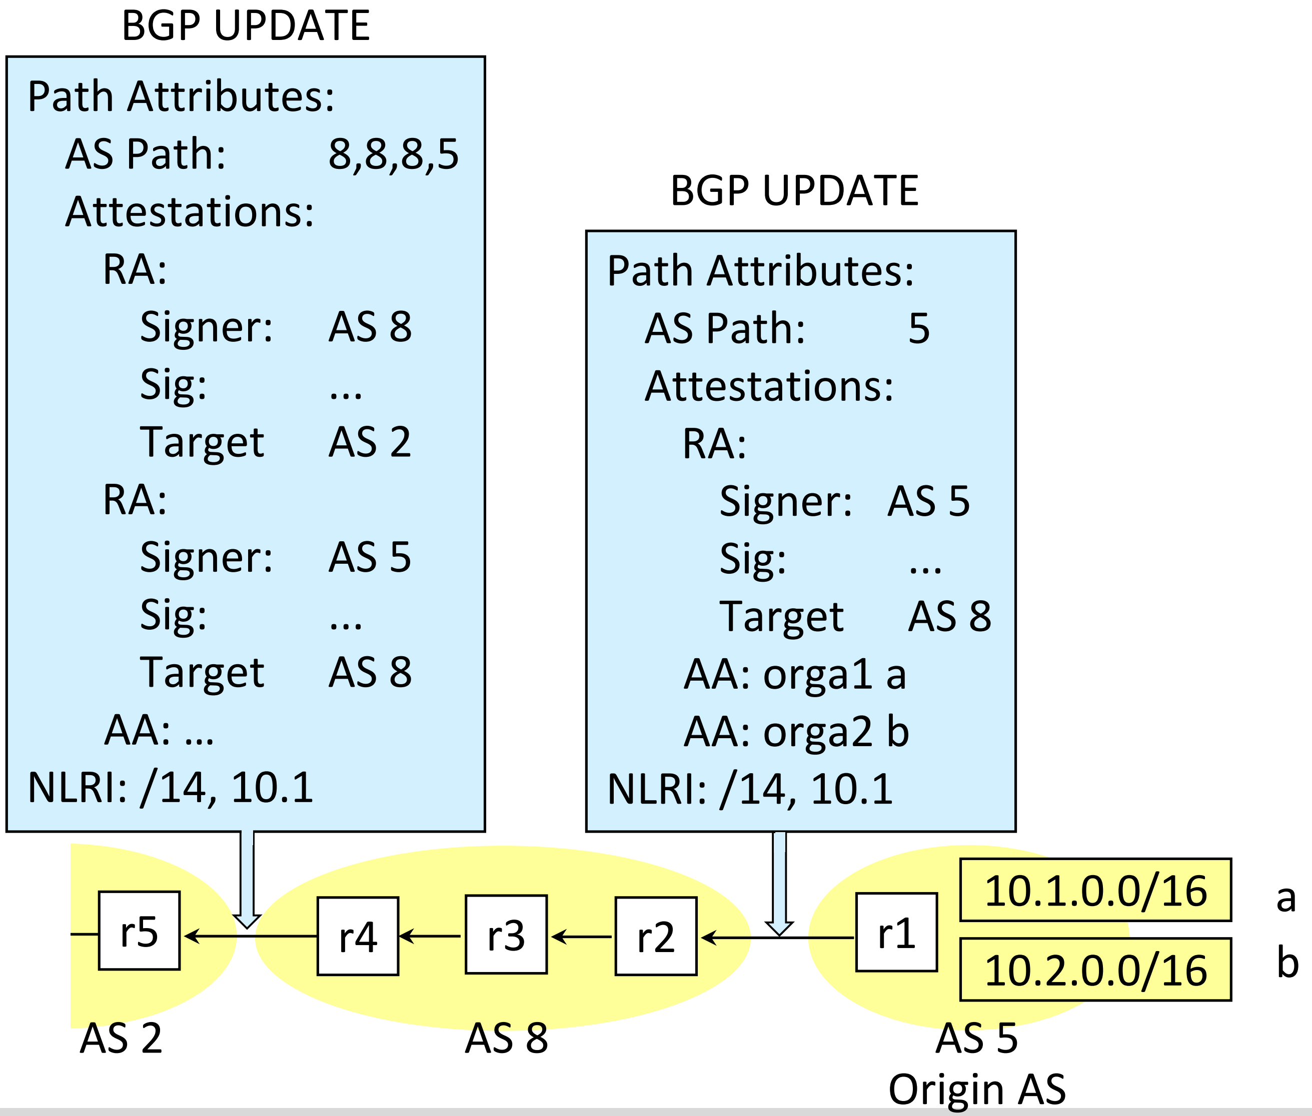

S-BGP/BGPsec: Propagation of an S-BGP UPDATE

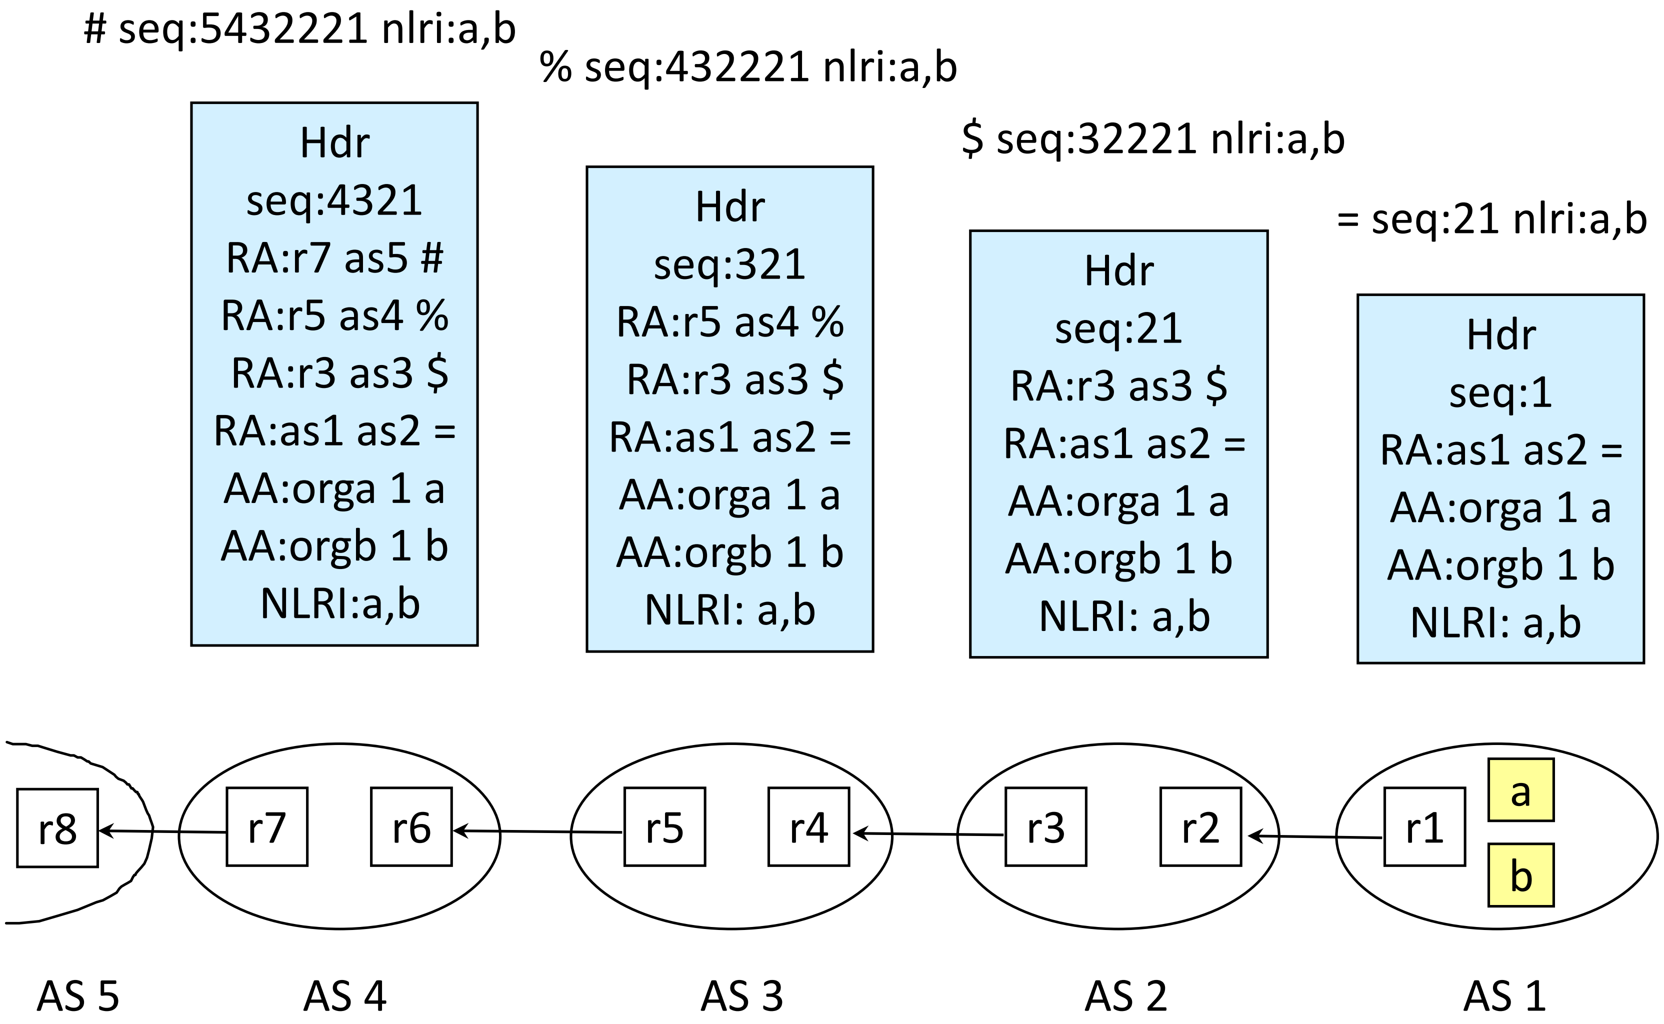

S-BGP: Encoding of Attestations

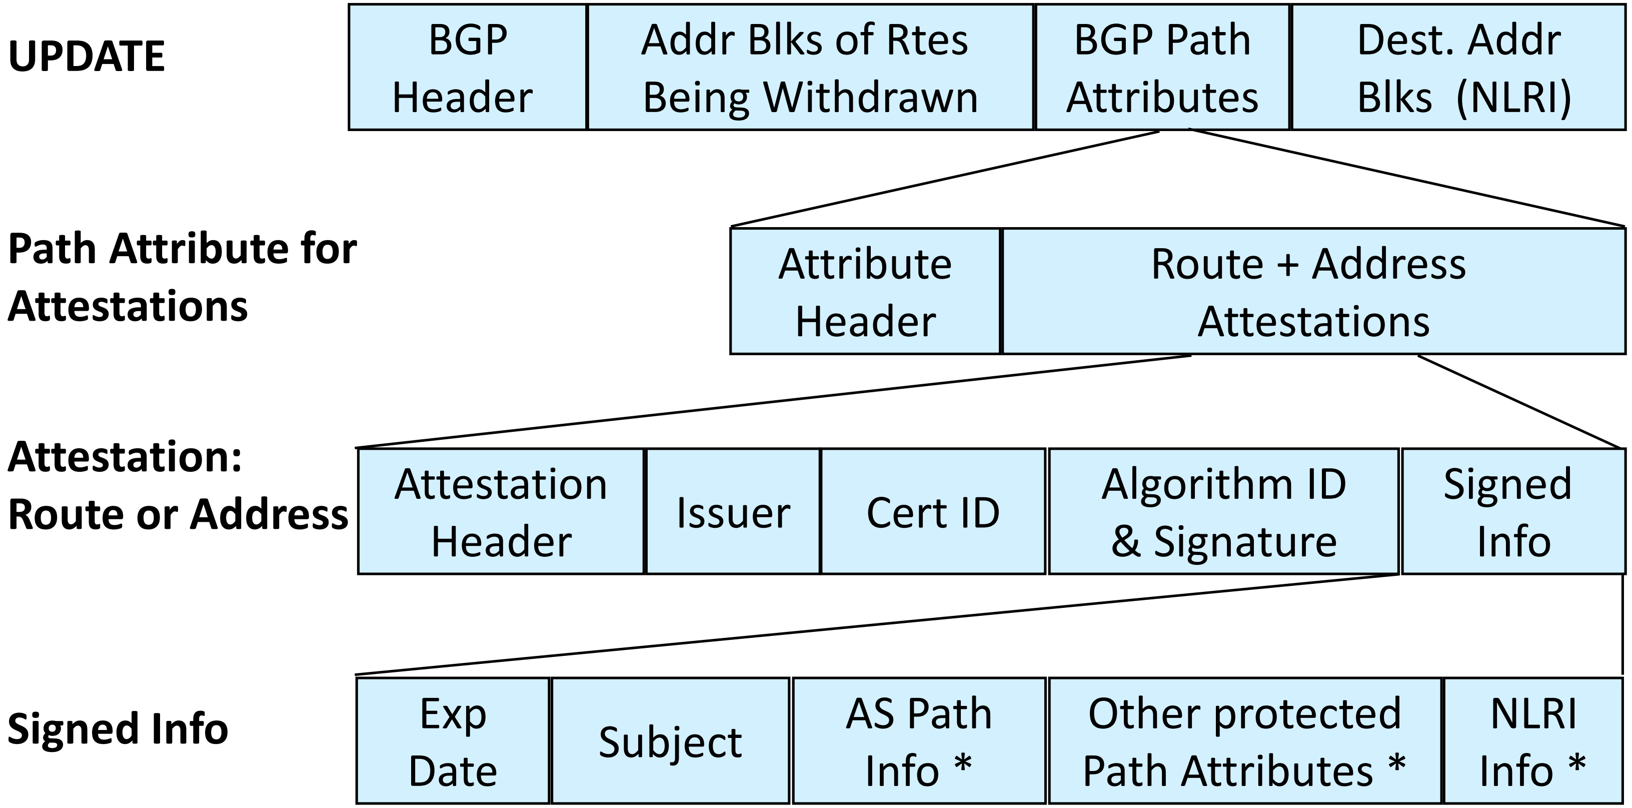

*explicit in the aggregation case, or if Path Attribute changes unpredictably

S-BGP:
Distributing Certificates, CRLs, & AAs

- Putting certificates & CRLs in UPDATEs:
  - would be redundant and make UPDATEs too big
  - same is true for address attestations

- Solution – use servers for these data items:
  - replicate for redundancy & scalability
  - locate at network access points (NAPs = points where multiple BGP speakers are interconnected with high speed LANs) for direct (non-routed) access

- Download options:
  - whole certificate/AA/CRL databases
  - queries for specific certificates/AAs/CRLs

