Cloud Provider Connectivity in the Flat Internet

Todd Arnold[†], Jia He[†], Weifan Jiang[†], Matt Calder[‡†], Italo Cunha[♮†], Vasileios Giotsas[#], Ethan Katz-Bassett[†]

[†] Columbia University in the City of New York

[‡] Microsoft

[♮] UFMG Universidade Federal de Minas Gerais

[#] Lancaster University

Evolving Internet Topology

- Traditional: Hierarchical
  - Tier-1 ISPs are global networks, and all other networks fall under at least one
  - Tier-2 ISPs are larger, regional networks
  - Tier-3 ISPs interconnect edge networks
  - Edge networks at the bottom
  - Networks pay higher tiers to transport their data (a.k.a. *transit*)

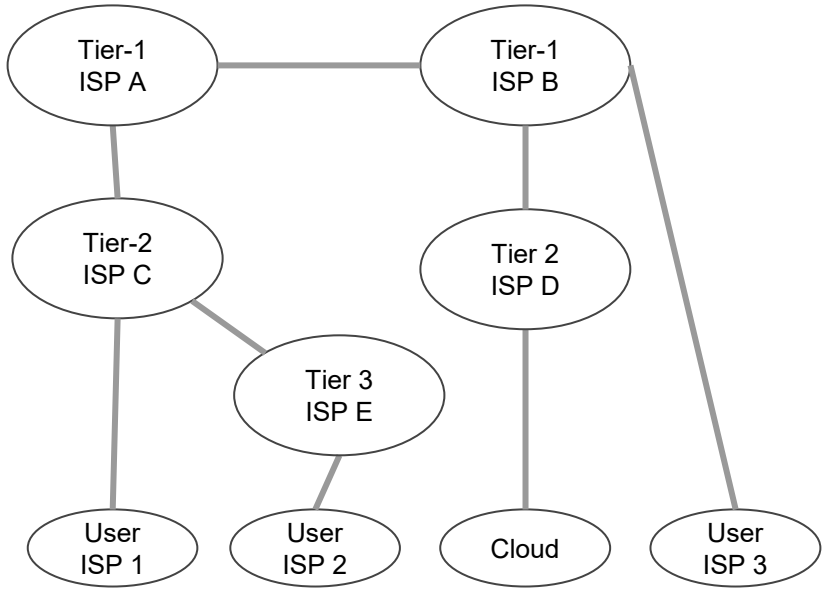

Evolving Internet Topology

- Traditional: Hierarchical
  - Tier-1 ISPs are global networks, and all other networks fall under at least one
  - Tier-2 ISPs are larger, regional networks
  - Tier-3 ISPs interconnect edge networks
  - Edge networks at the bottom
  - Networks pay higher tiers to transport their data (a.k.a. *transit*)

