University of Glasgow

Internet Protocols Laboratory

Content Distribution and Routing

Networked Systems (H)

Lecture 9

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow | This work is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License.
To view a copy of this license, visit http://creativecommons.org/licenses/by-nd/4.0/ or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

Lecture Outline

- Content Distribution
- Internet Routing
  - Inter-domain routing
  - Routing security
  - Intra-domain routing
- Wrap-up

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

University of Glasgow

Content Distribution Networks (CDNs)

- Load balancing and latency reduction
- Implementation using DNS
- Implementation using anycast routing

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Content Distribution Networks (CDNs)

- **Content distribution networks (CDNs)** provide scalability, load balancing, and low latency
  - Host content for their customers in web caches spread around the world in edge networks and data centres
  - Reduces load on the main servers – rather than keep a local copy of the files, they link to the CDN-hosted copy
  - Reduces latency to respond to requests
  - Reduces chances of successful denial-of-service attack
- Many commercial CDNs available
- Many large organisations run their own
  - **Hypergiants** such as Google, Facebook, Netflix, …

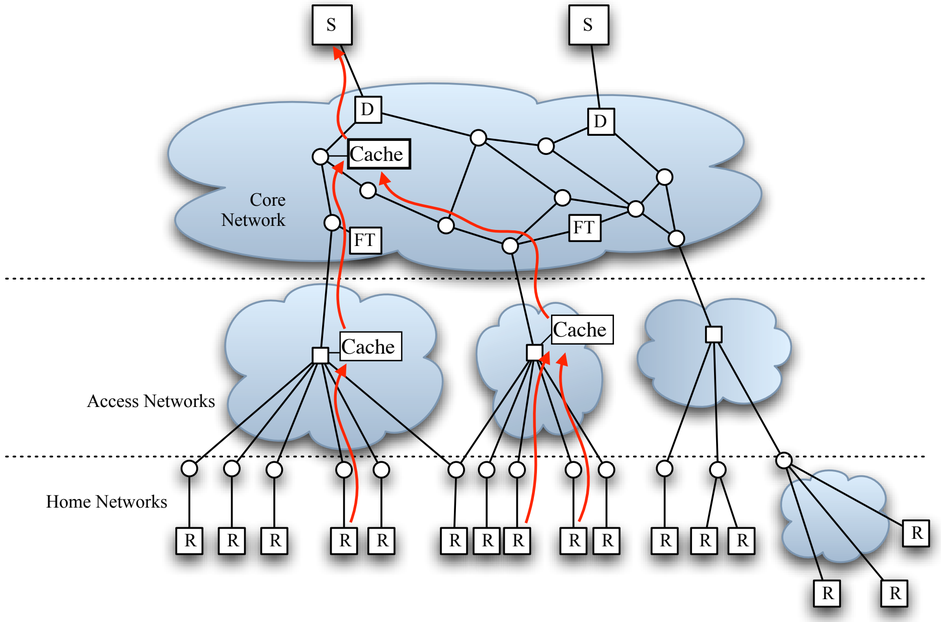

Akamai CLOUDFLARE fastly

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Using CDNs to Distribute Load

- CDNs distribute load by locating caches around the world and answering most requests locally
- Needs extensive investment, large-scale cooperation with ISPs, Internet Exchange Points, etc.
- Mutual benefit for an ISP to host the servers for a CDN
  - Reduces the load on the ISP's network
    - One copy of the content sent to cache; distributed many times
    - Avoids overload of link from ISP to outside world
    - e.g., Netflix distribute "tens of terabits per second" of video – not possible from a single data centre; needs local fanout
  - Increases the reach and robustness of the CDN

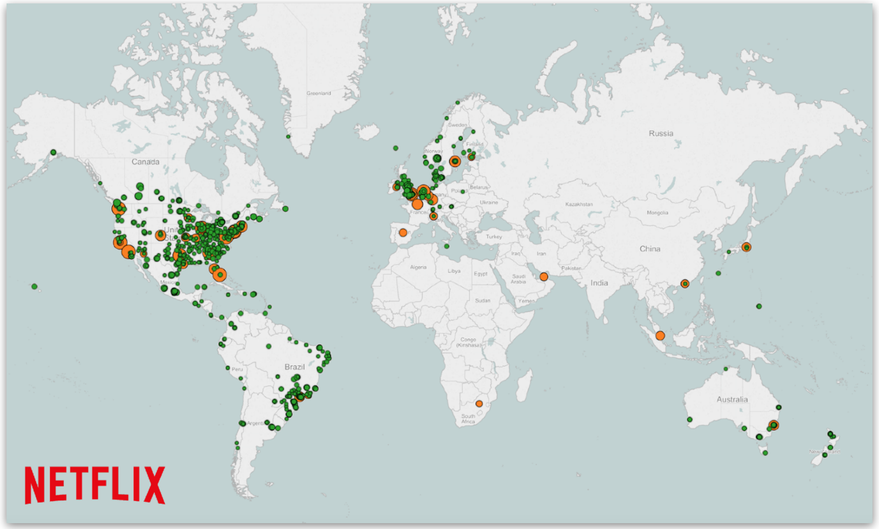

Netflix Open Connect CDN locations (source: https://media.netflix.com/en/company-blog/how-netflix-works-with-isps-around-the-globe-to-deliver-a-great-viewing-experience)

Akamai "more than 240,000 servers in over 130 countries"

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Using CDNs to reduce latency

- Key CDN benefit is reducing latency → content is delivered from nearby caches
  - e.g., requests to a US-based service are answered from a CDN cache in Europe, rather than having to cross long-distance path to the US
- Requires global distribution of CDN proxy caches
  - Are CDNs effectively serving the entire world?
  - Providing effective Internet access to developing regions requires more than just connectivity; also needs data centres and infrastructure to host the CDN nodes
- Effective for cacheable static content – video, software updates, images – may need edge compute infrastructure to support other applications

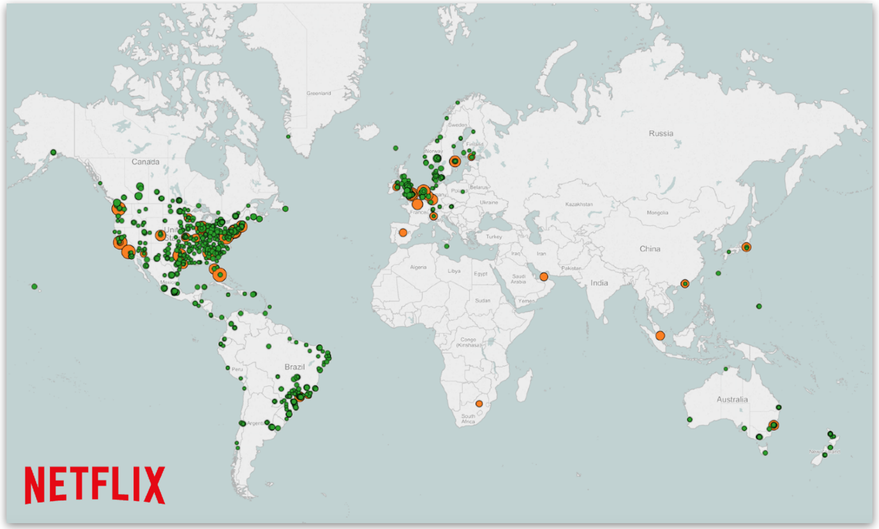

Netflix Open Connect CDN locations (source: https://media.netflix.com/en/company-blog/how-netflix-works-with-isps-around-the-globe-to-deliver-a-great-viewing-experience)

- CDN server locations broadly follow global income distribution; evidence of the digital divide?