- To minimize processing & storage overhead, network operations centers (NOCs) should validate certificates & AAs, and send processed extracts to routers
  - However, in this case trust is delegated to the NOC!

S-BGP: Validating a Route

- To validate a route from ASn, ASn+1 needs:
  - 1 address attestation from each organization owning an address block(s) in the network layer reachability information (NLRI),
  - 1 address allocation certificate from each organization owning address blocks in the NLRI,
  - 1 route attestation from every AS along the path (AS1 to ASn), where the route attestation for ASk specifies the NLRI and the AS Path up to that point (AS1 through ASk+1),
  - 1 certificate for each AS or router along the path (AS1 to ASn) to use to check signatures on the route attestations, and
  - of course, all the relevant certificate revocation lists (CRLs) must have been verified (in case a private key was compromised and the corresponding certificate must be revoked)

S-BGP: Performance Issues – Resources

- Certificates (generation and signing done offline):
  - Disk space for storing certificates
  - CPU resources for validating certificates
- CRLs (generation and signing done offline):
  - Disk space for storing CRLs
  - CPU resources for validating CRLs
- Attestations:
  - Routing Information Base (RIB) memory space for storing attestations
  - Disk space for faster recovery from router reboot (optional)
  - CPU resources for signing and validating attestations
  - Resources for transmitting attestations (to make this a dynamic system)
- Size of the problem (May 2017):
  - ~ 57,546 AS, ~ 670,590 owners of address prefixes
  - Resulting certificate database size: > 330 Mbyte (~ 450 byte / certificate)
  - CRLs would add to this (should not be too much)

Further S-BGP Issues

Karlsruhe Institute of Technology

- Remaining vulnerabilities:
  - Failure to advertise route withdrawal
  - Premature re-advertisement of withdrawn routes
  - Erroneous application of local policy
  -
- (Non-)Deployment:
  - The ideas of S-BGP are slowly seeping into reality through BGPsec – SIDR/RPKI
  - S-BGP mainly
    - shows the tasks to be accomplished regarding certification of IP address ownership, AS# ownership, and authorization to advertise certain routes
    - gives an impression on the scale of the effort that has to be invested in order to secure a global-scale Inter-AS routing protocol