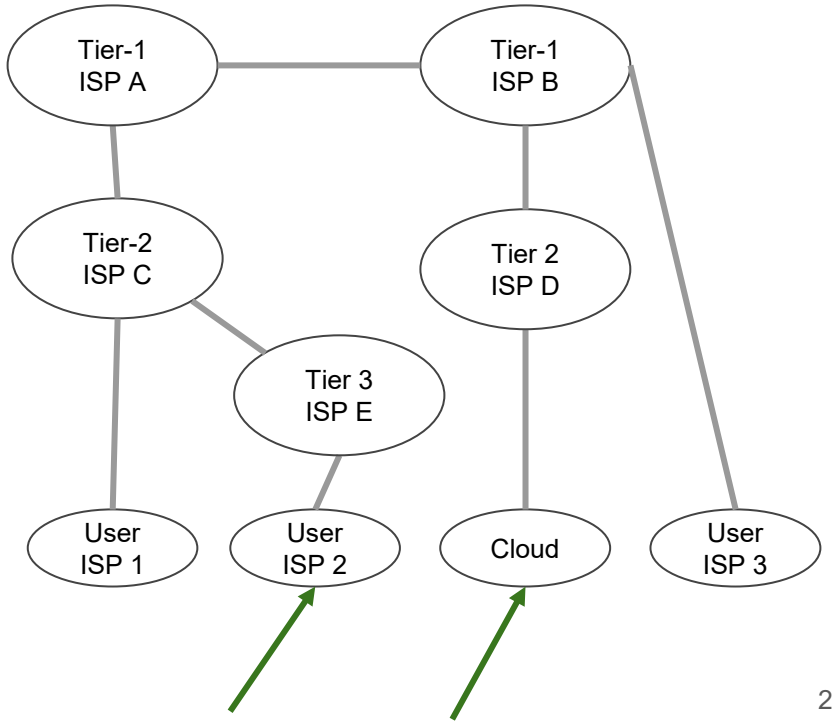

Evolving Internet Topology

- Traditional: Hierarchical
  - Tier-1 ISPs are global networks, and all other networks fall under at least one
  - Tier-2 ISPs are larger, regional networks
  - Tier-3 ISPs interconnect edge networks
  - Edge networks at the bottom
  - Networks pay higher tiers to transport their data (a.k.a. *transit*)

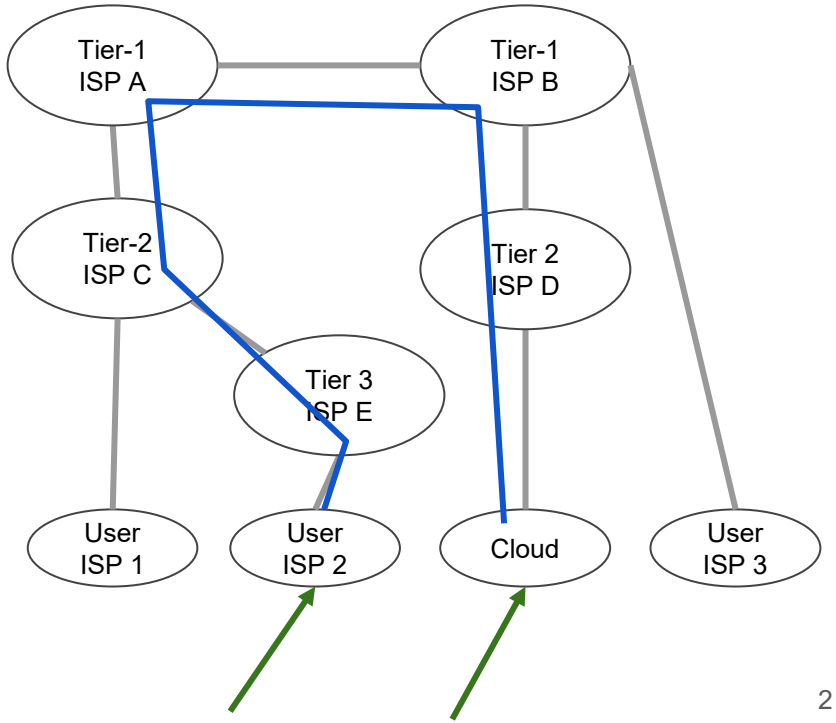

Evolving Internet Topology

- Traditional: Hierarchical
  - Tier-1 ISPs are global networks, and all other networks fall under at least one
  - Tier-2 ISPs are larger, regional networks
  - Tier-3 ISPs interconnect edge networks
  - Edge networks at the bottom
  - Networks pay higher tiers to transport their data (a.k.a. *transit*)
- Recent Years: Internet Flattening
  - Increased direct connectivity between networks at lower tiers

