

IXP Manager & Route Servers



Route Servers Video Tutorial Series – Part 1

Barry O'Donovan – [@barryo79](#), June 2021

Auto-Provisioning in IXP Manager

When a interface is added to IXP Manager, you (can) get:


- Route Collector BGP session auto-provisioned
- Secure Route Server BGP sessions auto-provisioned
- MRTG / graphing auto-provisioned
- Peer to peer graphs auto-provisioned
- Smokeping target for member's interface
- Nagios monitoring of member's interface
- AS112 BGP session
- ARPA DNS for IXP assigned address
- RIR AS-SET / ASN objects
- And more...

Route Servers


- a third-party brokering system providing multilateral interconnection via BGP
- BGP sessions required for full mesh peering:

$$\frac{n(n-1)}{2}$$

- 10 participants: 45 sessions
- 100 participants: 4,950 sessions



IXP full mesh peering relationships



IXP route server peering relationships

Route Servers

- routes announced by one [route server client](#) are forwarded to all other participating clients
- **does not** forward traffic and is not a router
- [rfc7947](#) and [rfc7948](#), [Hilliard et al](#):

"the overhead associated with dense interconnection can cause substantial operational scaling problems for participants of larger IXPs"

Route Servers

- Considered a production level service at IXPs
 - Stability, reliability, consistency
- Threats can be malicious or accidental
- Threats include:
 - Route leak (DFZ or targeted network)
 - Next hop hijacking
 - Route server software bugs

Route Servers - Hardware

For a IX carrying **significant** traffic, ideally:

- Two servers
- Dual PSU
- ≥ 4 core CPU, ≥ 16 GB RAM
- Storage requirements moderate (< 30 GB / VM)
- Virtualisation useful for maintenance, lab work, etc.
- Direct connection to peering LAN and separate management interface

Route Servers - Software

For a IX carrying **significant** traffic, ideally:

- Two different software stacks
 - Or cautious approach to software updates
- Idempotent Atomic Session-Based Configuration Merge
 - Essentially: supports complete replacement of configuration and reload
- Support for efficient handling of large prefix sets; RPKI; large communities; etc.
- Bird2 fully supported. Plan to support OpenBGPD.

Route Servers & IXP Manager

- INEX has operated route servers since 2007
- INEX CTO co-authored the RFCs
- We have given presentations, workshops, tutorials
- Always automated - never deployed manually
- Always secured with prefix filtering

This experience and knowledge has been distilled into IXP Manager

Subject: [inex-tech] Route server system now in beta

Date: Fri Nov 23 12:20:17 GMT 2007

From: Nick Hilliard <email@inex.ie>

Following the announcement at the last INEX members meeting that we were looking into running a route server system, we are now pleased to announce that we now have a route server system which is in stable beta.

As a brief summary, the route server system offers the following advantages:

- dramatically reduces the number of BGP sessions required to peer with other INEX members
- strict route filtering on inbound announcements means that only prefixes registered at RIPE by exchange members will be visible
- dual-hosted system offers high reliability
- community based filtering allows route server users to control which INEX members their prefixes are sent to

Well-Known Community Filters

Allowing members to control prefix propagation is a critical feature to drive route server usage.

```
router bgp 65503
  address family ipv4
    neighbor 192.0.2.8 remote-as 65501
    neighbor 192.0.2.8 route-map ix-router-server-out out
    neighbor 192.0.2.8 send-community
!
route-map ix-router-server-out
  set community 65501:65502
```

More information on <https://docs.ixpmanager.org/>

Well-Known Community Filters - Standard

Description	Community
Prevent announcement of a prefix to a peer	0:peer-as
Announce a route to a certain peer	rs-asn:peer-as
Prevent announcement of a prefix to all peers	0:rs-asn
Announce a route to all peers	rs-asn:rs-asn

- No one except X, Y and Z: 0:rs-asn rs-asn:X rs-asn:Y rs-asn:Z
- Everyone but X and Y: 0:X 0:Y

Well-Known Community Filters - Large Communities

Description	Community
Prevent announcement of a prefix to a peer	rs-asn:0:peer-as
Announce a route to a certain peer	rs-asn:1:peer-as
Prevent announcement of a prefix to all peers	rs-asn:0:0
Announce a route to all peers	rs-asn:1:0

- Standard filtering does not work with 32-bit ASNs
- A route server client should not mix standard 16-bit communities and large communities – please choose one or the other