SIDR: BGPsec & RPKI

- Secure Inter-Domain Routing
- Standardization for approach based on simplified S-BGP principles

- Several changes
  - Split in **BGPsec** (***routing attestations*** and mappings) RFC 8205 – 8211 and
  - Resource Public Key Infrastructure (**RPKI**) RFC 6810, a directory for ***Secure Origin Authentication***
  - Certificate information is replicated among distributed servers
  - Signatures are distributed by BGP UPDATES
    - Allows for BGPsec negotiation between routers
    - Support of "BGPsec islands"
  - Several optimizations with regard to efficiency

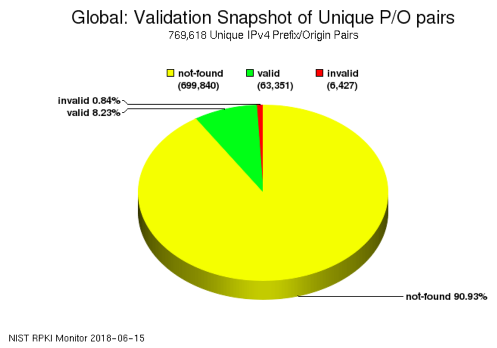

- Current status: RPKI is rolled out, but not widely used (yet)
  - See http://rpki.surfnet.nl/global.html or https://rpki-monitor.antd.nist.gov/

- No IANA root certificate

Securing BGP by state observation

- Drawbacks of seen cryptographic approaches
  - Computation and communication intensive
  - Usually public-key infrastructures or central databases needed
  - Incremental deployment with somewhat limited security gain

- Idea: Use available information to check credibility of BGP Update messages

- Interesting approaches:
  - Pretty Good BGP: Cautiously Adopting Routes
  - Topology-based Analysis

Pretty Good BGP: Cautiously Adopting Routes (1)

- Observation: Almost half of bogus origin/prefix associations last less than 24 hours

- Idea: Treat unfamiliar routes cautiously
  - Time for a secondary validation process
    (manual, Internet Alert Registry, or by others)
  - Exploits natural redundancy, as other older routes still exist

- First step: identifying normal routes
  - Routers store history of known origin/prefix pairs for h days (history period)
  - Database defines normal behavior

- Second step: detect anomalous routes
  - Received route updates compared with database
  - Updates altering the normal state
  - Marked suspicious for s days (suspicious period)
  - After s days, suspicious routes added to the history

Pretty Good BGP: Cautiously Adopting Routes (2)

Karlsruhe Institute of Technology

- Third step: avoiding suspicious routes
  - Suspicious routes get lowest possible preference
  - Routers select best trusted route (if possible)
  - False positives possible (less desirable route)
    - However, routing operates normally

- Drawbacks of approach: If new subprefixes are introduced (or generated by an attacker)
  - Routers will use known route to the larger address block during suspicious period
  - Leads to false positives: Potentially better path to new (valid) subprefix not used during suspicious period

- **All attacks** persisting longer than suspicious period are **successful**, as new routes are not tested.

Topology-based Analysis (1) [KMR03]

- Observation: Internet exhibits certain structure
  - Densely connected core nodes (backbone)
  - Periphery nodes with connection to the core and at most a few direct neighbors