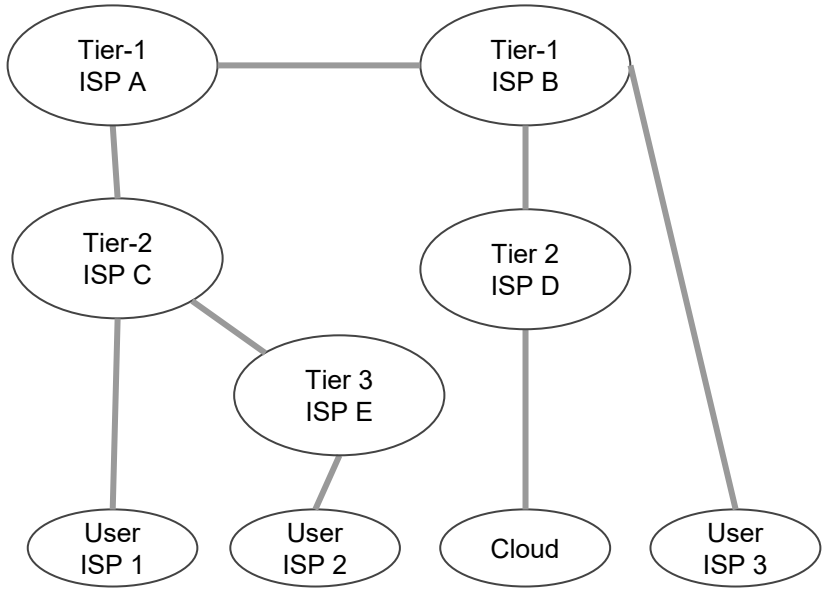

Evolving Internet Topology

- Traditional: Hierarchical
  - Tier-1 ISPs are global networks, and all other networks fall under at least one
  - Tier-2 ISPs are larger, regional networks
  - Tier-3 ISPs interconnect edge networks
  - Edge networks at the bottom
  - Networks pay higher tiers to transport their data (a.k.a. *transit*)
- Recent Years: Internet Flattening
  - Increased direct connectivity between networks at lower tiers

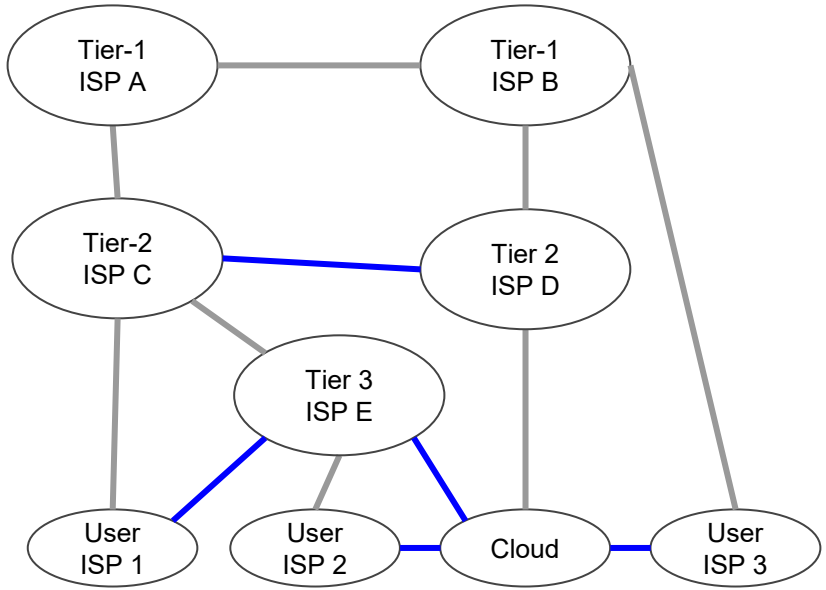

Evolving Internet Topology

- Traditional: Hierarchical
  - Tier-1 ISPs are global networks, and all other networks fall under at least one
  - Tier-2 ISPs are larger, regional networks
  - Tier-3 ISPs interconnect edge networks
  - Edge networks at the bottom
  - Networks pay higher tiers to transport their data (a.k.a. *transit*)
- Recent Years: Internet Flattening
  - Increased direct connectivity between networks at lower tiers