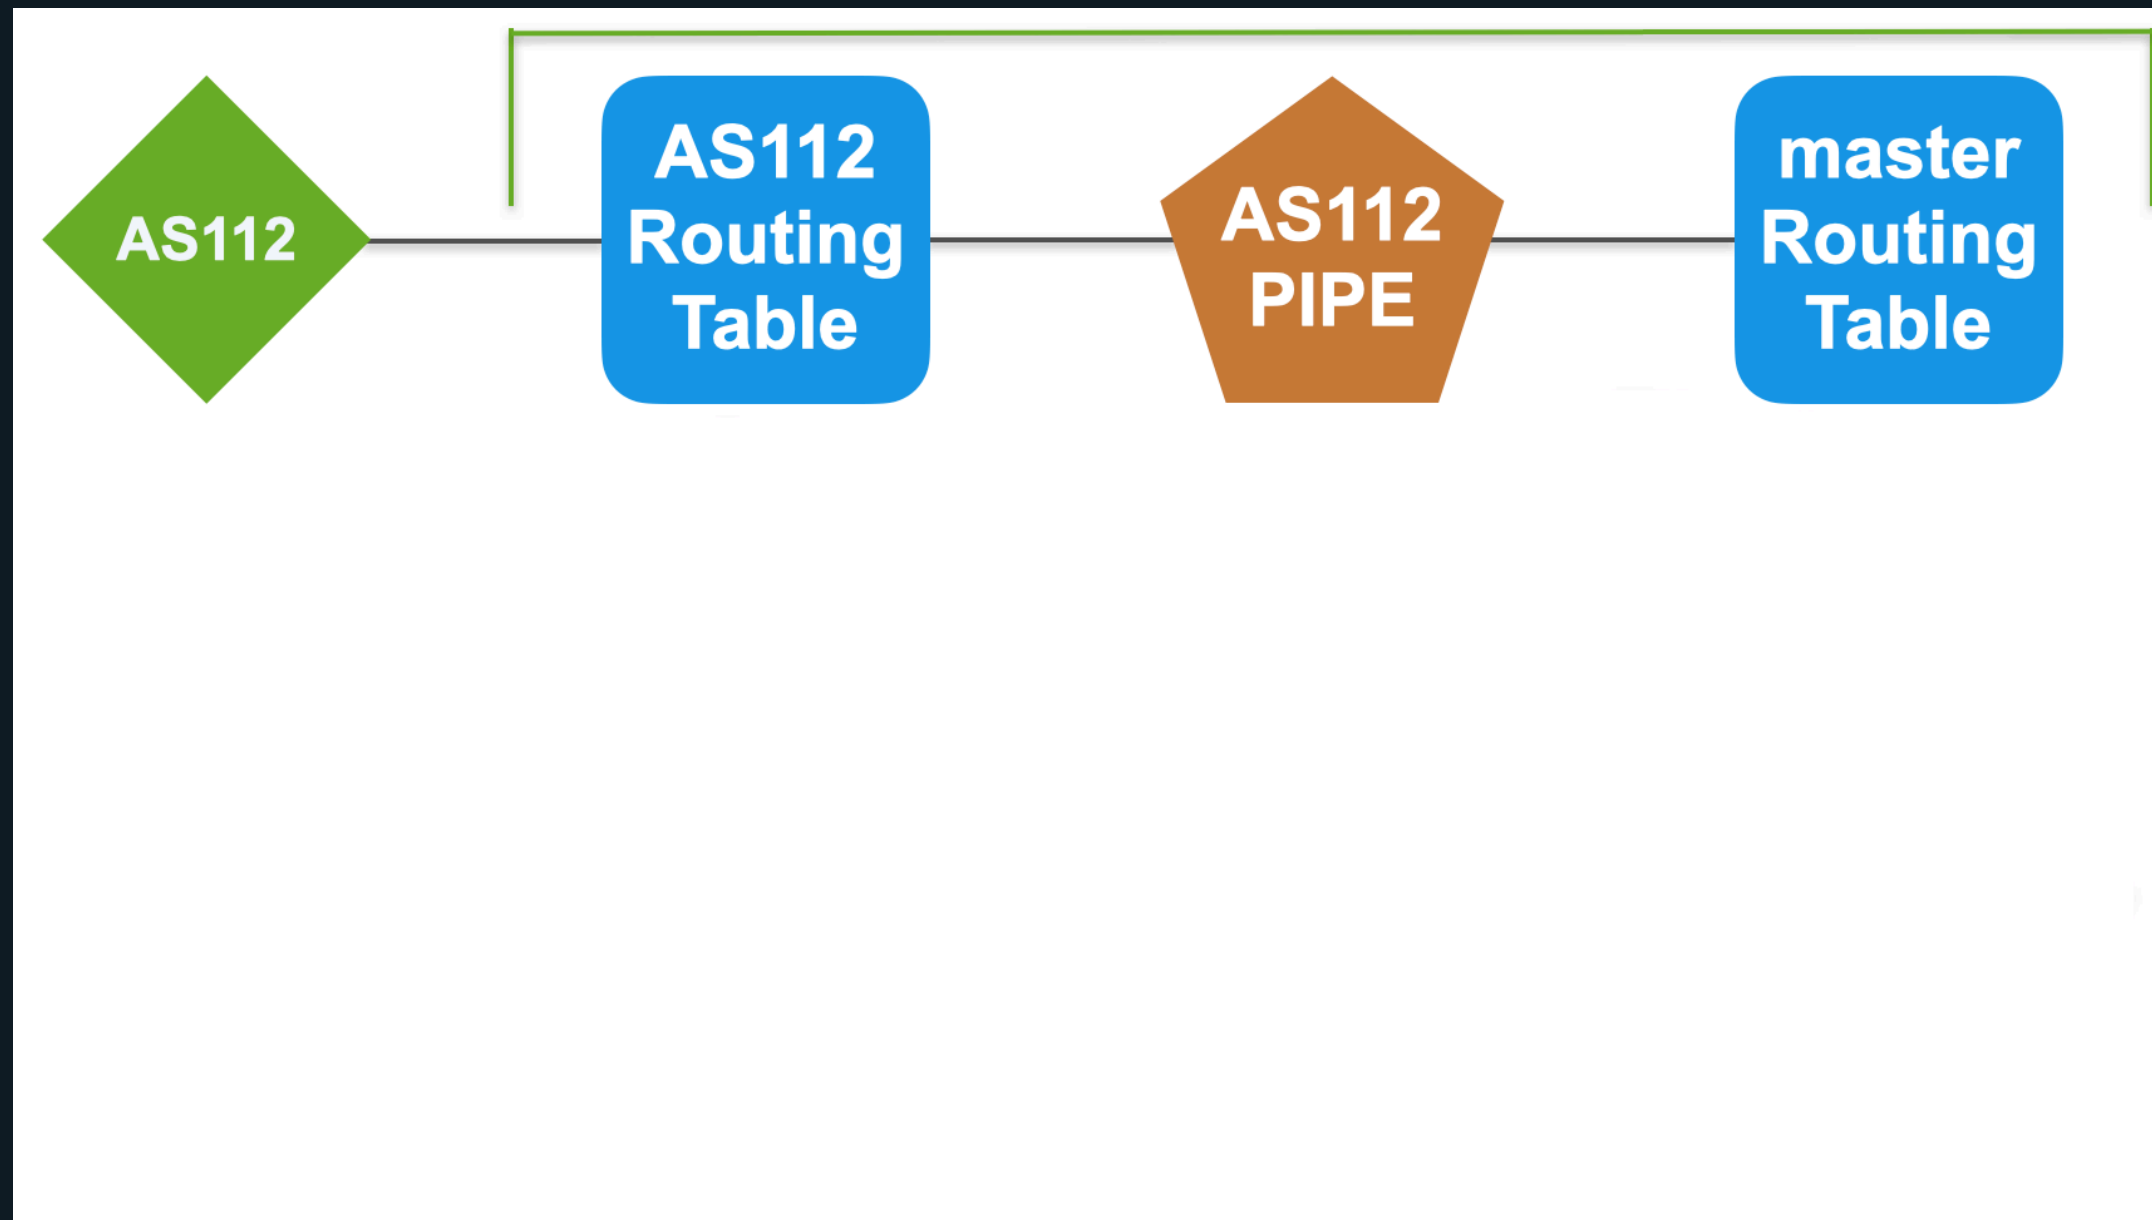
Well-Known Community Filters - Large Communities