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Locating Nearest CDN Node using DNS

- Locate nearest CDN node using the DNS
- **Every** resource hosted by the CDN is given a unique DNS name
  - Example: a website has a large movie file (e.g., https://example.com/images/kitten.m4v) that it wants to host on a CDN
  - CDN stores it at https://E3CC4EBC-343C-4F72-8C50-A059B4509C84.cdn.com/kitten.m4v where the hostname is unique and only used for this file
  - The original website links to the movie file the unique URL provided by the CDN
- DNS servers for the CDN returns different A or AAAA records for the hostname, depending on where it's requested from and what CDN caches have the data
  - Directs to local cache based on IP address of resolver making the query
  - Doesn't need to be particularly accurate – trying to figure out if you're in UK rather than US, to direct you to data centre in London rather than New York
  - Allows very fine-grained control, but puts high load on DNS

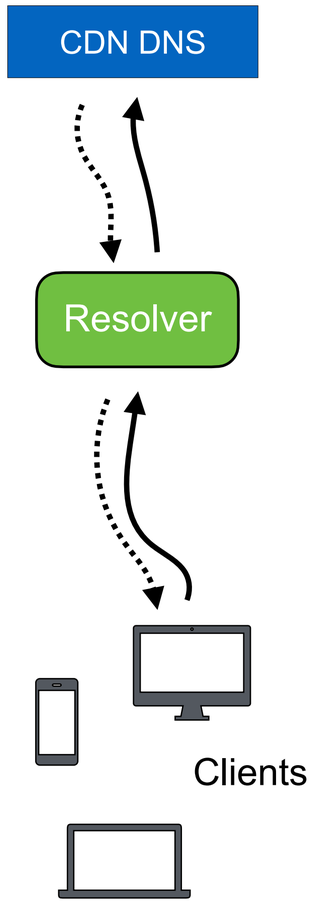

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Locating Nearest CDN Node using Anycast

- Locate nearest CDN node using **anycast routing**
- Each resource hosted by the CDN has a unique filename, but all use the same hostname
  - Example, the CDN might host https://example.com/images/kitten.m4v as https://cdn.com/E3CC4EBC-343C-4F72-8C50-A059B4509C84.m4v
  - The DNS name always maps to the same IP address – the IP address of a load balancer at the entrance to a data centre
- **The CDN uses multiple data centres – all assigned the same IP addresses**
- Each data centre advertises its addresses via BGP
- Inter-domain Internet routing ensures traffic goes to the closest data centre to source

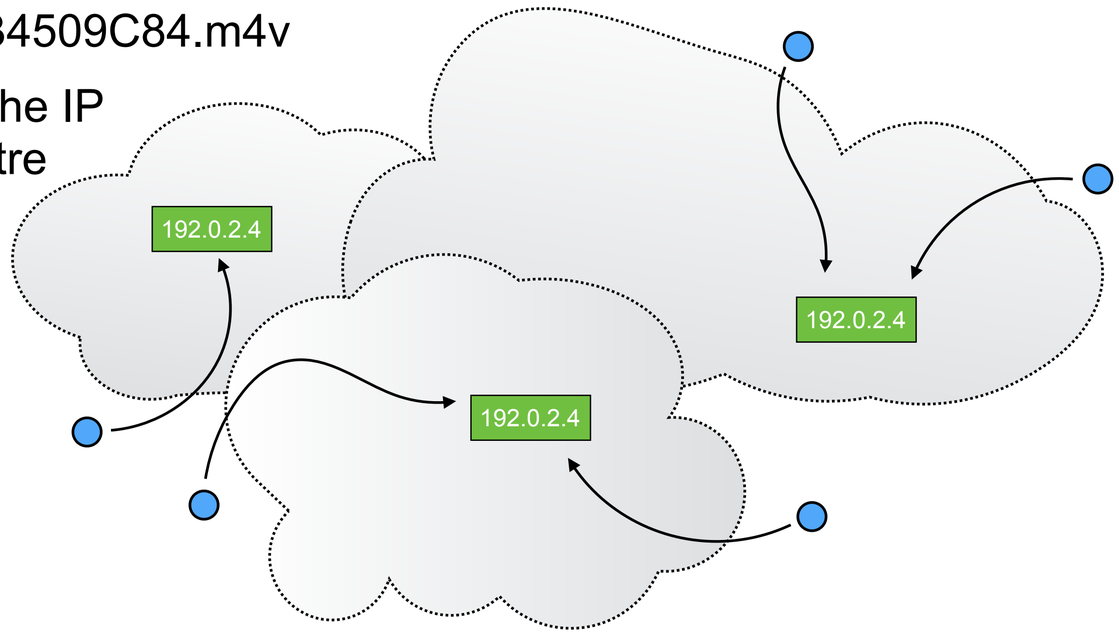

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

University of Glasgow

Inter-domain Routing

- Autonomous Systems and the AS graph
- Routing at the edge
- Routing in the core
- BGP

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Inter-domain Routing

- The Internet is a network of networks
  - Each network is an **Autonomous System (AS)**
  - Each network is a separate routing domain

- **Inter-domain routing** is the problem of finding the best path from the source network to the destination network
  - Treat each network as a node on the routing graph (the “AS topology graph”)
  - Treat connections between networks as edges in the graph
  - What is the best set of networks to choose to get from source to destination

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Autonomous Systems

- An **Autonomous System (AS)** is an independently operated network
  - An Internet service provider, or other organisation, that operates an IP network and wants to participate in BGP interdomain routing
  - Some organisations operate more than one AS
- Each AS is identified by a unique **AS Number** and owns one or more **IP address prefixes**
  - As of 14 February 2026, there were 78,176 AS numbers in-use, advertising 1,058,538 IP address prefixes

Source: Geoff Huston, APNIC, via https://www.potaroo.net/tools/asn32/ and https://www.cidr-report.org/as2.0/

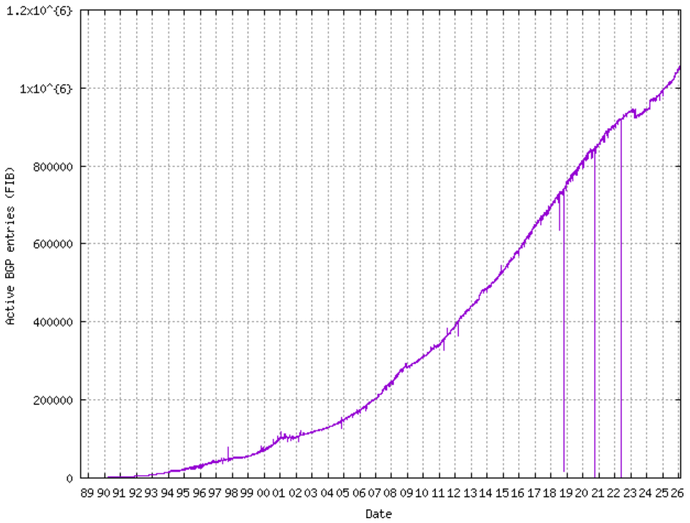

Number of IPv4 prefixes advertised in BGP

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

AS-level Internet Topology Graph

- AS-level topology – what are the inter-AS connections?
  - Each node is an autonomous system
    - Position around circle based on geographic location
    - Distance from centre based on number of links to other networks → more links, closer to the centre
  - Edges show links between autonomous systems
    - A **potential route** traffic can flow
    - Not all possible routes are in use

