




Mutual Aid
or an routed odyssey

Midwest Peering Summit

St. Paul, Minnesota - July 15th, 2025

Mutual aid



... is a form of voluntary reciprocal exchange of resources and services for mutual benefit, often within a community

Waves of innovation

Org structures

- Local/regional community, DBA
- Datacenter/meet-me room operator
- Non-profit'ish/sort-a-kinda commercial
 - *Big 3 Euro's, industry/trade associations*
- Franchise IX's
 - *Domestic incumbents “solved” market, expand to Africa, Asia, & South America*
- IX-as-a-*{ancillary}*-service
 - Megaport, PacketFabric, etc

Route Servers

- Limited / if any filtering
- Max. accepted prefixes
 - *Static* → *PeeringDB API*
- Blackhole sink
- IRR filtering
 - “Authenticated” IRR’s
- MANRS adoption
 - Known IRR and/or RPKI
- [Web interface](#) w/ annotations “reason(s)”
- Published RS filter policies, [ARouteServer](#)
- Tag prefixes w/ geographical identifier




Learned from Colo: scl02 (2972 Stender Way) (63055:1984:5)

Participant Type: Exempt (63055:1900:2)

Origin ASN is included in client's AS-SET (63055:64512:21)

{DNS} Caches

- PCH, ISC, ICANN, Netnod, RIPE
- Versign, Affilias, etc



Troubleshooting tools

- RouteViews, bgp.he.net, NLNOG RING, bgp.tools, etc

Summary: IX's all up in your route decisions

Mixed success

- Web portal - IXP manager or custom
- Mobile app and/or IX API
- Layer 2 transport (*“Private VLAN’s”*)
- Outreach / policy making outside of our constituents



Typical 2025 NA community IX

- Robust L2 fabric w/o STP, VXLAN/EVPN
- Multi-site fabrics (*campus and/or region*)
- Cross-connection MRC ... \$\$\$
- Difficult caches: Akamai, Google
- Growing caches: Netskrt, Valve, Apple
- New BYO caches: Microsoft MCC
- Growth from regional “access networks”
ISP’s - lots of M&A, state & federal grants
 - Cloud/content steady state - some pulling back from IX, pushing caches