Bonus

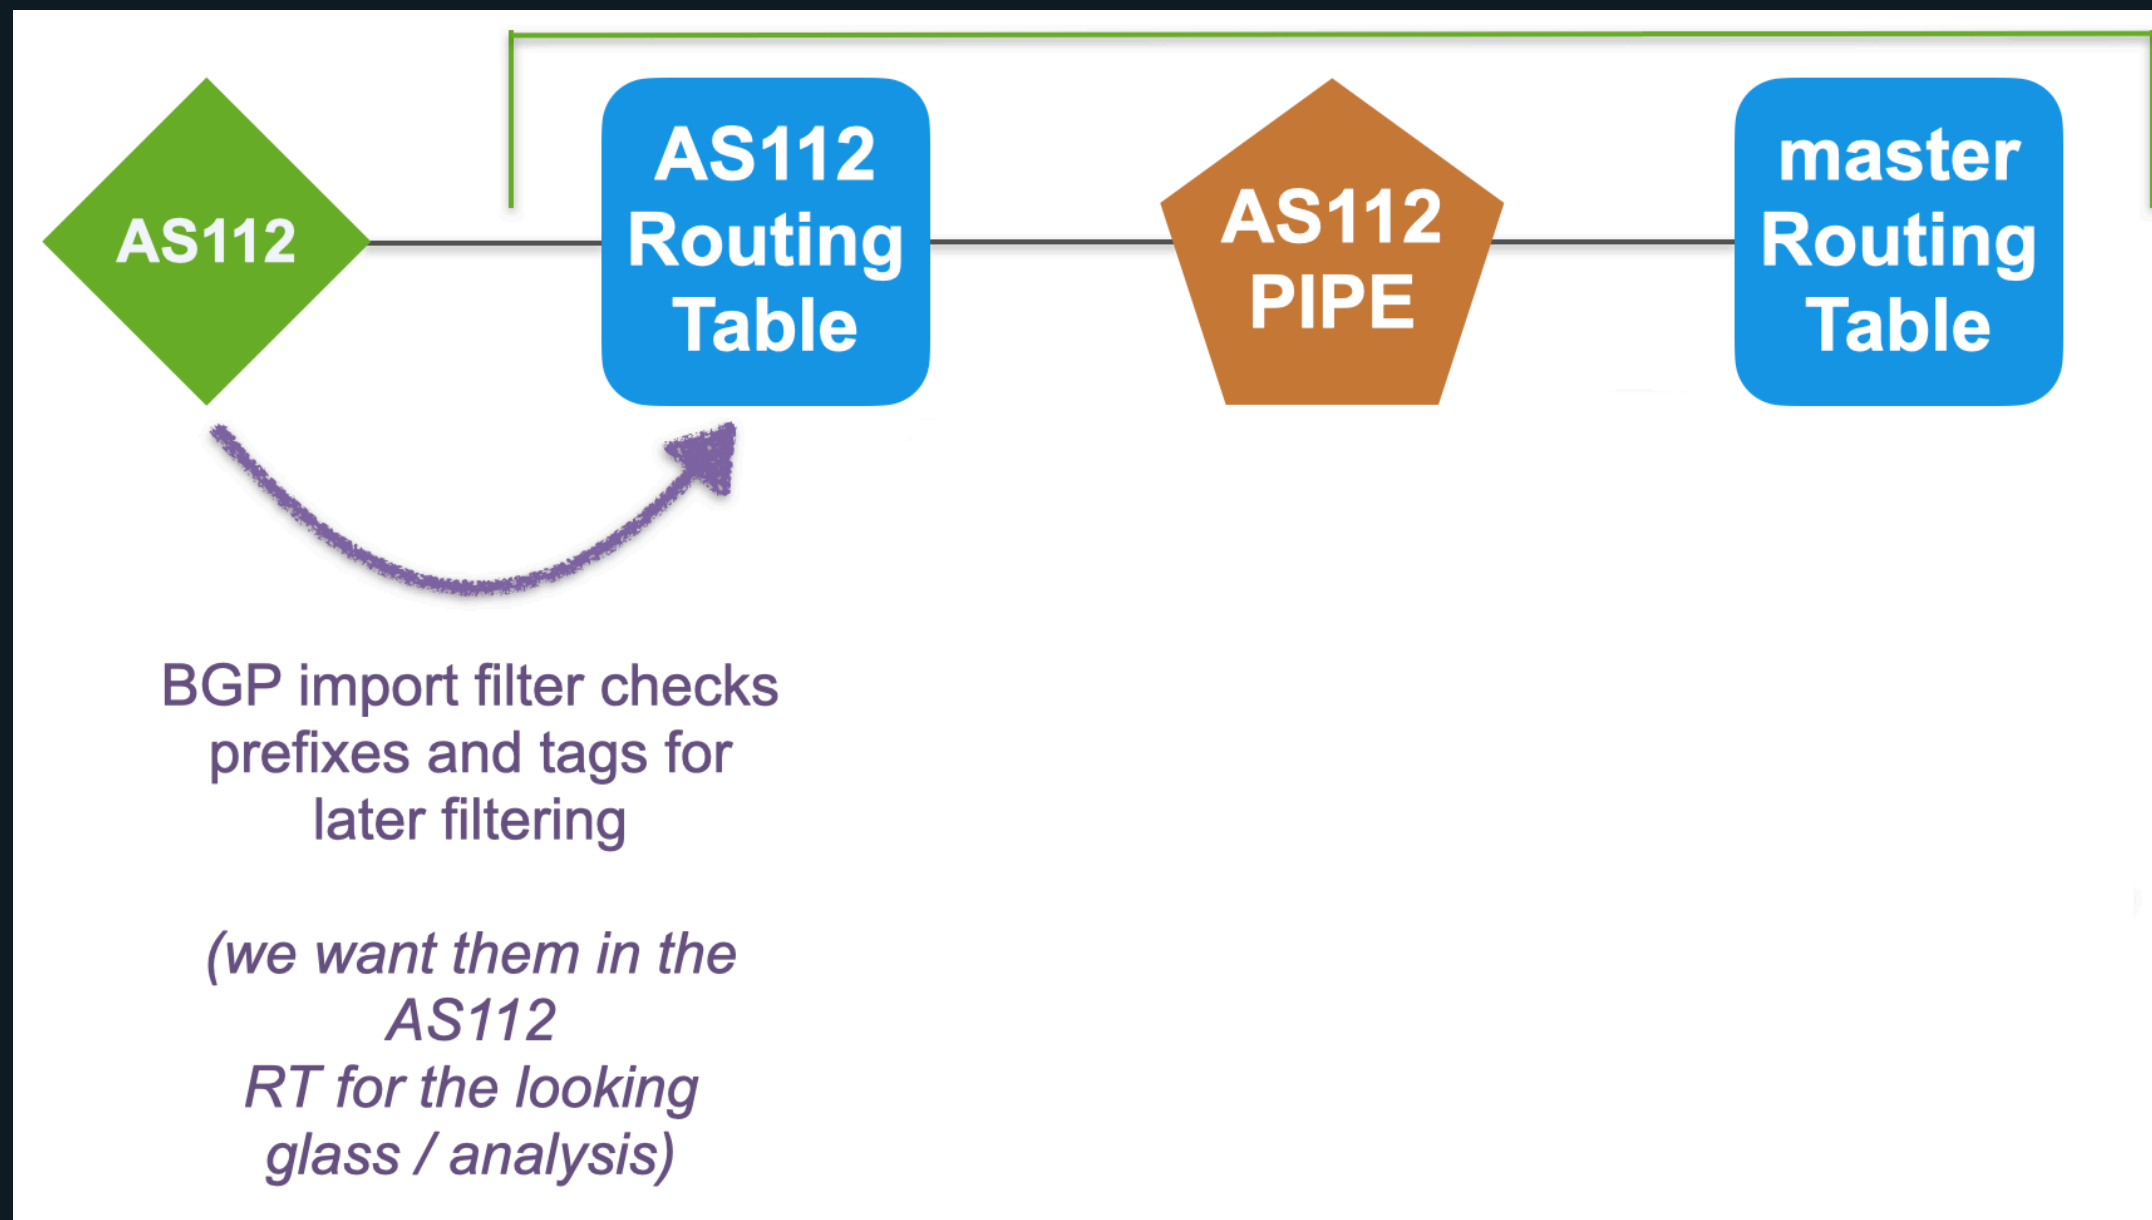
Description	Community
Prepend to peer AS once	rs-asn:101:peer-as
Prepend to peer AS twice	rs-asn:102:peer-as
Prepend to peer AS thrice	rs-asn:103:peer-as