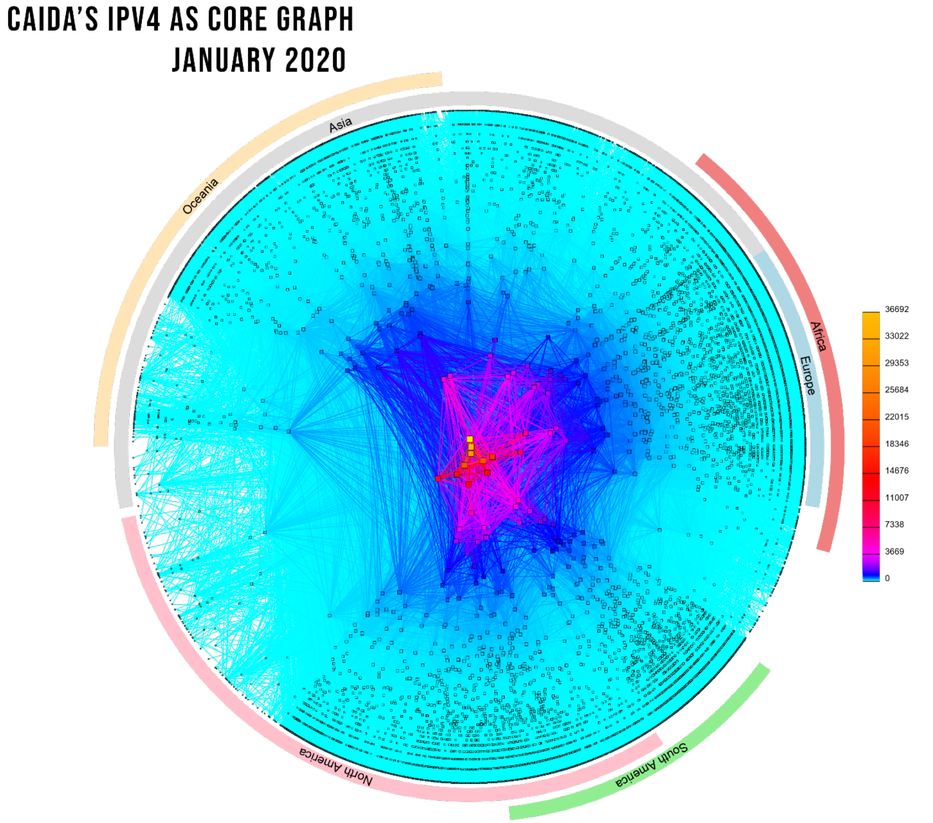

Source: CAIDA https://www.caida.org/projects/as-core/

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

AS-level Internet Topology Graph

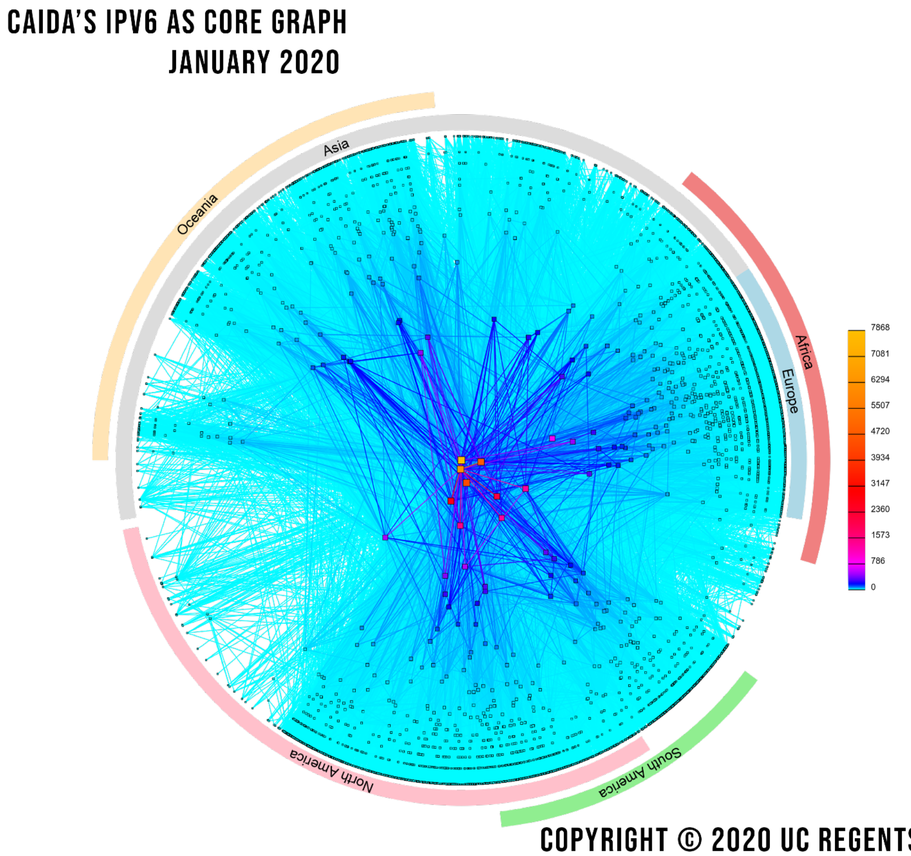

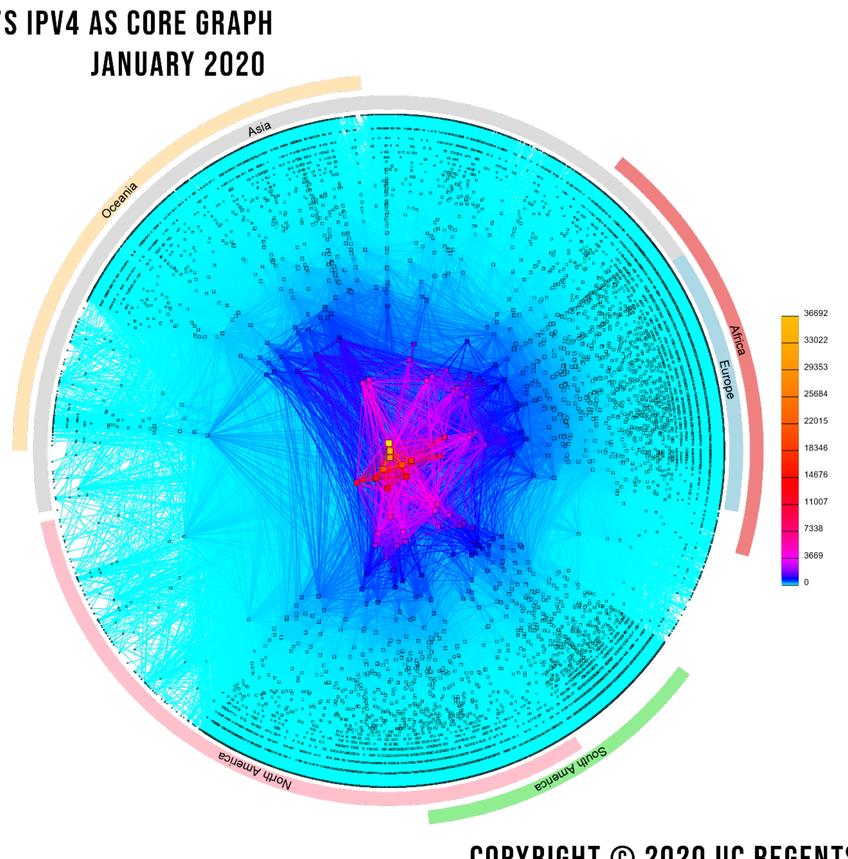

Source: CAIDA https://www.caida.org/projects/as-core/

Source: http://www.caida.org/research/topology/as_core_network/2015/

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Differences Between IPv4 & IPv6 Topologies

