

TeraFlow
SDN
by ETSI

Hackfest #4: Introduction to BGP & BGP-LS Speaker

Pablo Armingol(TID)

Juan Carlos Caja (TID)

What is BGP?

What is BGP?

- BGP is the routing protocol used to advertise routes between institutions and the KAREN network
- More than just a routing protocol, BGP routes contain many additional attributes
- Controlled by flexible “Policy” rules that limit what routes we will learn and what we will advertise
- BGP policy is traditionally used as an interpretation of commercial arrangements between carriers, ISP's etc
- **BGP can be as simple or as complex as you wish**

What RFC's define BGP?

RFC 1771 - A Border Gateway Protocol 4 (BGP-4)

RFC 1772 - Application of the Border Gateway Protocol in the Internet

RFC 1997 - BGP Communities Attribute

RFC 1965 - Autonomous System Confederations for BGP

RFC 1966 - BGP Route Reflection. An alternative to full mesh IBGP

RFC 2270 - Using a Dedicated AS for Sites Homed to a Single Provider

RFC 2283 - Multiprotocol Extensions for BGP-4

RFC 2385 - Protection of BGP Sessions via the TCP MD5 Signature Option

RFC 2439 - BGP Route Flap Damping

RFC 2545 - Multiprotocol Extensions for IPv6 Inter-Domain Routing

Routing Protocols

Routing Protocols


IGP - Interior Gateway protocol

- Figures optimal path from one node to another node in a network
- Examples include RIP, OSPF, ISTS etc...
- Runs under a single technical/administrative control (AS)
- Can support either IPv4 and/or IPv6

EGP - Exterior Gateway protocols

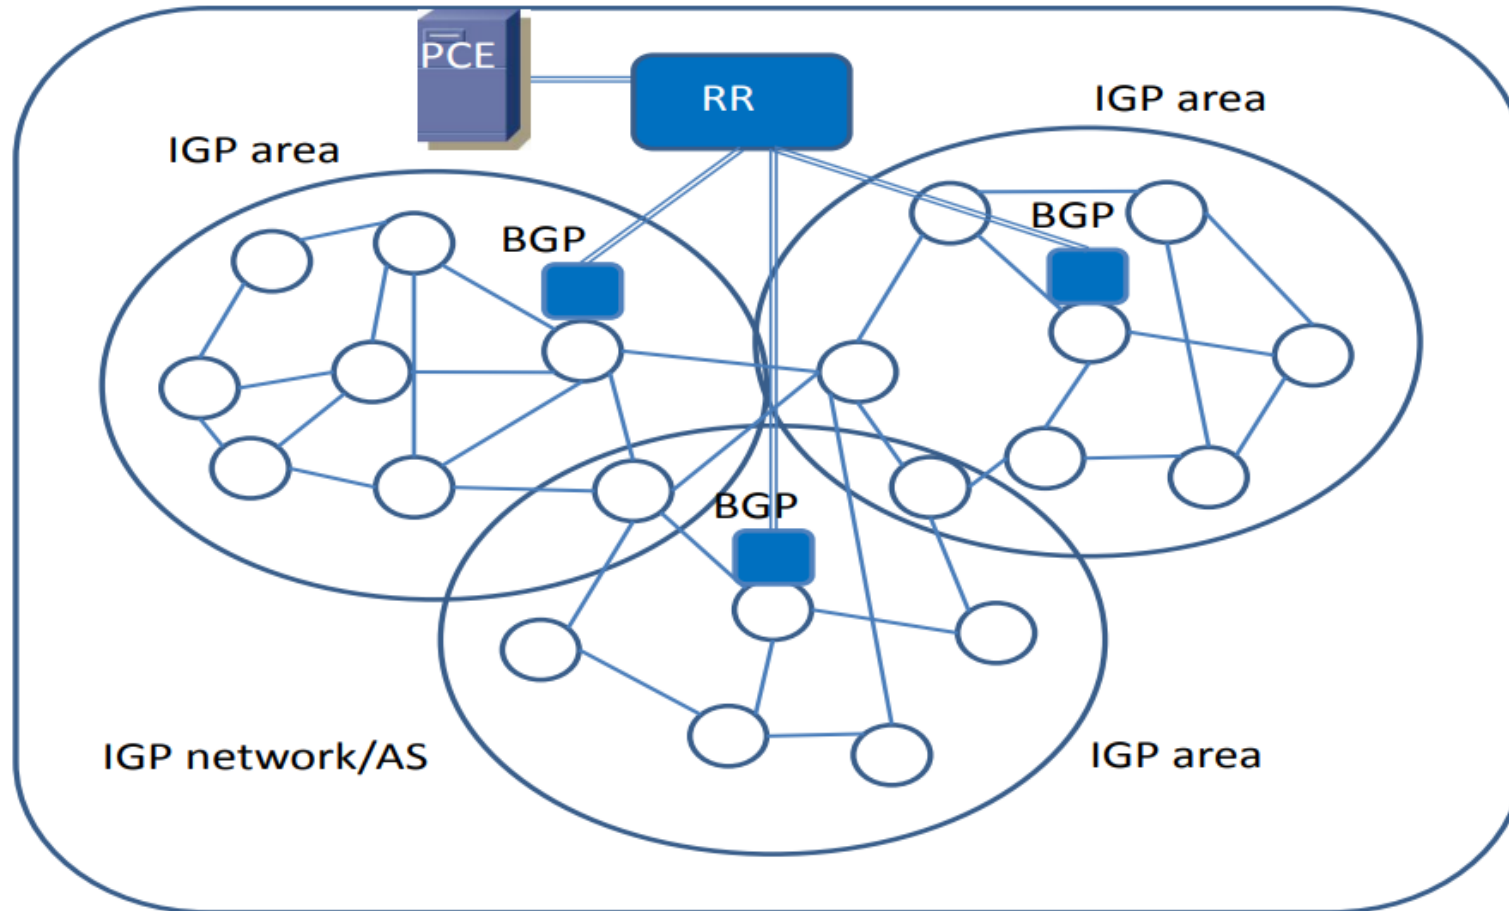
- Allows different AS's to exchange routing information to allow traffic across two different areas of control
- Only one EGP used in the KAREN network -> **BGP** .
- Policy allows Peers to control routes learnt between sites based on established agreements .
- Best practice is to only advertise the routes you wish people to actually reach.