NB: communities control propagation of a client's routes to other clients. Clients responsible for filtering inbound themselves.


IXP Manager Bird Topology - Import from Client



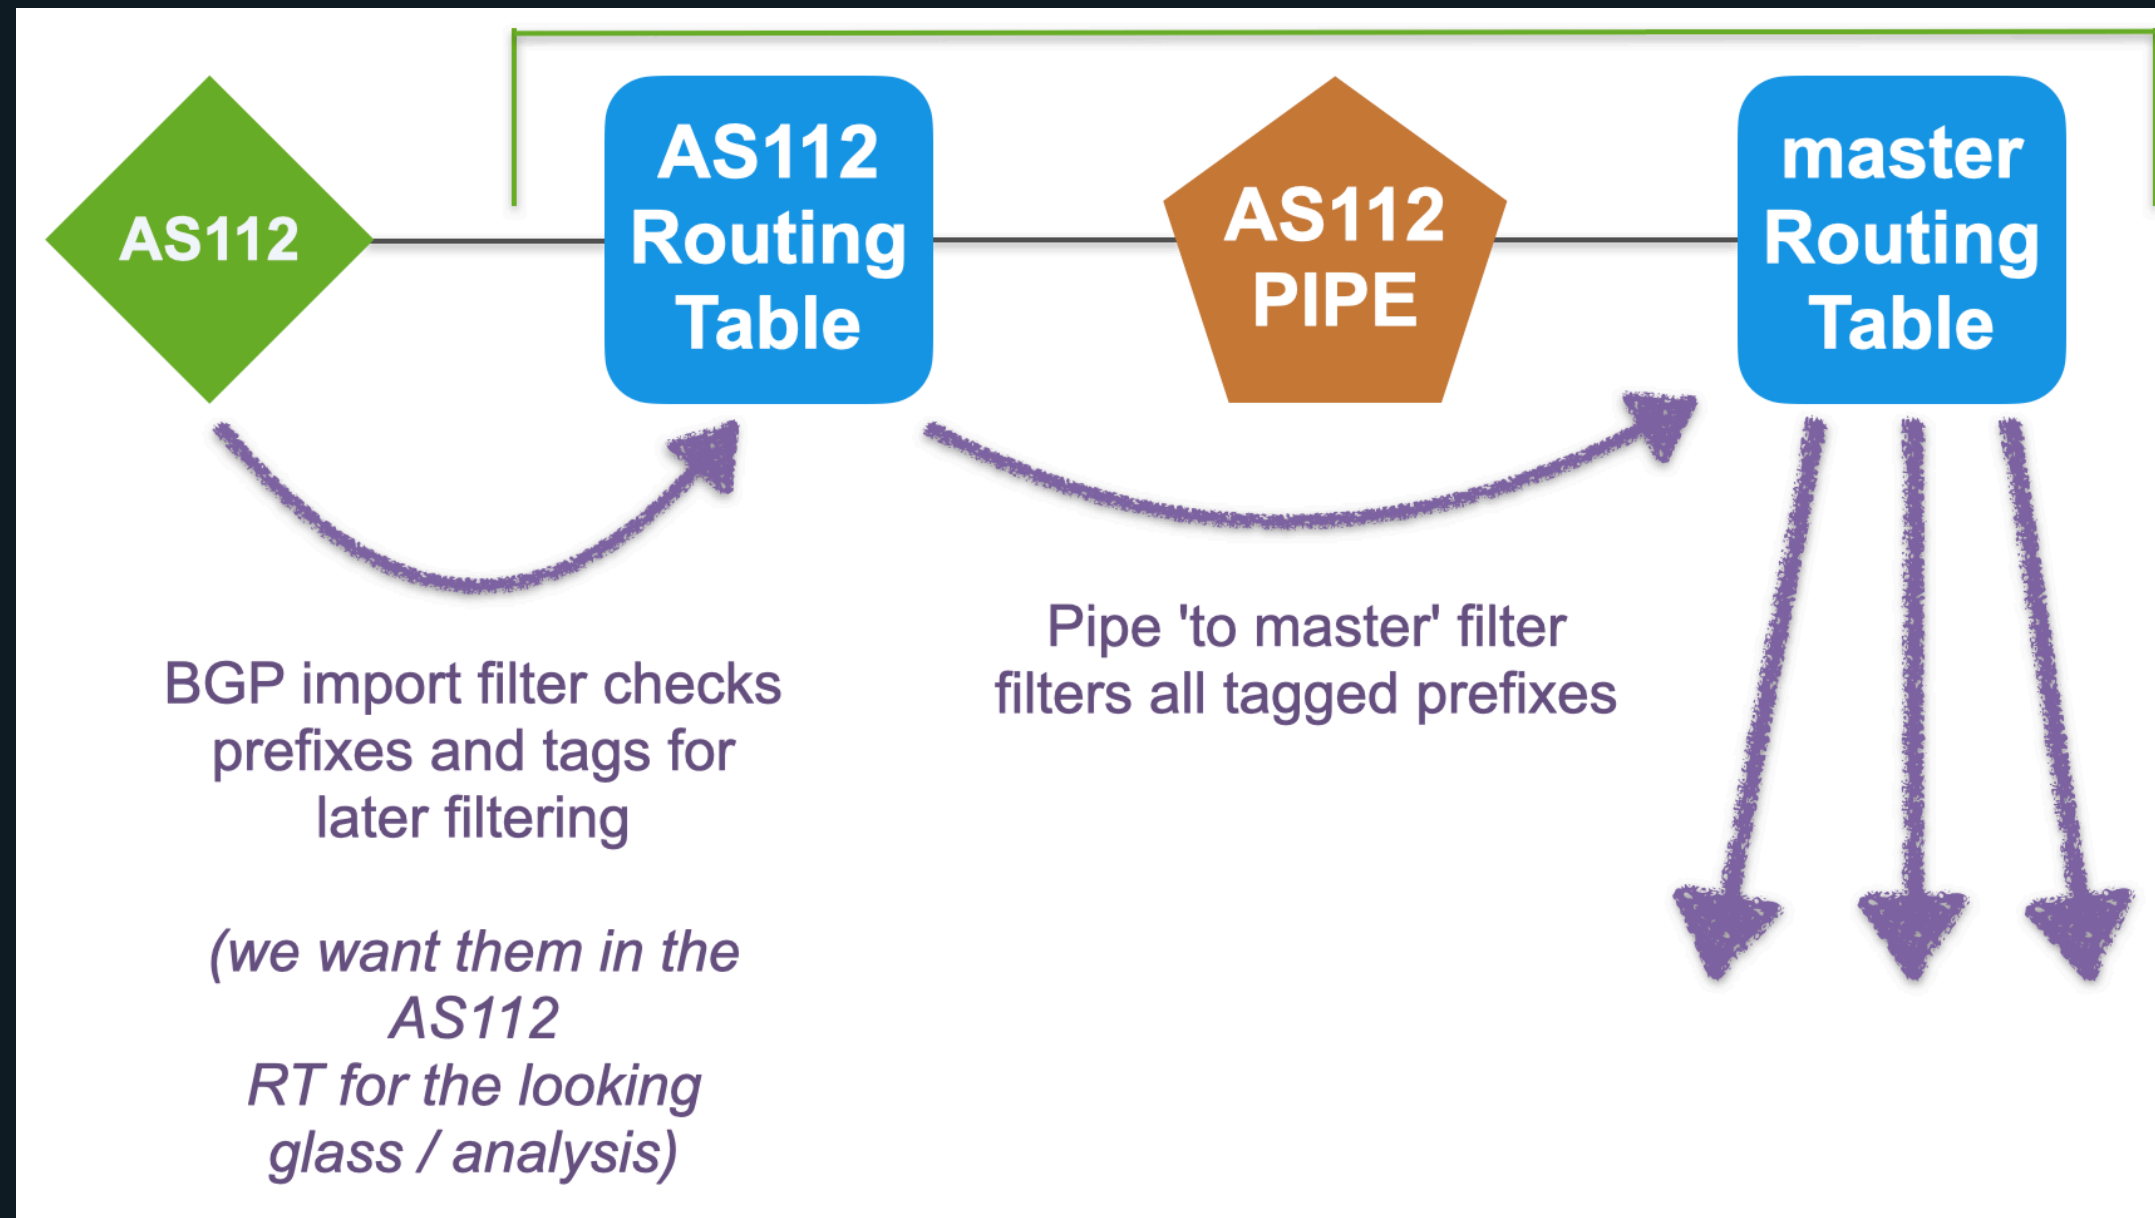
IXP Manager Bird Topology - Import from Client