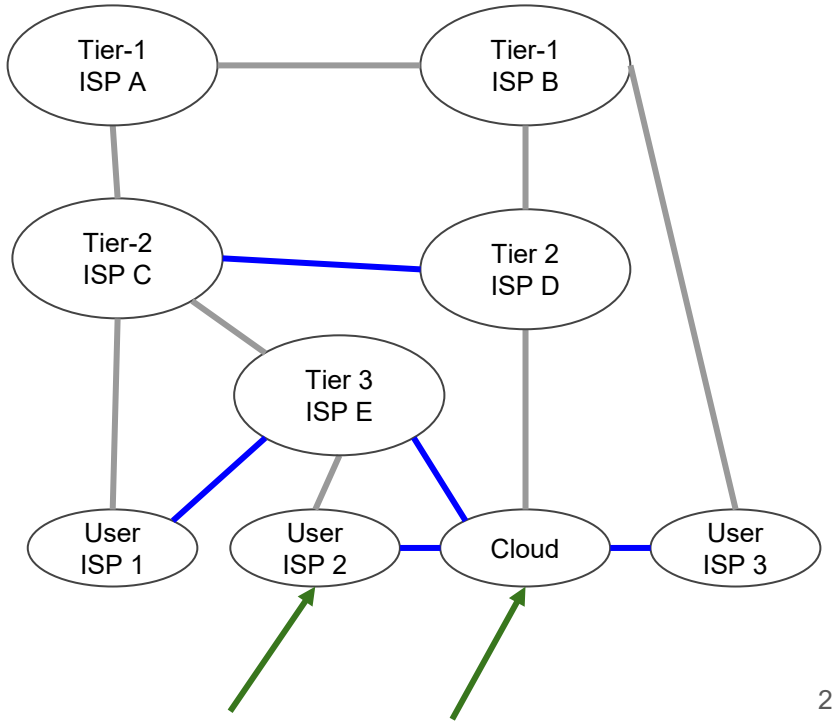

Motivation and Goals

- Majority of Internet traffic now occurs over direct connections
- Impact of flattening is not captured by traditional approaches
  - Invisible to traditional vantage points
  - Existing metrics of importance (*e.g.,* customer cone)
    - Do not reflect the rich peering interconnectivity of the flat Internet
    - Focus on how much transit an AS could provide rather than how much it does provide
- To understand this gap and capture the progress of Internet flattening
  - Uncover the missing links
  - Understand to what degree they enable the major cloud providers (Amazon, Google, IBM, and Microsoft) to bypass the traditional hierarchy

Methodology

AS topology graph from two sources

- CAIDA’s AS relationship dataset

Methodology

AS topology graph from two sources

- CAIDA’s AS relationship dataset
- Traceroutes from inside clouds
  - Used to identify directly connected neighbors to add to the topology
  - Map IP-to-ASN

Methodology

AS topology graph from two sources

- CAIDA's AS relationship dataset
- Traceroutes from inside clouds
  - Used to identify directly connected neighbors to add to the topology
  - Map IP-to-ASN
- Validation
  - Iterative process with feedback from both Microsoft and Google
  - Worked with Microsoft while we refined our methodology
  - Google's feedback validated our refinements
  - Microsoft: 11% FDR, 21% FNR, 3,565 neighbors
  - Google: 15% FDR, similar FNR, 7,554 neighbors
  - Amazon: 1,188 neighbors
  - IBM: 2,747

Calculating Hierarchy-free Reachability

- Calculate reachability propagating announcements through customers and peers, but not
  - Cloud's providers
  - Tier-1 ISPs
  - Tier-2 ISPs
- Reachability
  - If AS B receives a prefix announcement from the cloud, AS B is reachable by the cloud
  - Announcing AS called the *origin*
  - Uses augmented topology
  - Enforces common routing policies
- ***Hierarchy-free Reachability***
  - Count of reachable ASes when using peer links and not the hierarchical Internet

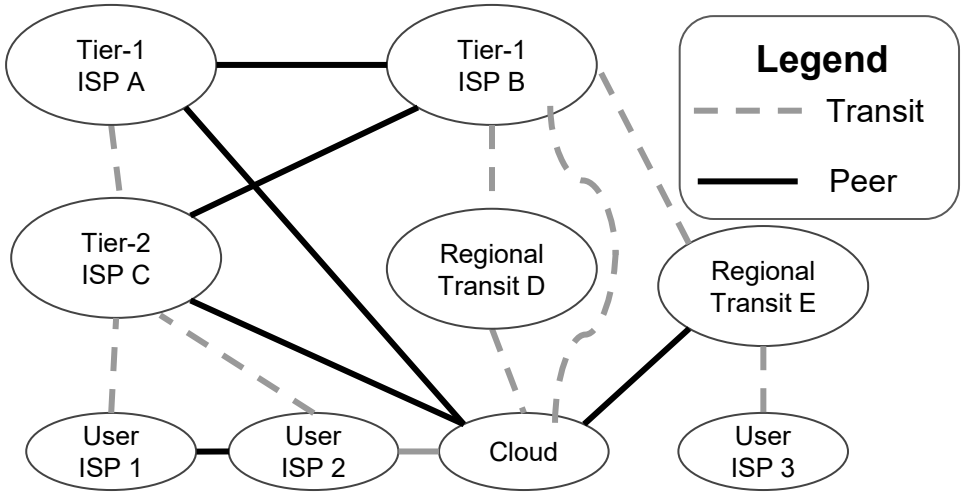

Do not propagate routes via
(1) Origin's transit providers
(2) Tier 1 ISPs
(3) Tier 2 ISPs