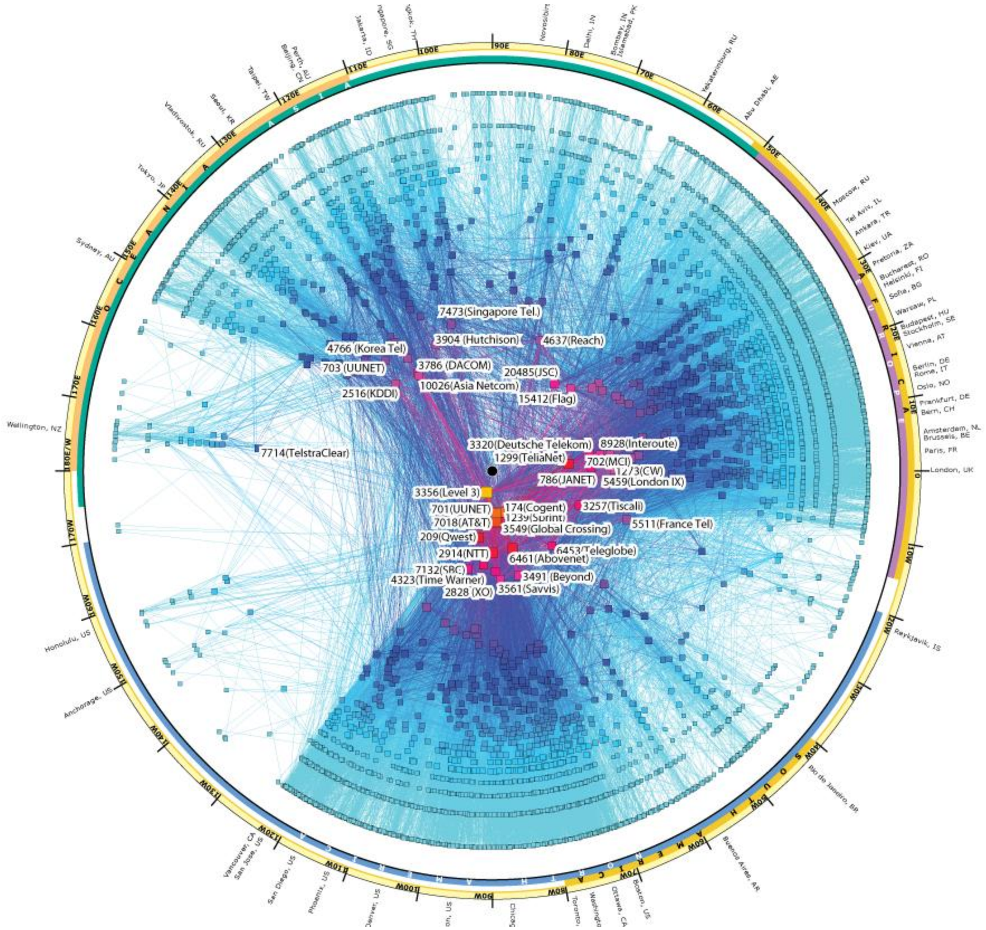

- Connectivity graph
  - Routers are nodes, direct links are edges
  - Can be approximated with information from route updates (combine several routers)

- Yellow and Red AS have many links (up to 1845), Blue AS have few links to other AS

- Attacks commonly modify or truncate path through backbone

Topology-based Analysis (2)

- Remove core nodes from Graph
  - Clusters of periphery nodes

- Routers with access to geographical data can determine the diameter of a cluster
  - Maximum geographical distance between two systems within a cluster
  - Diameter of most clusters is small (local networks connected to large providers)
  - Kruegel at al. use preprocessed information from the whois databases to determine geographical positions
  - Example excerpt of a whois record:

```
inetnum:       141.30.0.0 - 141.30.255.255
netname:       TUDR
descr:         Technische Universitaet Dresden
address:       Helmholtzstr. 10
address:       01069 Dresden
country:       DE
admin-c:       WW20
```