- IPv6 AS-level topology graph is sparse compared to IPv4, but there is significant IPv6 usage → IPv4 has 30 years head start on deployment; the topology can be expected to be more densely interconnected

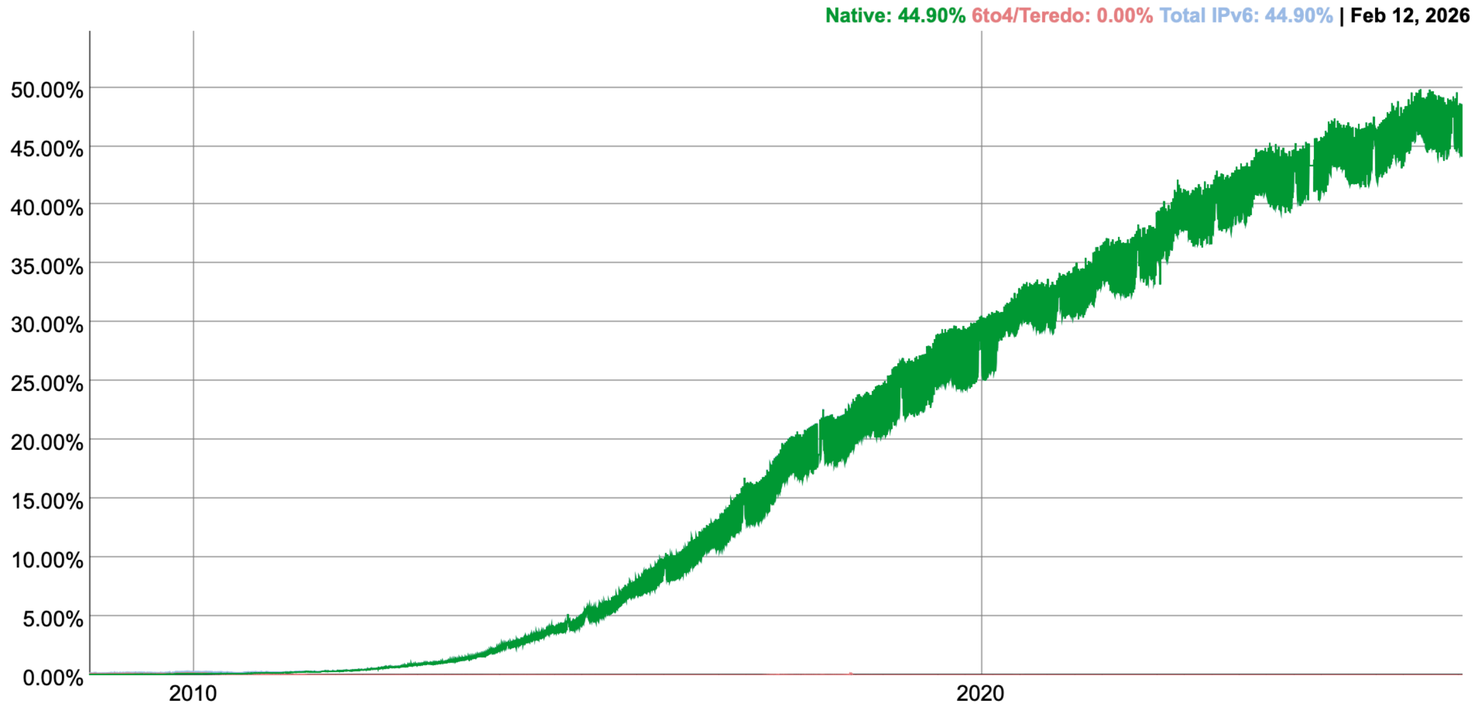

Source: https://www.google.com/intl/en/ipv6/statistics.html

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Routing at the Edge

- Devices at the edge tend to have simple **routing tables**
  - The devices on the local network
  - A **default route** for everything else
  - e.g., routing table for hosts in Glasgow SoCS:

```
Network:        Netmask:        Gateway:
130.209.240.0   255.255.240.0   eth0
default         0.0.0.0         130.209.240.48
```

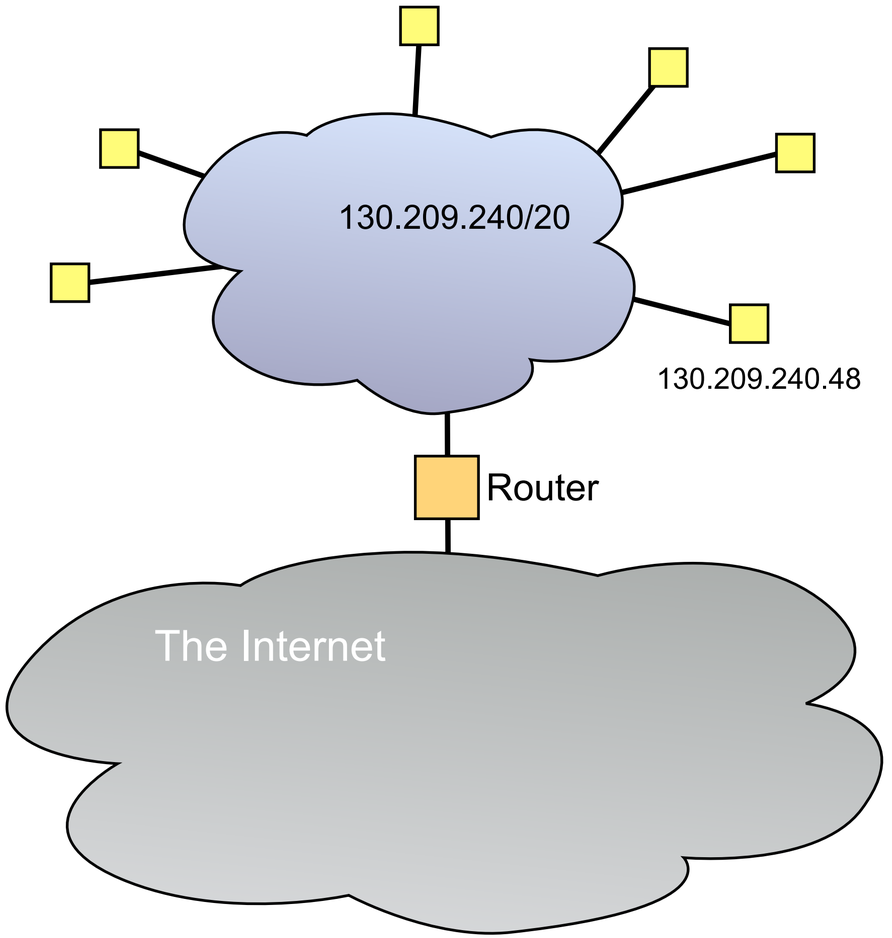

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Routing in the Core of the Internet (1/2)