Routing Protocols



BGP-LS

BGPLS



NLRI-Information

- Type 1 – Node NLRI
 - Node Identifier and Node Descriptor fields
 - OSPF Router ID or IS-IS System ID
- Type 2 – Link NLRI
 - Uniquely identify a link in the network
 - Link Descriptor field – TLV format
- Type 3 – IPv4 Topology prefix NLRI
- Type 4 – IPv6 Topology prefix NLRI
 - Uniquely identify a prefix originated by a node
 - Both IPv4 and IPv6 have the same prefix NLRI format

```
*> [V] [L2] [I0x0] [N[c100] [b192.168.2.2] [s0000.0000.0002.00]] /328  
0.0.0.0 0 i
```

```
*>  
[E] [L2] [I0x0] [N[c100] [b192.168.2.2] [s0000.0000.0001.00]] [R[c100] [b192.168.2.  
2] [s0000.0000.0004.01]] [L[i10.1.14.1] [n10.1.14.4]] /704  
0.0.0.0 0 i
```

```
*>  
[T] [L2] [I0x0] [N[c100] [b192.168.2.2] [s0000.0000.0003.00]] [P[p10.1.23.0/24]] /392  
0.0.0.0 0 i
```

Node Attribute

- Multi-Topology Id
- Node Flag Bits
- Opaque Node Attribute
- Node Name
- IS-IS Area Id
- IPv4 Router-Id (local node)
- IPv6 Router-Id (local node)

Link Attributes

- IPv4 / IPv6 Router-Id (local & Remote)
- IGP metric
- Administrative group
- Max. link bandwidth
- Max. reservable bandwidth
- Unreserved bandwidth
- TE default metric
- Link protection type
- MPLS protocol mask
- Shared risk link group
- Opaque link attribute Link name

BGP Protocol Overview

BGP Protocol Overview

- BGP Runs over TCP
- Any two routers that have formed a TCP connection to exchange BGP information are called “Peers” or “Neighbours”
- Once connection is made, Peers exchange their full BGP routing tables.
- Updates are then sent as the table changes or new routes are added to the network.
- BGP peers in the KAREN network should be capable of exchanging both IPv4 and IPv6 routes

BGP routes

BGP routes

- BGP routes contain more than just the advertised prefix
 - Origin
 - ASPath
 - Next Hop
 - Local Preference
 - Multiple Exit Discriminator
 - Community
- BGP Policy looks at the prefix as well as route attributes for decision making
-
- BGP attributes can often be changed to influence downstream policy

Autonomous System

Autonomous System

Autonomous System (AS)

- Group of routers belonging to a single administrative domain
- Viewed externally as a single, coherent interior routing domain
- Each AS runs their own chosen IGP

AS Numbers

- Public and private AS numbers are available for use
- Public numbers assigned locally by APNIC to each institution
- Larger tertiary institutions potentially already have one assigned

You will need a Publicly Assigned AS number to connect to KAREN!