Calculating Hierarchy-free Reachability

- Calculate reachability propagating announcements through customers and peers, but not
  - Cloud's providers
  - Tier-1 ISPs
  - Tier-2 ISPs
- Reachability
  - If AS B receives a prefix announcement from the cloud, AS B is reachable by the cloud
  - Announcing AS called the *origin*
  - Uses augmented topology
  - Enforces common routing policies
- ***Hierarchy-free Reachability***
  - Count of reachable ASes when using peer links and not the hierarchical Internet

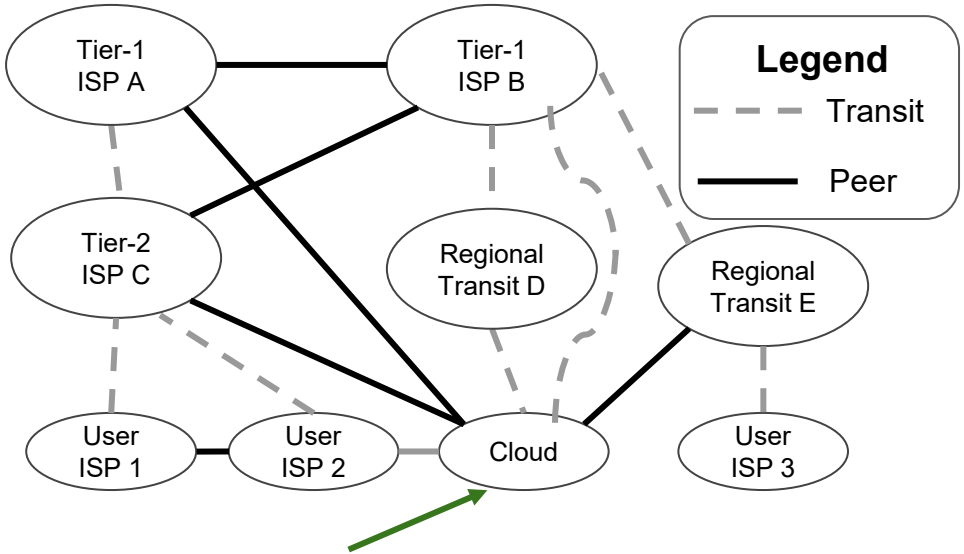

Do not propagate routes via
(1) Origin's transit providers
(2) Tier 1 ISPs
(3) Tier 2 ISPs

Calculating Hierarchy-free Reachability

- Calculate reachability propagating announcements through customers and peers, but not
  - Cloud's providers
  - Tier-1 ISPs
  - Tier-2 ISPs
- Reachability
  - If AS B receives a prefix announcement from the cloud, AS B is reachable by the cloud
  - Announcing AS called the *origin*
  - Uses augmented topology
  - Enforces common routing policies
- ***Hierarchy-free Reachability***
  - Count of reachable ASes when using peer links and not the hierarchical Internet