- Networks near the edge can often use a **default route** for much of their traffic
- Core networks need full routing table
  - The **default free zone** (DFZ)
  - With enough interconnections, they must store routing information about the broad network topology
  - Many possible paths – need information to choose the most appropriate

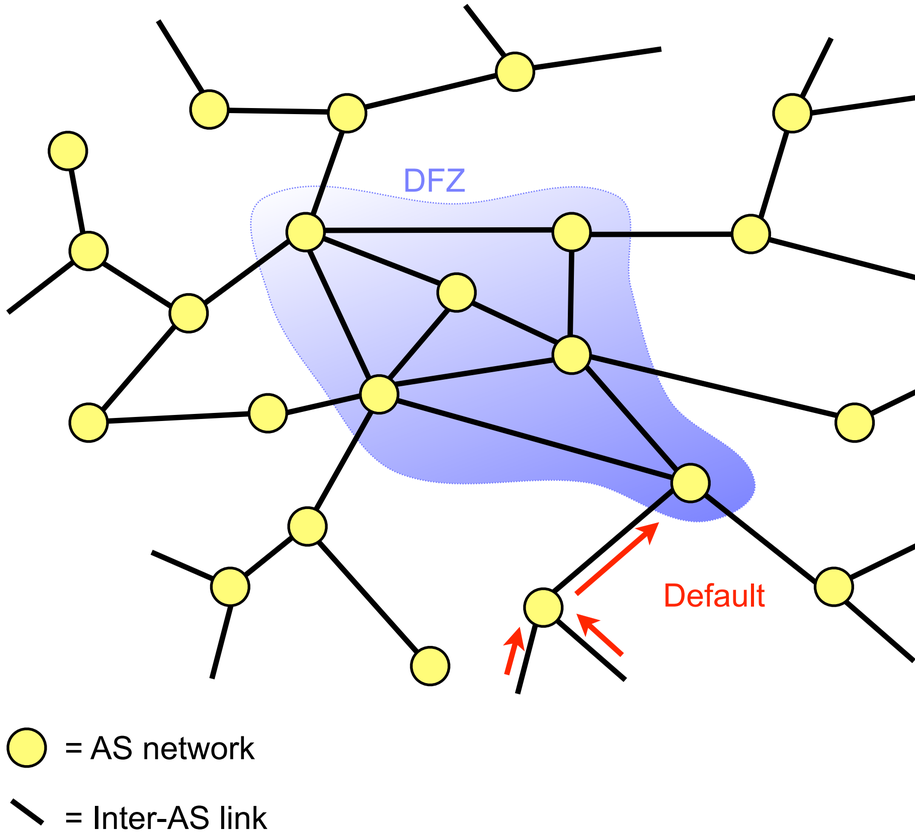

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Routing in the Core of the Internet (2/2)

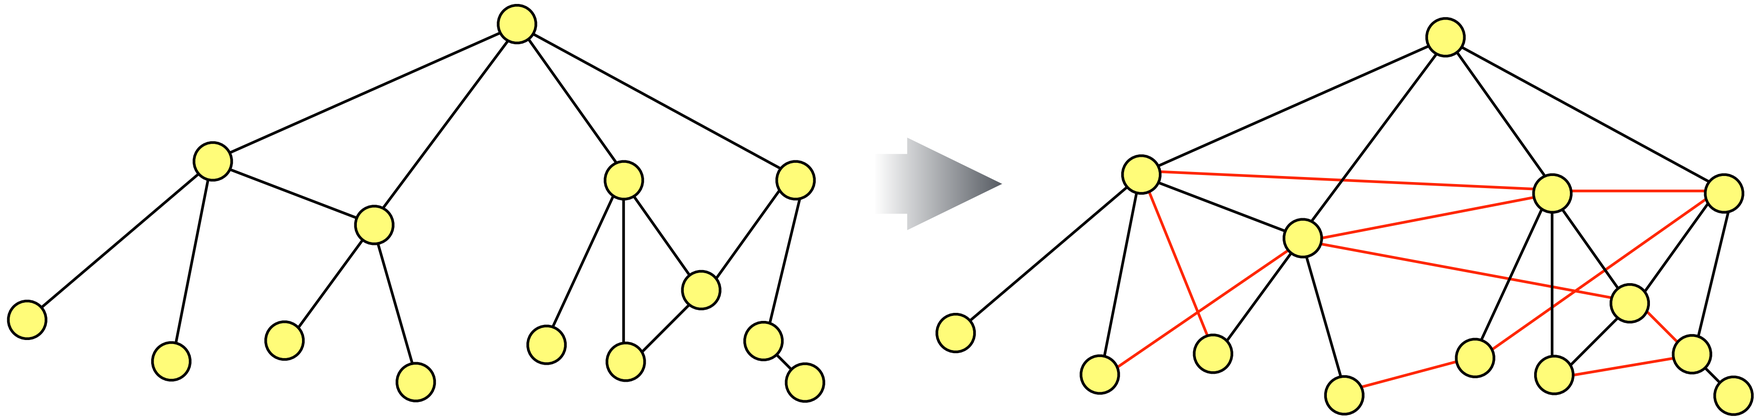

- AS-level topology flattening → increasingly rich interconnections
- **Internet Exchange Points (IXPs)** are becoming commonplace
  - Example: London Internet Exchange
  - > 950 different ASes interconnected
  - > 10 Tbps of traffic

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Routing Policy

- Interdomain routing is between competitors
  - Each AS is independently operated and may compete for customers
  - Business, political, and economic relationships influence routing
- Routing **must** consider policy
  - Policy restrictions on who can determine the network topology
  - Policy restrictions on which route traffic should follow between a particular source and destination
  - Policy restrictions might prioritise non-shortest path routes
    - Do you prioritise cost, bandwidth, or latency when choosing a route?
    - Traffic between certain networks may be prohibiting from passing through other networks
    - Traffic between certain regions should avoid other regions

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Border Gateway Protocol (BGP)

- Internet routing uses the **Border Gateway Protocol (BGP)**
  - External BGP (**eBGP**) used to exchange routing information between ASes
    - An eBGP session is used by an AS to exchange information about the AS-level topology with a neighbouring AS
    - Runs over a TCP connection between two routers, one in each AS
    - Exchanges knowledge of the AS-level topology of the network, filtered according to policy
    - Used to compute appropriate **interdomain routes** to each destination AS
  - Internal BGP (**iBGP**) propagates this information to routers within an AS
    - An **intra-domain routing protocol** determines routes within the AS
    - iBGP distributes information to internal routers on how to reach external destinations

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Routing Information Exchanged in eBGP

- eBGP routers advertise lists of IP address prefixes and associated AS-level paths
- Combined to form a routing table