IXP Manager Bird Topology - Import from Client




IXP Manager Bird Topology - Import from Client



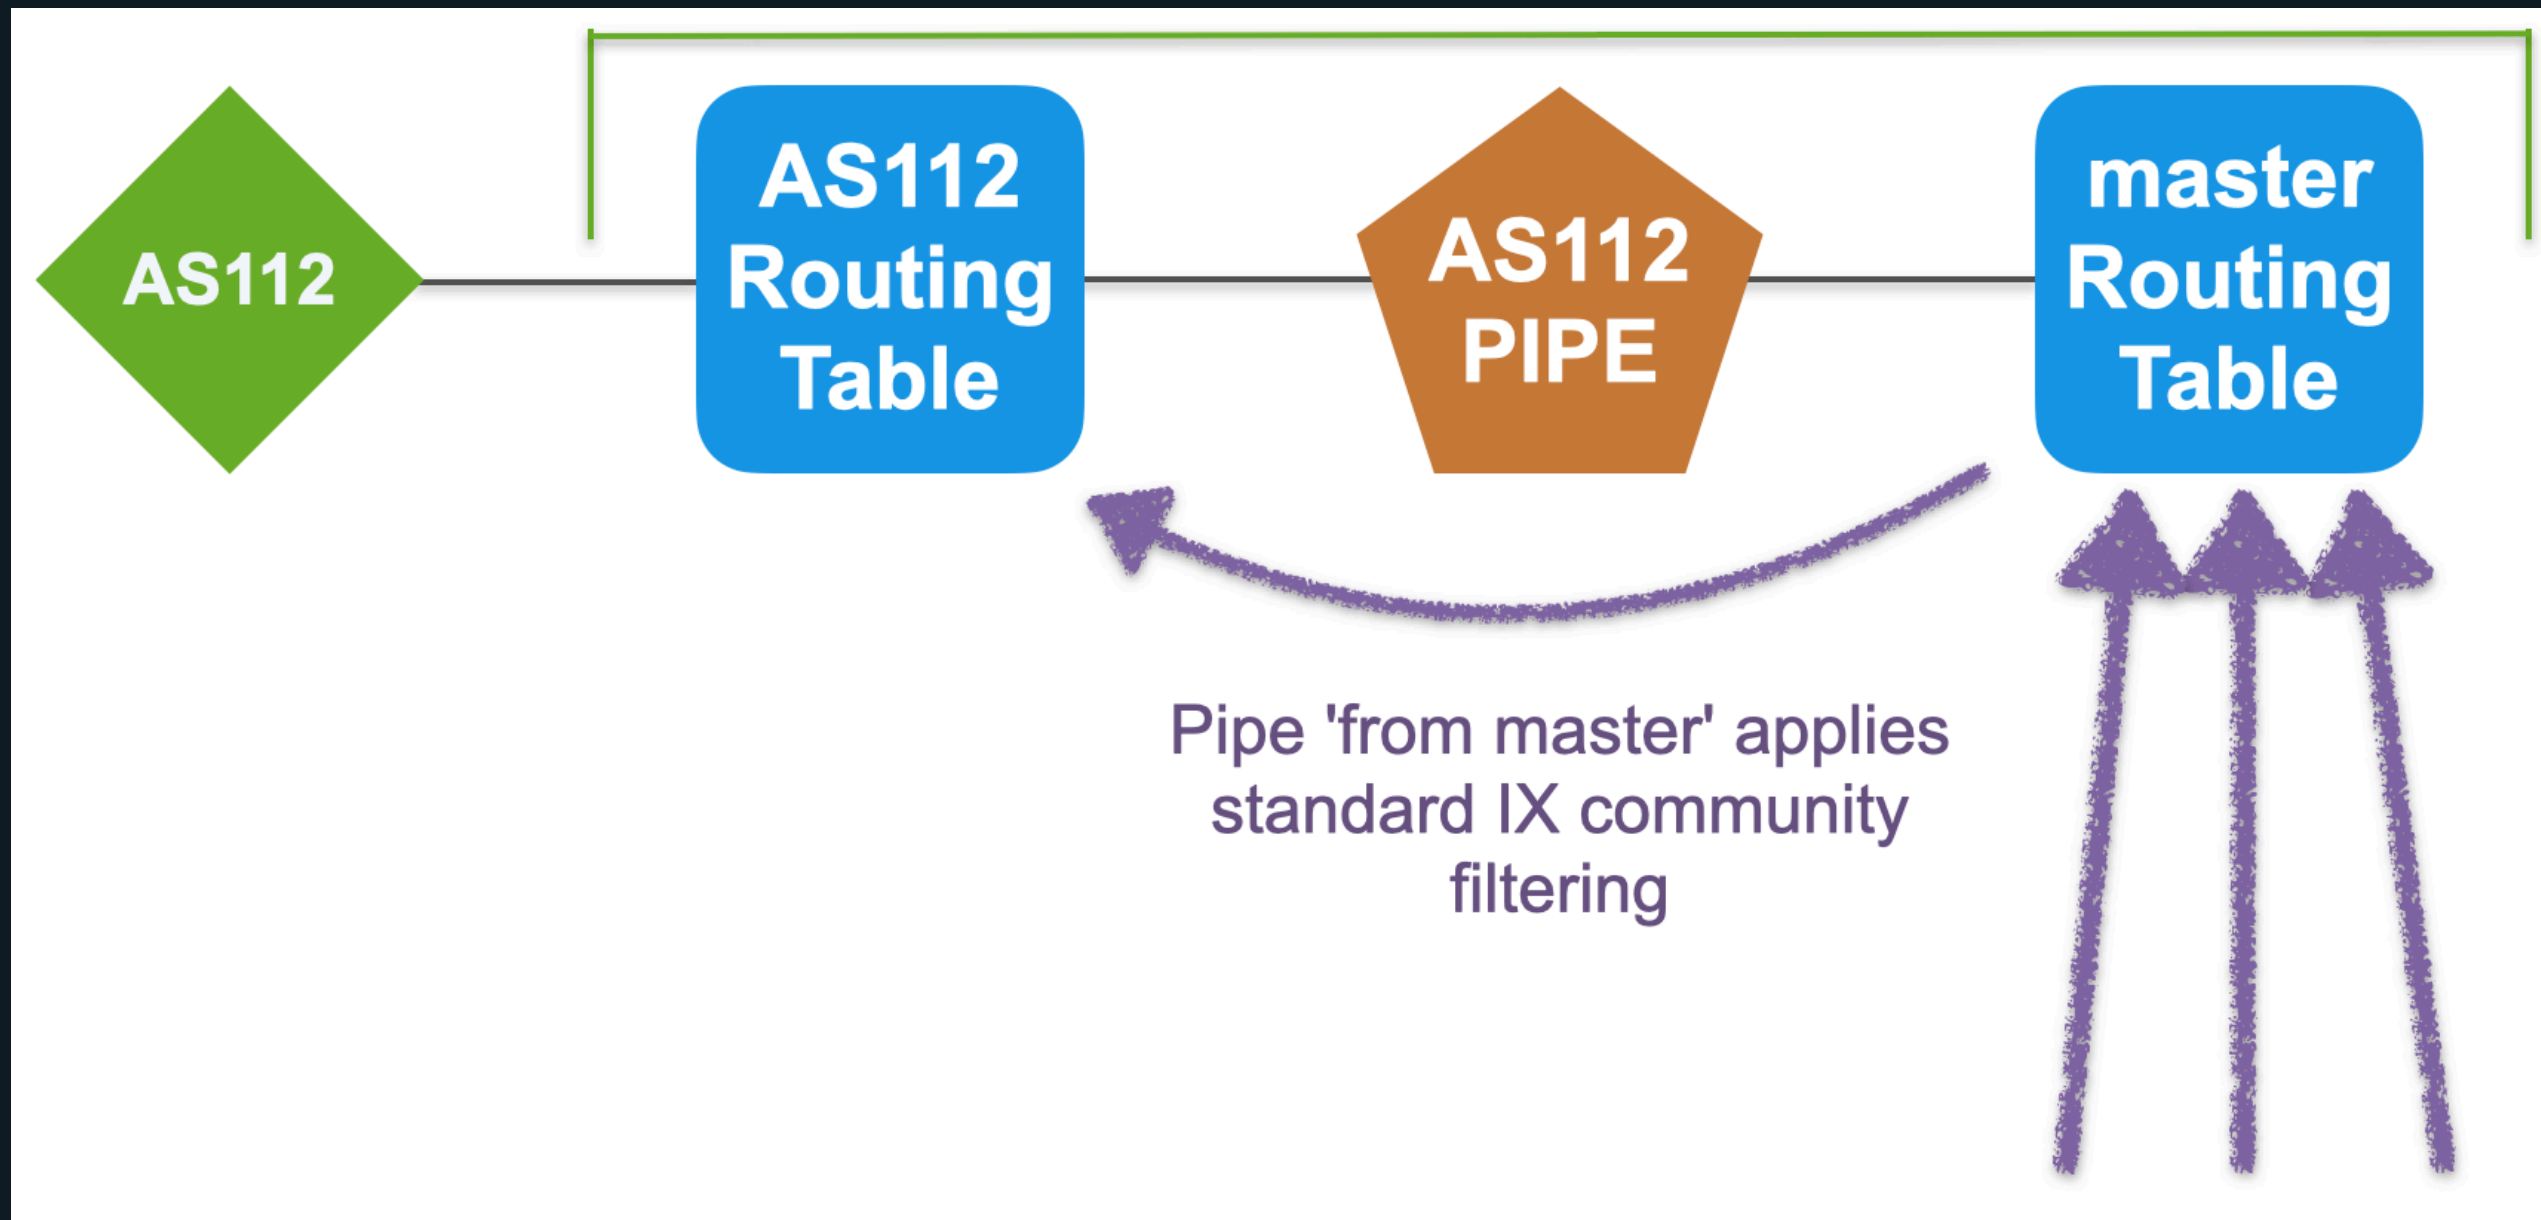
IXP Manager Bird Topology - Export to Client