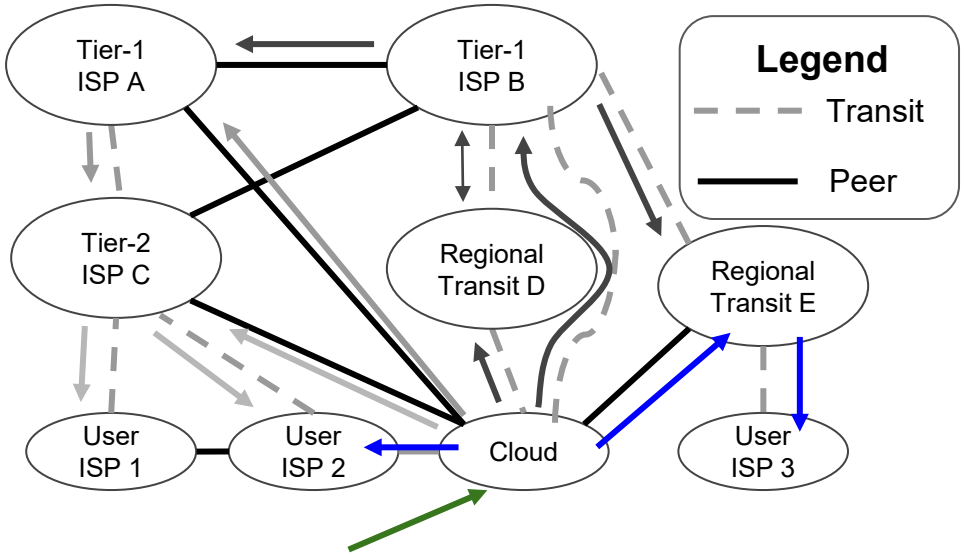

Do not propagate routes via
(1) Origin's transit providers
(2) Tier 1 ISPs
(3) Tier 2 ISPs

Calculating Hierarchy-free Reachability

- Calculate reachability propagating announcements through customers and peers, but not
  - Cloud's providers
  - Tier-1 ISPs
  - Tier-2 ISPs
- Reachability
  - If AS B receives a prefix announcement from the cloud, AS B is reachable by the cloud
  - Announcing AS called the *origin*
  - Uses augmented topology
  - Enforces common routing policies
- ***Hierarchy-free Reachability***
  - Count of reachable ASes when using peer links and not the hierarchical Internet

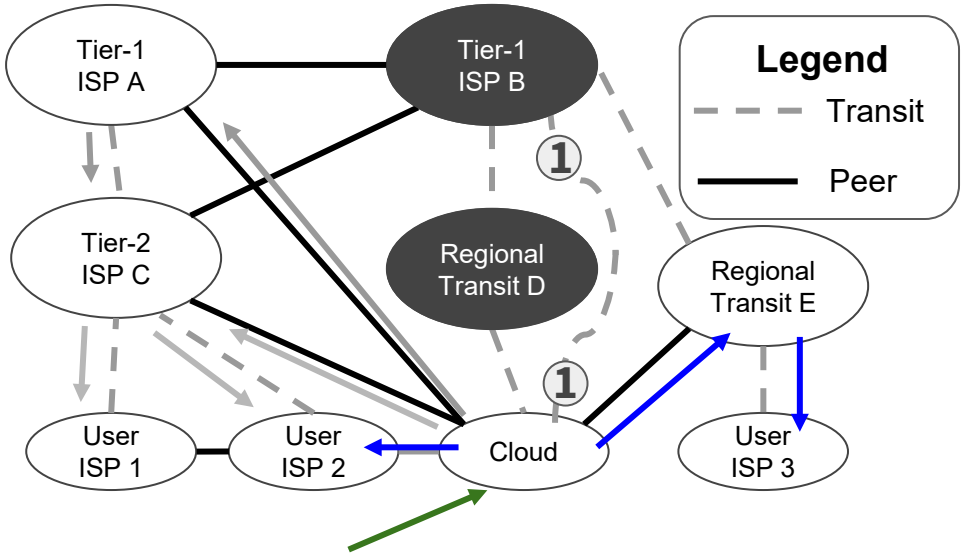

Do not propagate routes via
(1) Origin's transit providers
(2) Tier 1 ISPs
(3) Tier 2 ISPs

Calculating Hierarchy-free Reachability

- Calculate reachability propagating announcements through customers and peers, but not
  - Cloud's providers
  - Tier-1 ISPs
  - Tier-2 ISPs
- Reachability
  - If AS B receives a prefix announcement from the cloud, AS B is reachable by the cloud
  - Announcing AS called the *origin*
  - Uses augmented topology
  - Enforces common routing policies
- ***Hierarchy-free Reachability***
  - Count of reachable ASes when using peer links and not the hierarchical Internet