Topology-based Analysis (3) – Path modification attacks

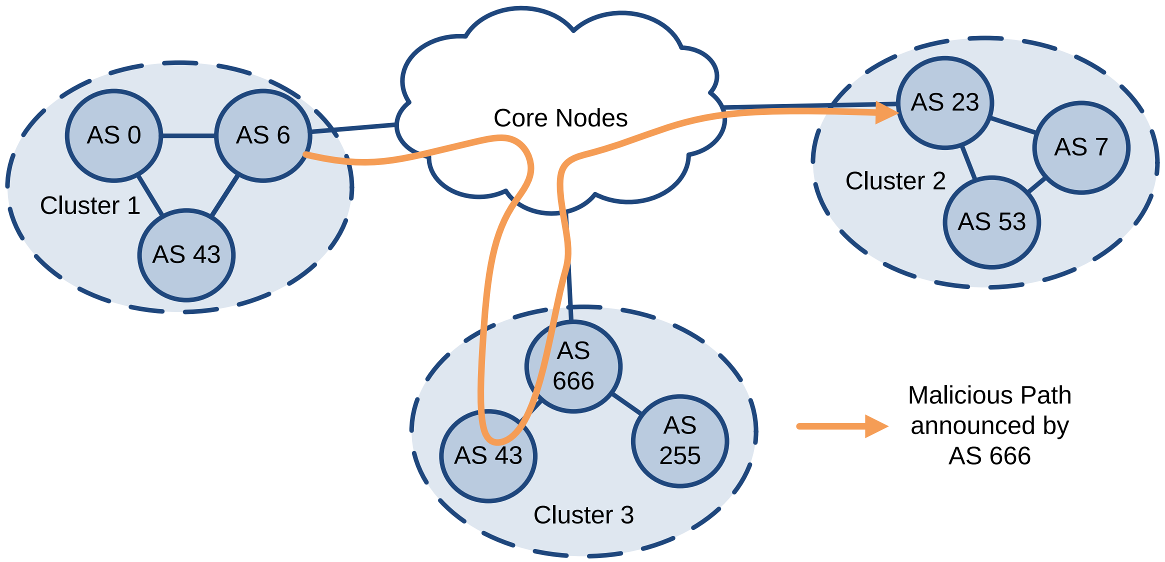

Valid routes must satisfy constraints:

- A valid route has only one single subsequence of core nodes
  - ➡ Identify "*path modification attacks*"
  - In the example the sequence goes through core nodes before AS 43 and after AS 666, hence considered invalid

Topology-based Analysis (3) – Path truncation attacks

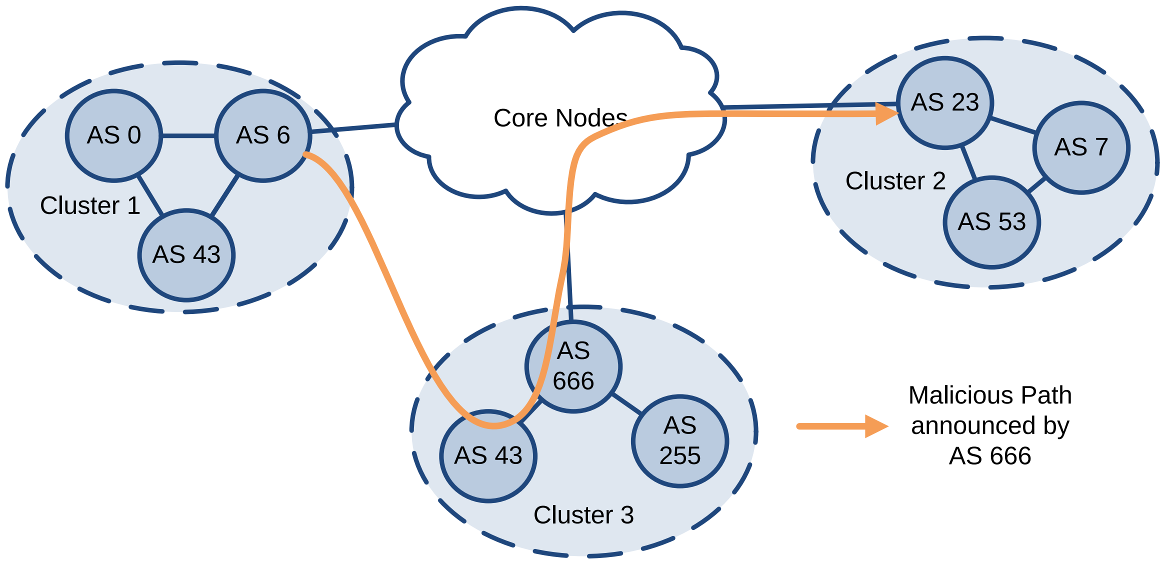

Valid routes must satisfy constraints:

- All consecutive pairs of periphery nodes in a route must be in a cluster or close geographical range (a 300km threshold proposed for the Internet)
  - ➡ Identify “*path truncation attacks*”
  - In the example the direct link between AS 6 and AS 43 is a violation of the constraint

Which types of adversaries seen above can/can‘t this detect/prevent?

Summary

- Several general routing threats exist
- Internet relies on Inter-AS routing
- BGP has been designed without adversaries in mind
- Plethora of problems arise, simple attacks possible
- Potential solutions are
  - Hacks
  - Based on crypto (S-BGP, SIDR)
  - …more hacks ;-)