	Prefix	Next Hop	AS Path
…			
*	12.10.231.0/24	194.68.130.254	5459 5413 5696 7369 i
*		158.43.133.48	1849 702 701 6113 5696 7369 i
*		193.0.0.242	3333 286 5696 7369 i
*		204.212.44.128	234 266 237 3561 5696 7369 i
*>		202.232.1.8	2497 5696 7369 i
*		204.70.4.89	3561 5696 7369 i
*		131.103.20.49	1225 3561 5696 7369 i
*	62.224.0.0/19	134.24.127.3	1740 3561 2578 8371 i
*		194.68.130.254	5459 2529 3561 2578 8371 i
*		158.43.133.48	1849 702 701 3216 3216 3216 8371 8371 i
*		193.0.0.242	3333 286 3561 2578 8371 i
*		144.228.240.93	1239 3561 2578 8371 i
*		204.212.44.128	234 266 237 3561 2578 8371 i
*		202.232.1.8	2497 3561 2578 8371 i
*		205.238.48.3	2914 3561 2578 8371 i
*>		204.70.4.89	3561 2578 8371 i
*		131.103.20.49	1225 3561 2578 8371 i
…			

- Hosts with IP addresses in the range 12.10.231.0 - 12.10.231.255 are in AS 7369.
- That AS is best reached via AS 2497, then AS 5696.
- Packets destined for those addresses should be sent to address 202.232.1.8 next, from where they will be forwarded.
- The **longest prefix** covering a given address is used when selecting how to route traffic to that address

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

AS Topology Graph

An example fragment of the AS topology graph:

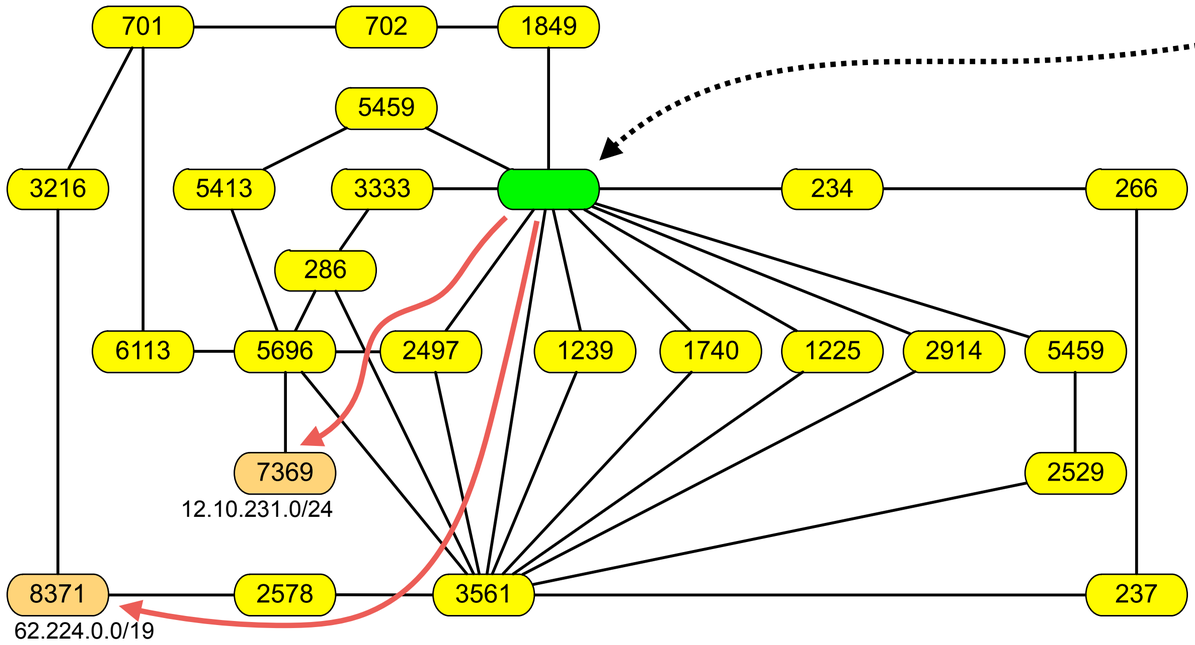

	Prefix	Next Hop	AS Path
…			
*	12.10.231.0/24	194.68.130.254	5459 5413 5696 7369 i
*		158.43.133.48	1849 702 701 6113 5696 7369 i
*		193.0.0.242	3333 286 5696 7369 i
*		204.212.44.128	234 266 237 3561 5696 7369 i
*>		202.232.1.8	2497 5696 7369 i
*		204.70.4.89	3561 5696 7369 i
*		131.103.20.49	1225 3561 5696 7369 i
*	62.224.0.0/19	134.24.127.3	1740 3561 2578 8371 i
*		194.68.130.254	5459 2529 3561 2578 8371 i
*		158.43.133.48	1849 702 701 3216 3216 3216 8371 8371 i
*		193.0.0.242	3333 286 3561 2578 8371 i
*		144.228.240.93	1239 3561 2578 8371 i
*		204.212.44.128	234 266 237 3561 2578 8371 i
*		202.232.1.8	2497 3561 2578 8371 i
*		205.238.48.3	2914 3561 2578 8371 i
*>		204.70.4.89	3561 2578 8371 i
*		131.103.20.49	1225 3561 2578 8371 i
…			

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Routing Policy in eBGP

- Each AS chooses what routes to advertise to its neighbours
  - Usual to drop some routes from the advertisement – depends on the chosen routing policy
  - Common approach: the Gao-Rexford rules:

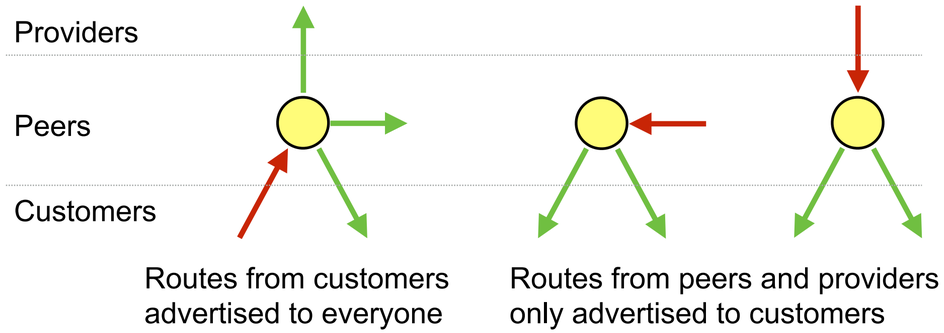


Lixin Gao


Jennifer Rexford

  - Gao-Rexford rules specify **what routes are advertised**, not how traffic is forwarded – difference between **control plane** and **data plane**
  - Resulting graph is valley-free DAG: traffic flows up to common provider, then down
  - Designed to optimise **profit**, not routing efficiency – don't advertise expensive routes

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

BGP Routing Decision Process (1/2)

- Each AS exchanges routing information with neighbours
  - AS-level topology information, filtered based on policies applied by each AS; gives partial view of AS-level topology
- Each AS applies policy to choose what routes to use:
  - Choose what route to install in forwarding table for each destination prefix
  - BGP may not find a route, even if one exists, since may be prohibited by policy
  - Routes are often not the shortest AS path – shortest policy-compliant path
- Each AS applies policy to determine what routes to advertise to neighbours
  - Filter out routes the neighbour should not use (e.g., Gao-Rexford rules, and other reasons)
  - De-prioritise certain other routes
  - Tag routes for special processing where business agreements exist
  - Mapping business goals to BGP policies is poorly documented → many operational secrets

Table 2: Simplified BGP decision process [6, 24]. This table was also provided with the survey.