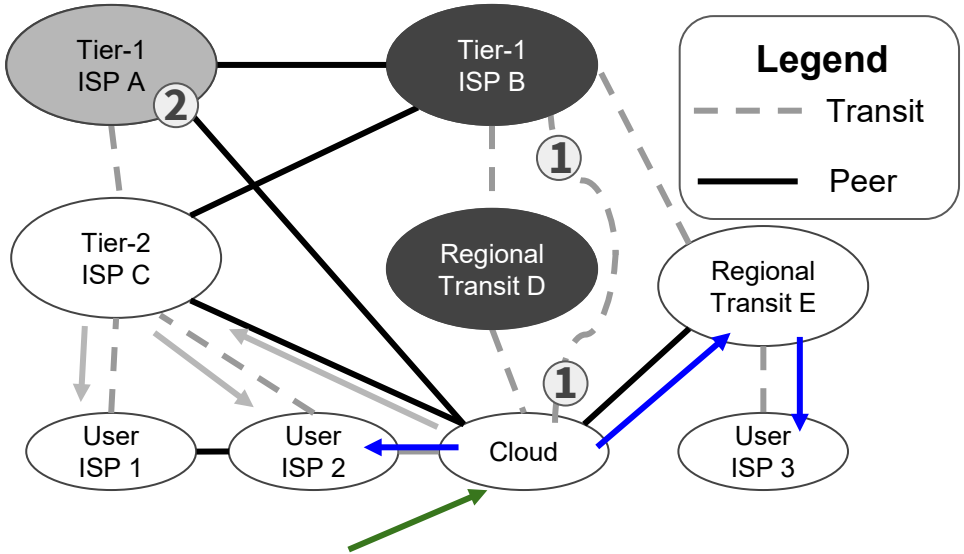

Do not propagate routes via
(1) Origin's transit providers
(2) Tier 1 ISPs
(3) Tier 2 ISPs

Calculating Hierarchy-free Reachability

- Calculate reachability propagating announcements through customers and peers, but not
  - Cloud's providers
  - Tier-1 ISPs
  - Tier-2 ISPs
- Reachability
  - If AS B receives a prefix announcement from the cloud, AS B is reachable by the cloud
  - Announcing AS called the *origin*
  - Uses augmented topology
  - Enforces common routing policies
- ***Hierarchy-free Reachability***
  - Count of reachable ASes when using peer links and not the hierarchical Internet

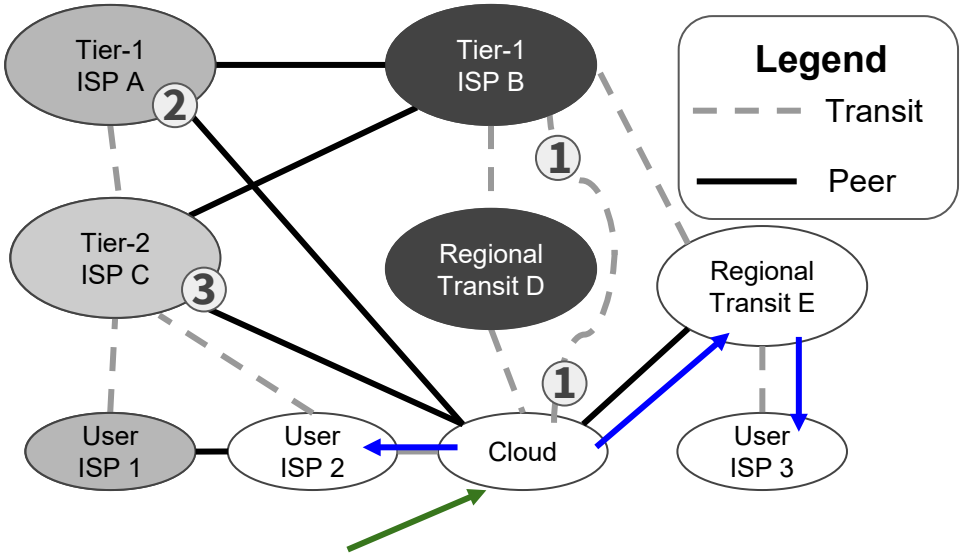

Do not propagate routes via
(1) Origin's transit providers
(2) Tier 1 ISPs
(3) Tier 2 ISPs