BGP Policy

BGP Policy

BGP Policy controls what BGP routes are installed in the routers routing table and what routes are advertised to your peers

Use BGP policy when

- You don't want to import all learned routes into the routing table
- You don't want to advertise all known routes to neighboring routers
- You want BGP to receive routes from another protocol (Redistribution)
- You want to modify information (BGP Attributes) associated with routes

BGP Policy configuration varies with each router vendor and platform

BGP Policy can be as simple or as complex as you wish

Ensure your routers policy implementation is as flexible as possible

BGP Policy

Default Policy

- BGP when left to its own devices will fall back to default policy for the import and export of routes.
- Import all routes learned from BGP neighbours are installed in the routing table
- Export all routes learned from BGP neighbours to all BGP
- If you aren't running IBGP in your network then you will need policy to redistribute IGP routes

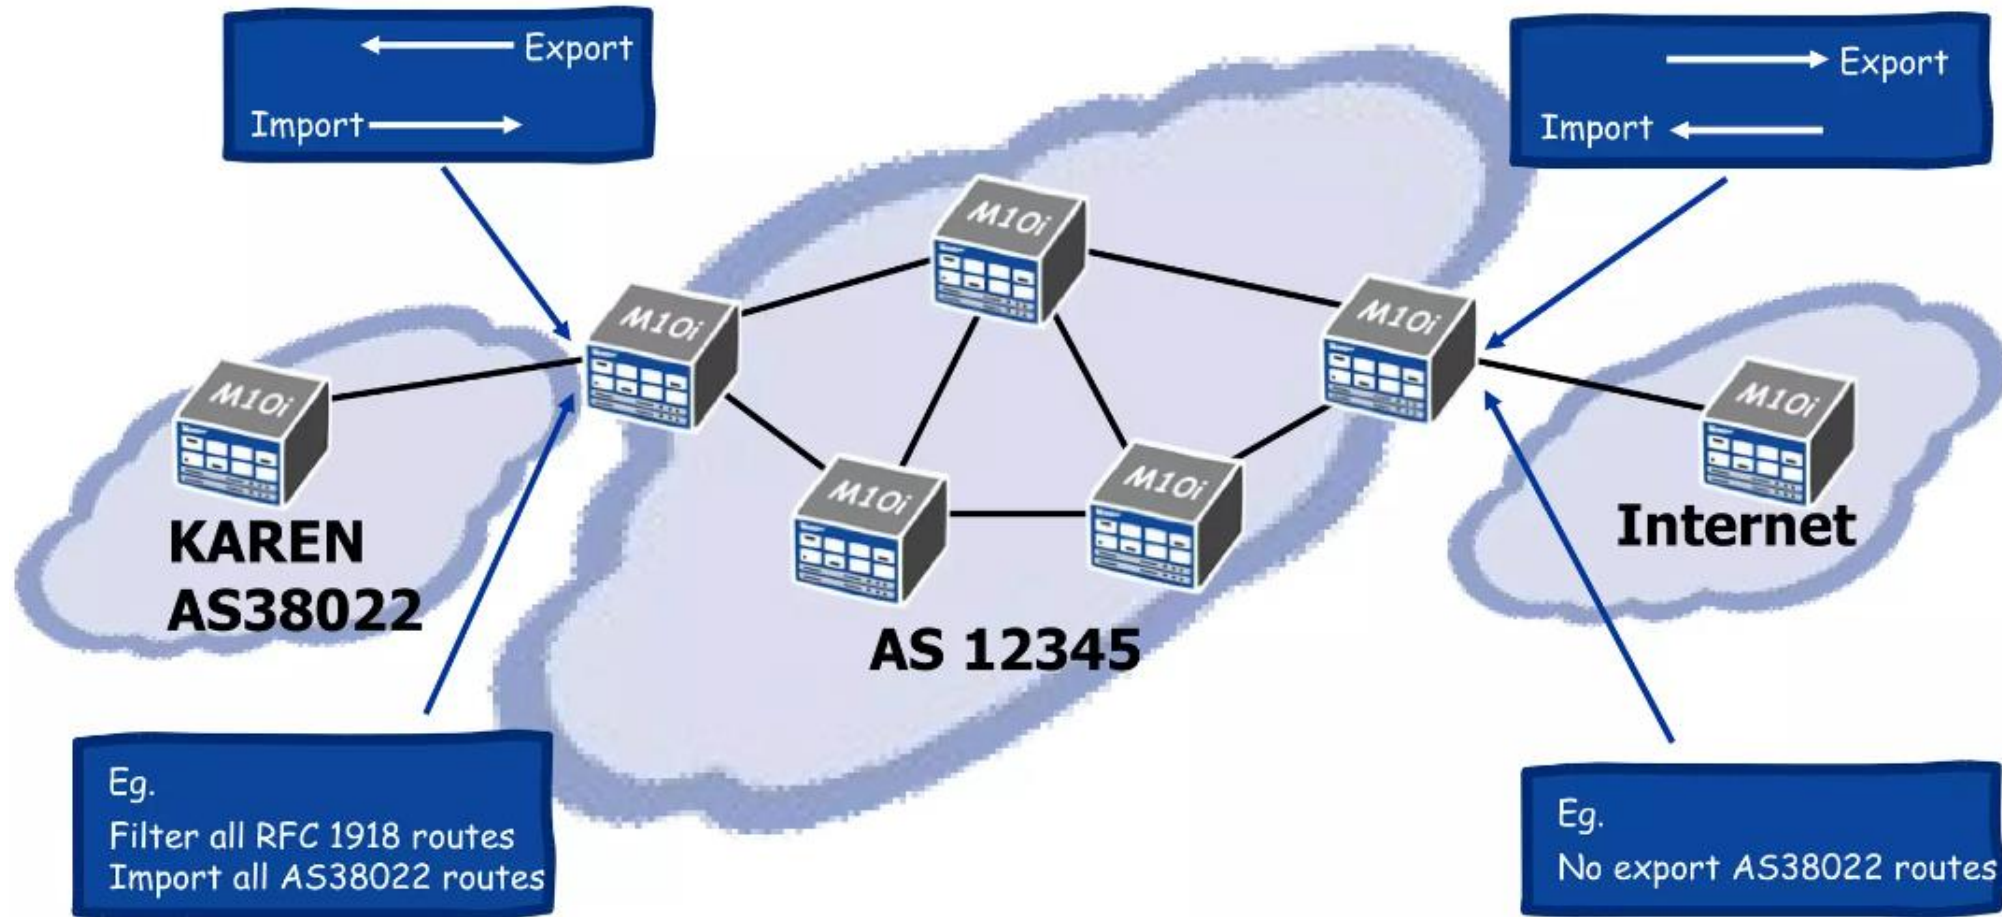
Import Policy

- Apply an import routing policy to control the routes that the routing protocol process uses to determine active routes