2024 SFMIX scene

Content moves to {off IX} peering and/or caches

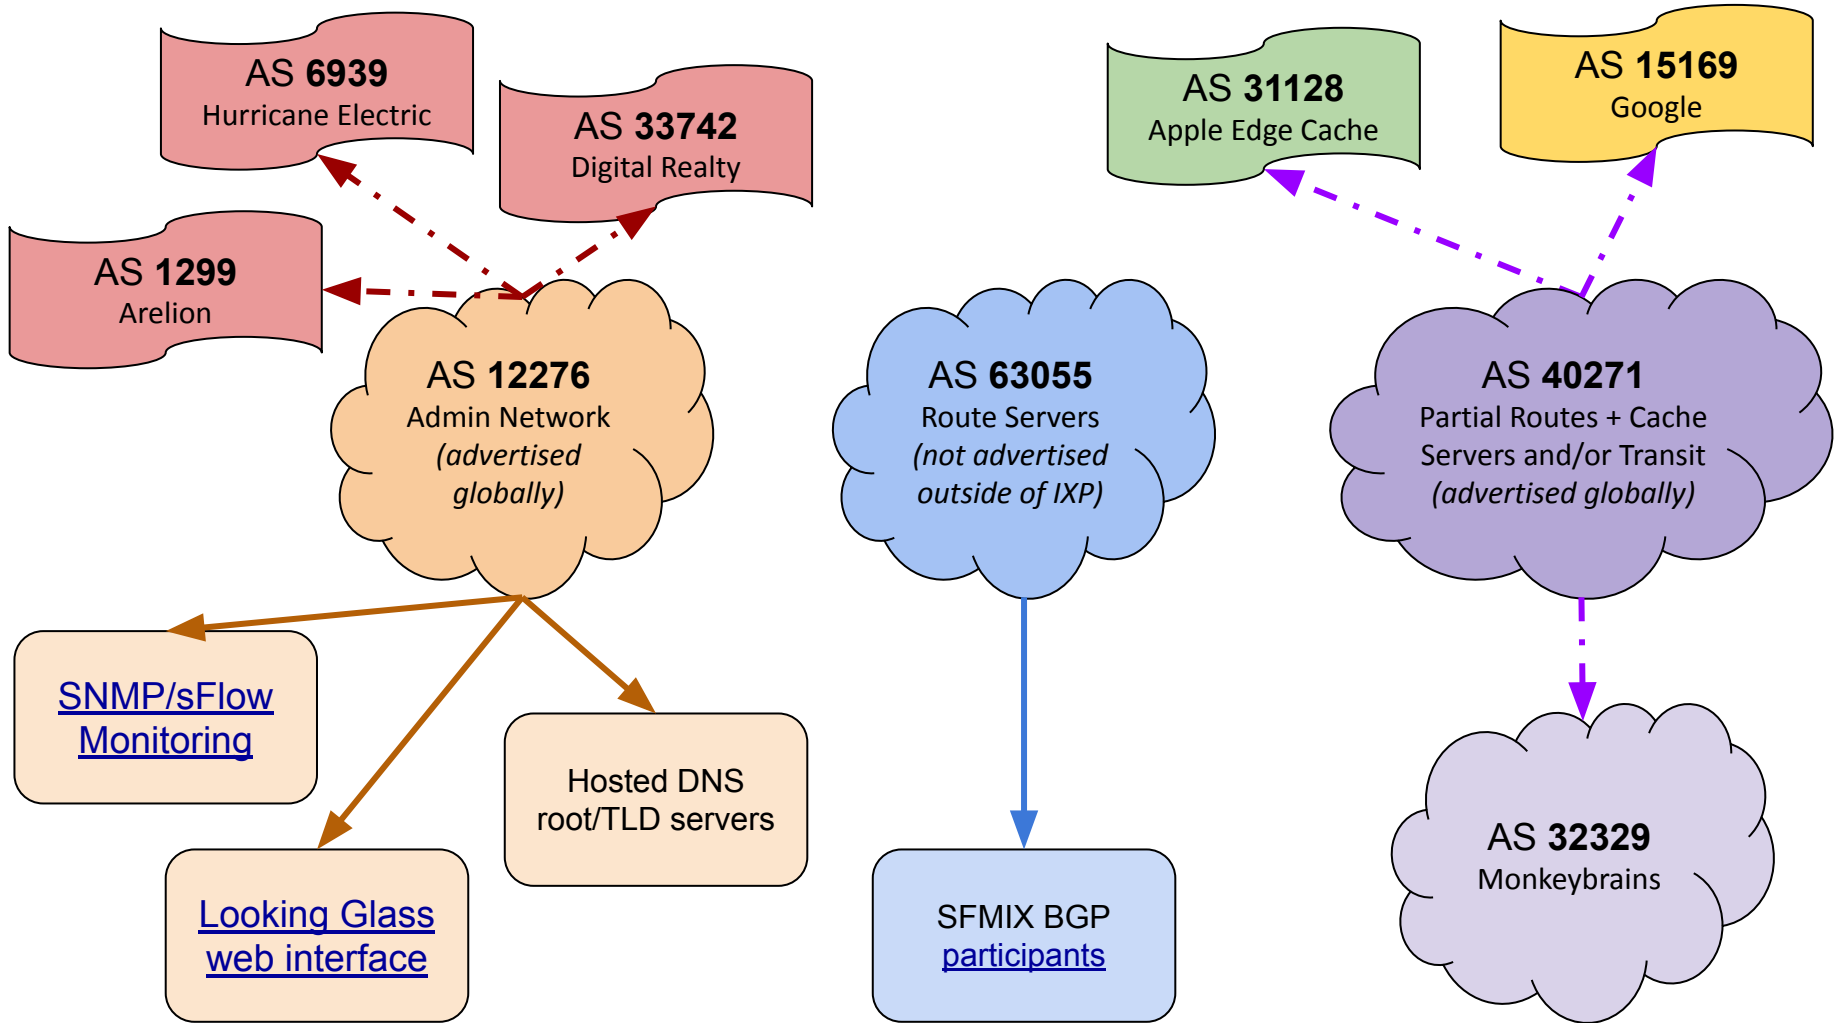
- Regional ISP lost major content peer over IX
 - ISP triggered undisclosed de-peer criteria
 - Content de-peer RS's a year prior
 - Peer is valuable to ISP, vocal customers
 - Content offers PNI, no mutual location
 - ISP escalates to IX for help
- IX not geared to transport PNI's
 - No SLA, tech support overhead
 - Compete against membership
- Content peer did offer PNI to the IX, *if an sourced with an ASN not on the IX...*

2024 SFMIX scene

Proof of concept

- Spin up “aggregator” ASN, get some v4/v6 PI
- Recycle existing Arista, limitations found
- Desire to re-use existing cross-x where possible, add VLAN tag - give more value
- Using EOL Juniper MX (*rib-groups w/ BGP signalling for selective VRF leaking*)
- Started adding caches: Apple, Netskrt, Valve
 - (*refused to join IX as normal member and/or not present in region*)

You get an AS



VRF leak'age

Content “🍏 cache”:

- Receive opted-in-ISP routes
- Not interested in other caches
- Static default route (*no DFZ*)

ISP “A”

- Wants cache + private peering

ISP “B”

- Wants cache + full table?
- No {duplicate} private peering

ISP “C”

- Wants private peering
- No caches



Current status



- Value *{for subset of IX users}* is real:
 - Single ISP filled 20Gbps day 1, scaled back POC due to capacity constraints
 - Convention center: *“This solves so many problems, been waiting for this..”*
 - ISP: *“We already trust you, the IX is more reliable than our transits, do it..”*
 - Recent survey (*~17% of total IX user base*) interest, commits cover CapEx in *~2 yrs*
- Board guidance very KISS - flat rate, 10Gbps \$300/MRC, no SLA, cover costs

Current status

- Negotiation for 2x 400Gbps capable routers
 - Bonus: Inter-op w/ IX VXLAN underlay
 - Challenge: 400G optic form factors
- Strong desire for separation, control-plane
 - Delivery over existing IX cross-x
- Unanswered: Caches need bursty transit..
 - Dark fiber supplies & DC's: don't care
 - One transit participant: *If >50Gbps, gotta start charging ya for cabinet* - totally fair



**“what you’re talking about doing is
literally the smartest thing I’ve
heard in IX land in forever”**

- someone in the audience

matt@sfmix.org

thank you