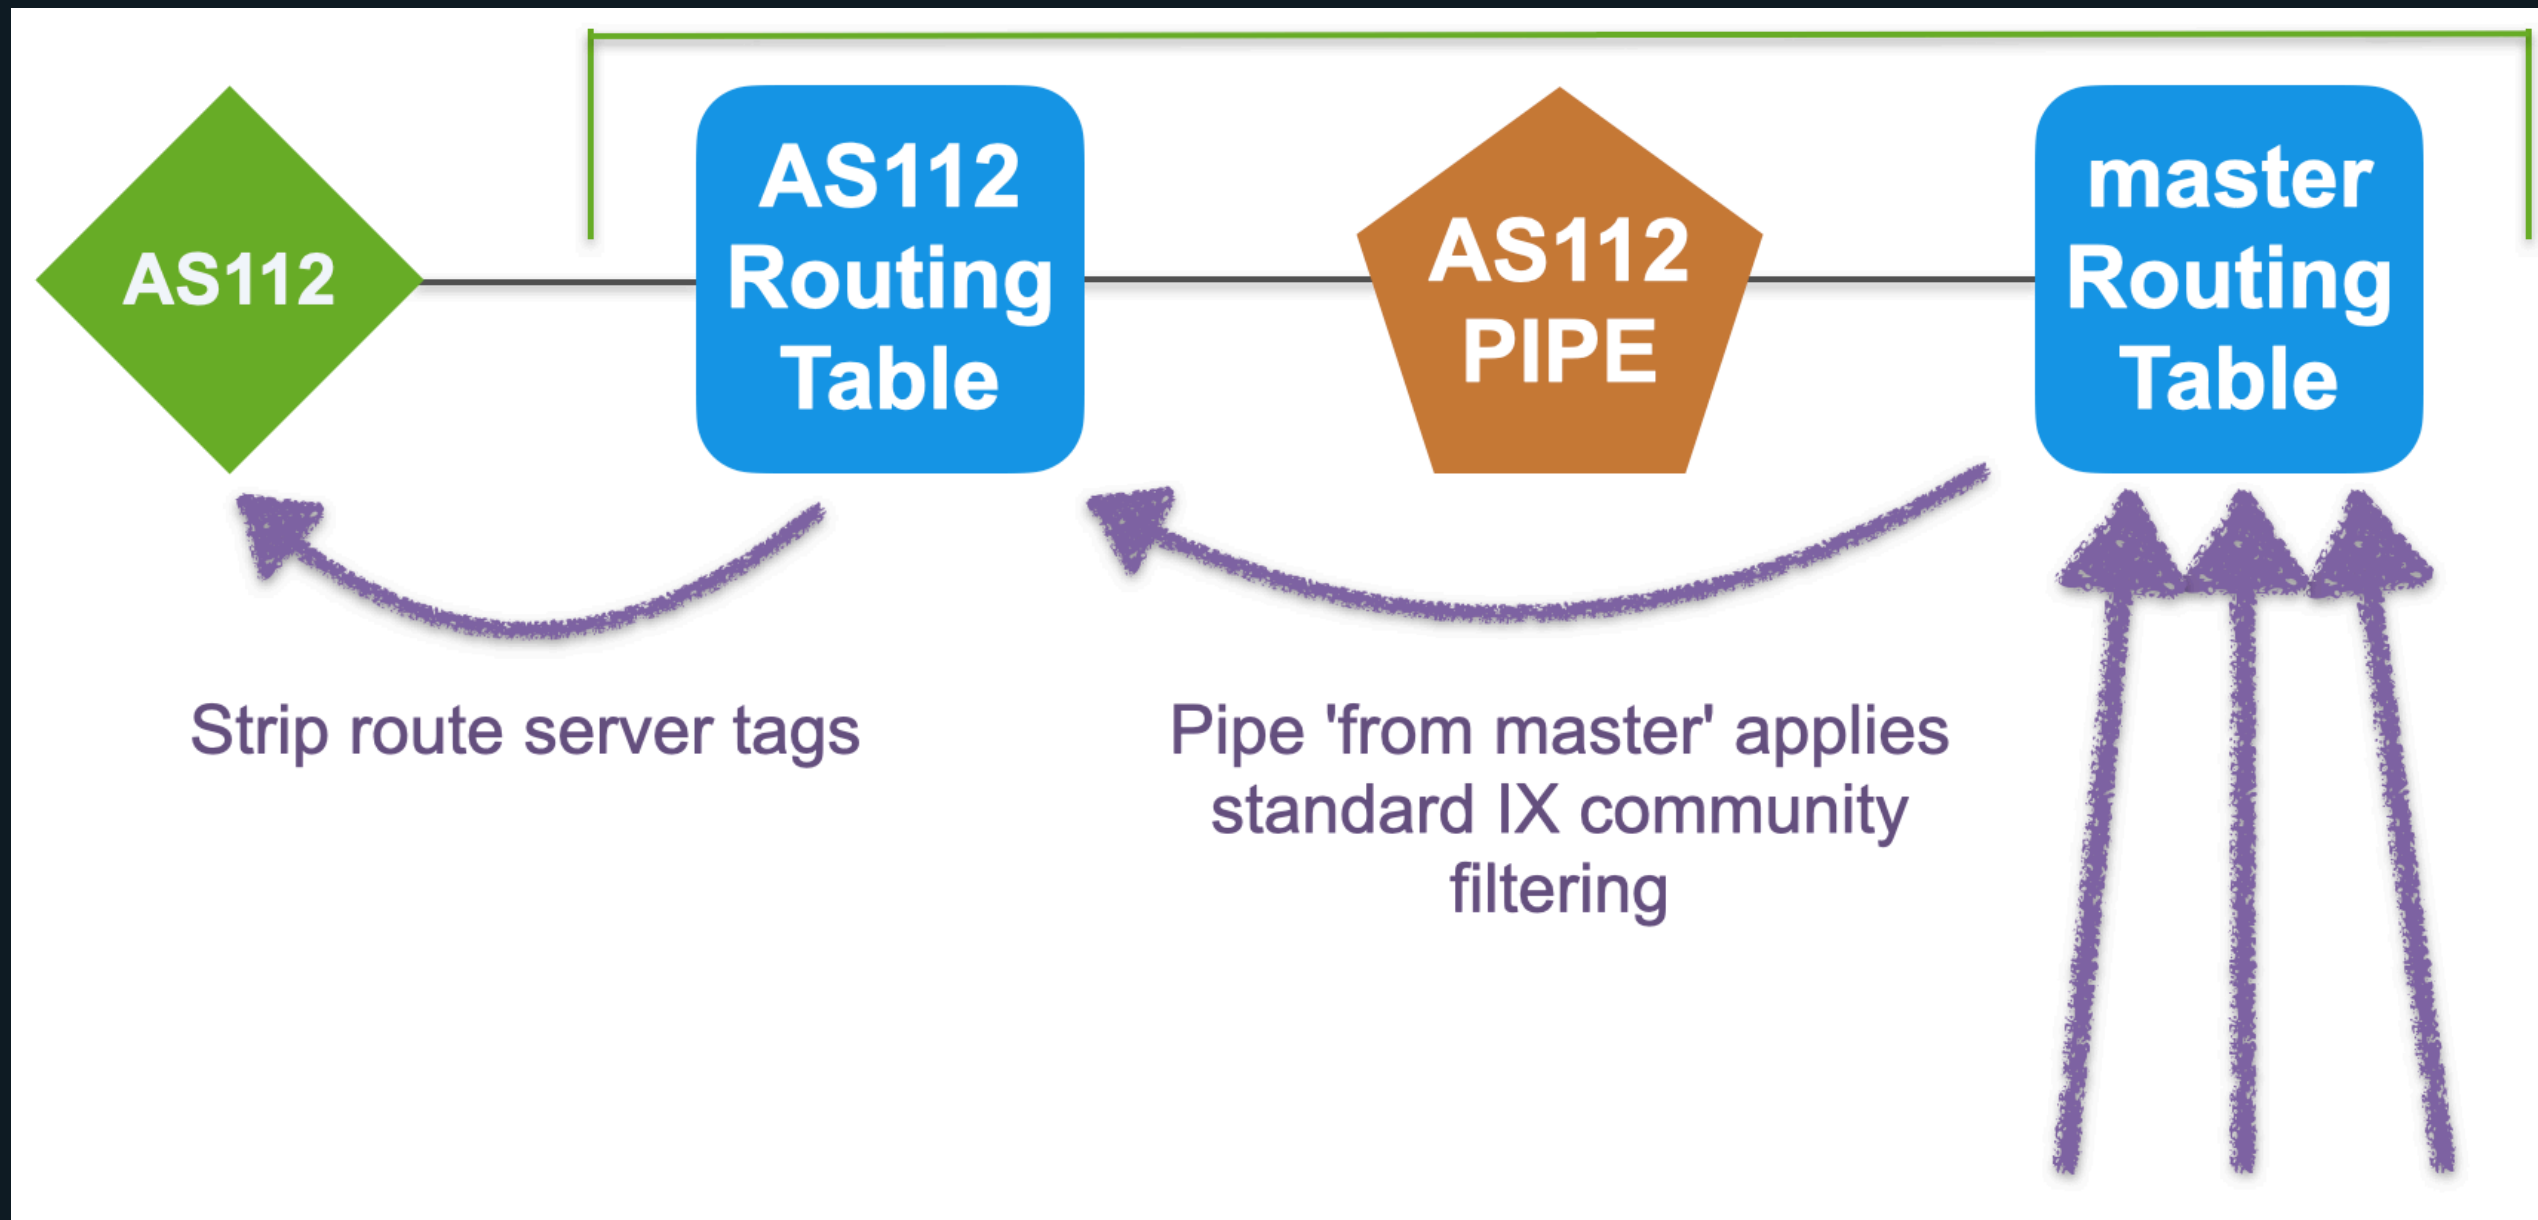
IXP Manager Bird Topology - Export to Client



IXP Manager Bird Topology - Export to Client



IXP Manager Bird Topology - Export to Client



Coming in Part 2:

Configuring Route Servers **with** **IXP Manager**

Thanks for watching!

- <https://www.ixpmanager.org/>
- <https://docs.ixpmanager.org/>
- <https://www.barryodonovan.com/>
- [@barryo79](#) on Twitter
- barry.odonovan@inex.ie