#	Criteria
1	Highest LocalPref
2	Lowest AS Path Length
3	Lowest origin type
4	Lowest MED
5	eBGP-learned over iBGP-learned
6	Lowest IGP cost to border router (hot-potato routing)
7	If both paths are external, prefer the path that was received first (i.e., the oldest path) [6]
8	Lowest router ID (to break ties)

Source: Phillipa Gill, Michael Schapira, and Sharon Goldberg, “A Survey of Interdomain Routing Policies”, ACM CCR, V44, N1, January 2014, p29-34

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

BGP Routing Decision Process (2/2)

- BGP exchanges routes between autonomous systems
  - Policy, economic, and political concerns outweigh shortest path in many cases
  - Routing is between competitors – there is little trust and little public data
  - The BGP decision process is public – the input data is not, making it difficult to evaluate how routing decisions are made

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

University of Glasgow

Routing Security

- Internet Routing Security
- RPKI
- MANRS

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Internet Routing is not Secure

- Any AS participating in BGP routing can announce any address prefix – whether or not they own it
  - Sometimes this is accidental
  - Sometimes this is malicious
- Result is that traffic is misdirected: **BGP hijacking**
  - Security risk if traffic can be misdirected
  - Accidental hijacking is a serious stability problem for the network

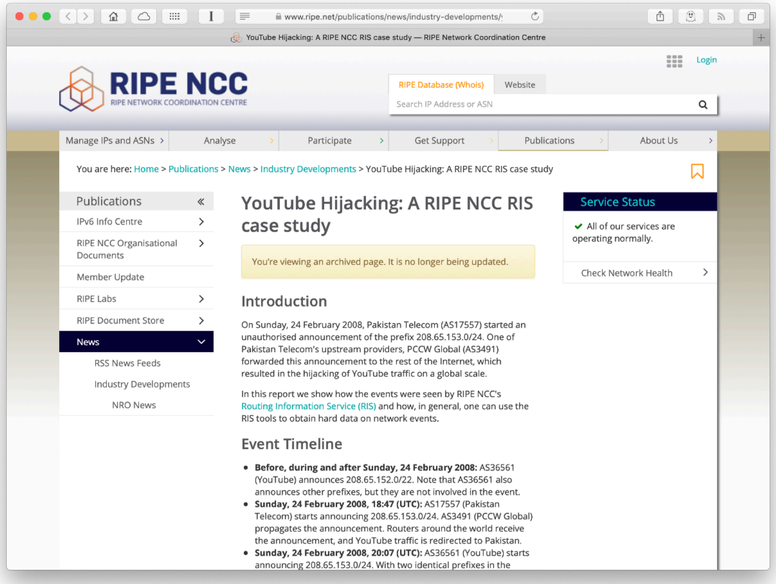

https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Resource Public Key Infrastructure (RPKI)

- RPKI is an attempt to secure Internet routing
- Allows an AS make a signed **Route Origin Authorisation** (ROA), a digital signature, announced via BGP, proving the AS is authorised to announce a IP address prefix
  - Hierarchical delegation: RIR → provider AS → customer AS
  - Allows a router to validate whether a prefix announcement is authorised; hijacked prefixes will not have a valid signature
  - Becomes part of BGP policy: prefer signed prefixes
- As of December 2024, more than half of IPv4 addresses are covered by prefixes with valid digital signatures
- More information: https://isbgpsafeyet.com/



RPKI is Coming of Age

A Longitudinal Study of RPKI Deployment and Invalid Route Origins

Taejoong Chung
Rochester Institute of Technology

Emile Aben
RIPE NCC

Tim Bruijnzeels
NLNetLabs

Balakrishnan Chandrasekaran
Max Planck Institute for Informatics

David Choffnes
Northeastern University

Dave Levin
University of Maryland

Bruce M. Maggs
Duke University and
Akamai Technologies

Alan Mislove
Northeastern University

Roland van Rijswijk-Deij
University of Twente and
NLNetLabs

John Rula
Akamai Technologies

Nick Sullivan
Cloudflare

ABSTRACT

Despite its critical role in Internet connectivity, the Border Gateway Protocol (BGP) remains highly vulnerable to attacks such as prefix hijacking, where an Autonomous System (AS) announces routes for IP space it does not control. To address this issue, the Resource Public Key Infrastructure (RPKI) was developed starting in 2008, with deployment beginning in 2011. This paper performs the first comprehensive, longitudinal study of the deployment, coverage, and quality of RPKI.

We use a unique dataset containing *all* RPKI Route Origin Authorizations (ROAs) from the moment RPKI was first deployed, more than 8 years ago. We combine this dataset with BGP announcements from more than 3,300 BGP collectors worldwide. Our analysis shows the after a gradual start, RPKI has seen a rapid increase in adoption over the past two years. We also show that although misconfigurations were rampant when RPKI was first deployed (causing many announcements to appear as invalid) they are quite rare today. We develop a taxonomy of invalid RPKI announcements, then quantify their prevalence. We further identify suspicious announcements indicative of prefix hijacking and present case studies of likely hijacks.

Overall, we conclude that while misconfigurations still do occur, RPKI is "ready for the big screen," and routing security can be increased by dropping invalid announcements. To foster reproducibility and further studies, we release all RPKI data and the tools we used to analyze it into the public domain.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IMC '19, October 21–23, 2019, Amsterdam, Netherlands

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6948-0/19/10…$15.00
https://doi.org/10.1145/3355369.3355596

1 INTRODUCTION

The Border Gateway Protocol (BGP) is *the* mechanism that allows routers to construct routing tables across the Internet. Unfortunately, the original BGP protocol lacked many security features (e.g., authorization of IP prefix announcements), making BGP vulnerable to attacks such as prefix hijacking [3, 5, 7, 14] and route leaks [5]. To defend against these threats, the Resource Public Key Infrastructure (RPKI) was developed in April 2008 as part of the IETF in the SIDR Working Group [54]. Beta deployments followed in the years after, until all Regional Internet Registries (RIRs) started production deployment of RPKI in January 2011.

At its core, RPKI is a hierarchical Public Key Infrastructure (PKI) that binds Internet Number Resources (INRs) such as Autonomous System Numbers (ASNs) and IP addresses to public keys via certificates. The corresponding private keys can be used by certificate holders to make attestations about these INRs—most importantly, Route Origin Authorization (ROA) objects. ROAs allow a certificate holder to authorize an ASN to announce certain IP prefixes, and are signed using the private key of a certificate covering the IP space.

Each of the five RIRs operate their own RPKI *trust anchor* (equivalent to a root certificate in other PKIs), the private key of which is used to sign such certificates. The RIRs also offer hosted services to their members, enabling them to obtain RPKI certificates and generate ROAs.

RPKI objects including certificates, ROAs, and supporting structures such as manifests and certificate revocation lists (CRLs) are published in so-called RPKI *repositories*. RPKI validation software—called Relying Party (RP) software—retrieves objects from these repositories and performs cryptographic validation of the content, ultimately producing a set of valid ROAs. A *validating router* can then use this set to verify incoming BGP announcements. If the router finds a BGP announcement to be in conflict with the set of valid ROAs, it should reject the announcement as (by definition) the origin AS is not authorized to announce the IP prefix(es).