- Affects routes that BGP receives from a neighbour
- Modify BGP attributes

Export Policy

- Apply an export routing policy to control the routes that a BGP router advertises to its neighbour
- Modify BGP attributes

BGP Policy



BGP-LS Speaker

BGP-LS Speaker



Open Source for Smart Networks and Services

Add BGPLS Speaker

New bgpls speaker

ip

port

as

[+ Add Speaker](#)

BGP-LS Speaker


Topology

1 speakers found in context *admin*

+ Add BGPLS Speaker

Speaker IP Address	Speaker As Number	
192.168.159.20	100	

BGP-LS Speaker



Open Source for Smart Networks and Services



TeraFlow

[Home](#)

[Device](#)

[Link](#)

[Service](#)

[Slice](#)

[Policy Rules](#)

[Grafana](#)

[Debug](#)

[Load Generator](#)

[BGPLS](#)

[About](#)

Selected Context(**admin**)/Topology(**admin**)


BGPLS Topology

1 speakers found in context *admin*

[+ Add BGPLS Speaker](#)

[✎ Edit BGPLS Speakers](#)

Node name	Local IGP Id	Remote IGP Id	Learnt from
CSR-3	3		192.168.159.20 + Add Device
CSR-4	4		192.168.159.20 + Add Device



TeraFlow
SDN
by ETSI

Thank You all for
your participation!