Hierarchy-free Reachability Results

Takeaway

- Cloud providers have higher reachability than most networks, including the Tier 1 and Tier 2 ISPs
- They are able to reach the majority of networks even when bypassing their transit providers, Tier 1 ISPs, and Tier 2 ISPs.
- Thousands of networks benefit from flattening

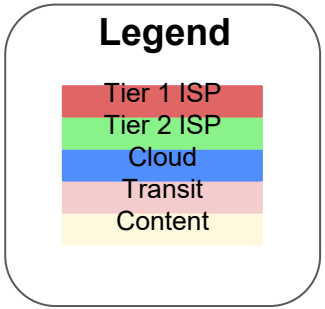

#	Network (AS)	Reachability (total, %)
1	Level 3 (3356)	61,154 (90.2%)
2	HE (6939)	58,981 (87.0%)
3	Google (15169)	58,922 (86.9%)
4	Microsoft (8075)	57,357 (84.6%)
5	IBM (36351)	55,714 (82.2%)
6	Cogent (174)	55,049 (81.2%)
7	Zayo (6461)	54,489 (80.4%)
8	Telia (1299)	54,324 (80.1%)
9	GTT (3257)	53,388 (78.7%)
10	SG.GS (24482)	53,157 (78.4%)
11	COLT (8220)	52,256 (77.1%)
12	G-Core Labs (199524)	51,820 (76.4%)
13	NTT (2914)	51,374 (75.8%)
14	Wikimedia (14907)	51,204 (75.5%)
15	Core-Backbone (33891)	51,110 (75.4%)
16	WV FIBER (19151)	51,083 (75.3%)
17	TELIN PT (7713)	50,919 (75.1%)
18	Amazon (16509)	50,867 (75.0%)
19	Swisscom (3303)	50,758 (74.9%)
20	IPTP (41095)	50,606 (74.6%)

Hierarchy-free Reachability Results

Takeaway

- Cloud providers have higher reachability than most networks, including the Tier 1 and Tier 2 ISPs
- They are able to reach the majority of networks even when bypassing their transit providers, Tier 1 ISPs, and Tier 2 ISPs.
- Thousands of networks benefit from flattening

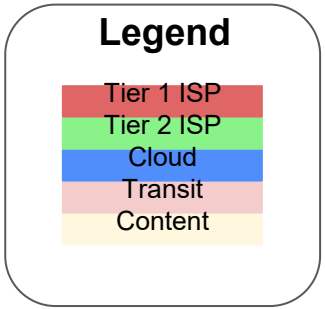

#	Network (AS)	Reachability (total, %)
1	Level 3 (3356)	61,154 (90.2%)
2	HE (6939)	58,981 (87.0%)
3	Google (15169)	58,922 (86.9%)
4	Microsoft (8075)	57,357 (84.6%)
5	IBM (36351)	55,714 (82.2%)
6	Cogent (174)	55,049 (81.2%)
7	Zayo (6461)	54,489 (80.4%)
8	Telia (1299)	54,324 (80.1%)
9	GTT (3257)	53,388 (78.7%)
10	SG.GS (24482)	53,157 (78.4%)
11	COLT (8220)	52,256 (77.1%)
12	G-Core Labs (199524)	51,820 (76.4%)
13	NTT (2914)	51,374 (75.8%)
14	Wikimedia (14907)	51,204 (75.5%)
15	Core-Backbone (33891)	51,110 (75.4%)
16	WV FIBER (19151)	51,083 (75.3%)
17	TELIN PT (7713)	50,919 (75.1%)
18	Amazon (16509)	50,867 (75.0%)
19	Swisscom (3303)	50,758 (74.9%)
20	IPTP (41095)	50,606 (74.6%)

Conclusions

- Emulated connectivity using an AS-level topology graph constructed from
  - BGP data
  - Traceroutes
  - Validated cloud neighbor lists
- Hierarchy-free Reachability quantifies the extent of Internet flattening and how little networks rely on the Internet hierarchy
- Show that thousands of networks benefit from flattening
  - Other metrics do not capture these insights
  - The cloud providers rely less on the hierarchy than most other networks