While RPKI sounds straightforward, in practice it can be complex, creating many opportunities for mistakes. For example, an AS may sub-allocate an IP prefix to a customer AS without updating its ROAs or mistakenly include the wrong range of IP prefixes in a

Taejoong Chung, *et al.*, "RPKI is coming of age: A longitudinal study of RPKI deployment and invalid route origins", ACM Internet Measurement Conference, Amsterdam, October 2019.https://dx.doi.org/10.1145/3355369.3355596

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Mutually Agreed Norms for Routing Security

- **Mutually Agreed Norms for Routing Security (MANRS)** project aims to improve routing security (https://www.manrs.org)
  - Filtering
  - Anti-spoofing
  - Coordination
  - Global validation via RKPI
  - Combination of signed route advertisements via RKPI to catch malicious attacks, with best practises for route filtering, anti-spoofing, and coordination between ISPs to protect against misconfigurations

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

University of Glasgow

Intra-domain Routing

- Distance vector and link state algorithms

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Intra-domain Unicast Routing

- BGP gives information on the path to reach other networks
- How to route traffic within an AS?
  - Single trust domain
    - No policy restrictions on who can determine network topology
    - No policy restrictions on which links can be used
  - Desire efficient routing → shortest path
    - Make best use of the network you have available
  - Two approaches
    - Distance vector – the Routing Information Protocol (RIP)
    - Link state – Open Shortest Path First routing (OSPF)

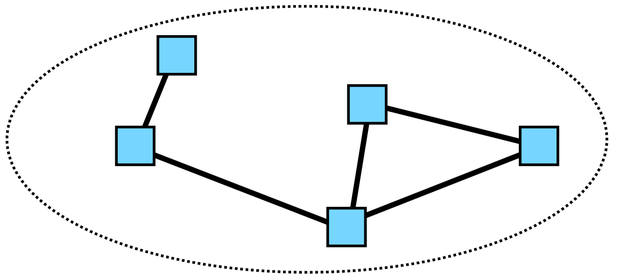

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Distance Vector Routing

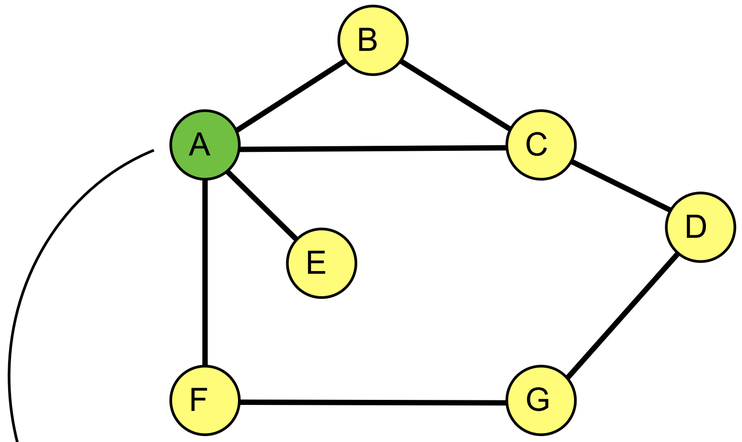

Routing Table at Node A

Destination	Distance	Next Hop
B	1	B
C	1	C
D	2	C
E	1	E
F	1	F
G	2	F

- Nodes maintain a vector containing the distance to every other node in network
  - Periodically exchanged with neighbours, eventually every node know distance and next hop to every other node
- Packets forwarded on shortest path to destination
- Not widely used in practice:
  - Slow to converge in normal operation
  - Count-to-infinity problem makes this worse on failure in networks containing loops

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Link State Routing

- Nodes know the links to their neighbours, and cost of using those links – **link state** information
  - Flood this information throughout the network on startup, and whenever a link is added or removed
    - Send address of node sending the update, the list of directly connected neighbours and costs for link usage, and message sequence number
    - Sequence numbers prevent loops
  - Each node then directly calculates shortest path to every other node, uses as routing table
- Flooding link state data from all nodes ensures all nodes know the entire topology
  - Each node uses Dijkstra's shortest-path algorithm to calculate optimal route to every other node
  - Optimal is assumed to be the shortest path, by cost

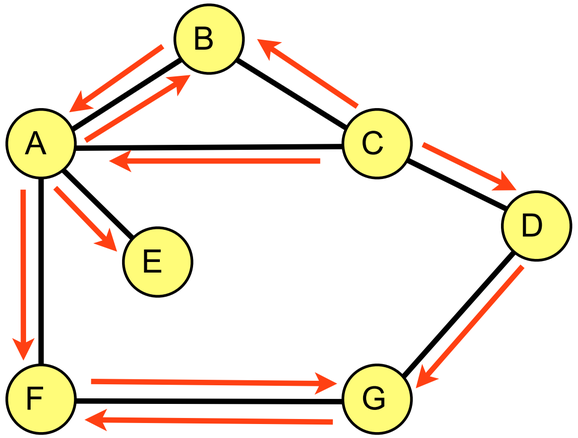

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Distance Vector vs. Link State

Distance vector routing:

- Simple to implement
- Routers only store the distance to each other node: O(n)
- Suffers from slow convergence

Link State routing:

- More complex
- Requires each router to store complete network map: O(n^2)
- Much faster convergence

Slow convergence times make distance vector routing unsuitable for large networks

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Challenges in Intra-domain Routing



- How to rapidly recover from link failures?
  - Construction work surprisingly good at breaking network cables; trawlers do similar damage to undersea cables
  - Good network designs have multiple paths from source to destination
  - Must quickly notice failure and switch to backup path
    - If the link is used for telephony or video streaming, this must happen in $<1/60^{th}$ of a second else it interrupts the speech or video playout
  - **Fast failover** – pre-calculate alternative paths to allow rapid switchover
- How to balance load across multiple paths?
  - **Equal cost multipath routing** – if two paths with the same cost, alternate traffic between them
  - Care needed: TCP uses a triple duplicate ACK to signal loss so is insensitive to small amount of packet re-ordering, but treats large amounts of reordering as loss and slows down
  - RTP applications don't care about order, as long as packets arrive before deadline

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

University of Glasgow

Wrap-Up

- Assessment and Revision Lecture
- Level 4 and MSci Project

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Assessment and Revision Lecture

- Marks for assessed coursework will be available shortly
- Final exam, worth 80% of the marks for the course, will be held in April/May
  - Three questions, focussed on testing your understanding of networked systems
  - **Tell me what you think** – online, open book, exams focus more on deeper understanding than on bookwork
  - Past papers are on Moodle – the style of questions in the 2020 exam is most representative

- Expect a **revision lecture** just after Easter – time for Q&A plus (non-assessed) discussion of the future evolution of the network

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Level 4 and MSci Projects

- If you are interested in Level 4 or MSci projects relating to networked systems, please contact me – colin.perkins@glasgow.ac.uk
  - Keen to work with motivated students to develop projects on topics including:
    - Understanding the Internet standards development process
    - Improving Internet protocol standards and specifications
    - Transport protocols, real-time applications, and QUIC
    - New networking APIs in high-level languages (e.g., Rust)
    - Network measurement
  - Strong focus on interactions with IETF standards and research communities:
    - Some potential projects are strongly technical
    - Others focus heavily on the standardisation process

Colin Perkins | https://csperkins.org/ | Copyright © 2026 University of Glasgow

University of Glasgow

The End

