











# **Border Gateway Protocol**

 A Routing Protocol used to exchange routing information between different networks

Exterior gateway protocol

- Described in RFC4271 RFC4276 gives an implementation report on BGP RFC4277 describes operational experiences using BGP
- IETF Working Groups
   IDR (Internet-Domain Routing: <u>http://datatracker.ietf.org/wg/idr/</u>)
   SIDR (Secure IDR: <u>http://datatracker.ietf.org/wg/sidr/</u>)
- The Autonomous System is the cornerstone of BGP
   It is used to uniquely identify networks with a common routing policy



## Autonomous System Number (ASN)

<ul> <li>Two ranges</li> <li>0-65535</li> <li>65536-4294967295</li> </ul>	(original 16-bit range) (32-bit range - RFC4893)	
<ul> <li>Usage: 0 and 65535</li> </ul>	(reserved)	
1-64495	(public Internet)	
64496-64511	(documentation - RFC5398)	
64512-65534	(private use only)	
23456	(represent 32-bit range in 16-bit world)	
65536-65551	(documentation - RFC5398)	
65552-4294967295	(public Internet)	
<ul> <li>32-bit range represer</li> <li>Defines "asplain" (tra</li> </ul>	ntation specified in RFC5396	
© 2010 Cisco and/or its affiliates. All rights reserved.		Cisco Confidential 9







































### Origin

- · Conveys the origin of the prefix
- Historical attribute
   Used in transition from EGP to BGP
- Transitive and Mandatory Attribute
- Influences best path selection
- Three values: IGP, EGP, incomplete
   IGP generated by BGP network statement
   EGP generated by EGP
   incomplete redistributed from another routing protocol











#### Multi-Exit Discriminator "metric confusion"

- MED is non-transitive and optional attribute Some implementations send learned MEDs to iBGP peers by default, others do not
   Some implementations send MEDs to eBGP peers by default, others do not
- Default metric varies according to vendor implementation Original BGP spec (RFC1771) made no recommendation
  - Some implementations handled absence of metric as meaning a metric of 0
  - Other implementations handled the absence of metric as meaning a metric of 2<sup>32</sup>-1 (highest possible) or 2<sup>32</sup>-2 Potential for "metric confusion"







































Curr	ent capabilities are:		
0	Reserved	[RFC3392]	
1	Multiprotocol Extensions for BGP-4	[RFC4760]	
2	Route Refresh Capability for BGP-4	[RFC2918]	
3	Outbound Route Filtering Capability	[RFC5291]	
4	Multiple routes to a destination capability	[RFC3107]	
5	Extended Next Hop Encoding	[RFC5549]	
64	Graceful Restart Capability	[RFC4724]	
65	Support for 4 octet ASNs	[RFC6793]	
66	Deprecated		
57	Support for Dynamic Capability	[ID]	
68	Multisession BGP	[ID]	
69	Add Path Capability	[ID]	
70	Enhanced Route Refresh Capability	[ID]	
See	www.iana.org/assignments/capability-codes		















































![](_page_39_Figure_1.jpeg)

![](_page_39_Figure_2.jpeg)

![](_page_40_Figure_1.jpeg)

![](_page_40_Figure_2.jpeg)

	Connectivity	Multi-Level Hierarchy	Policy Control	Scalability	Migration Complexity
Confed- erations	Anywhere In the Network	Yes	Yes	Medium	Medium to High
Route Reflectors	Anywhere In the Network	Yes	Yes	Very High	Very Low

![](_page_41_Figure_2.jpeg)

![](_page_42_Figure_1.jpeg)

![](_page_42_Figure_2.jpeg)

![](_page_43_Figure_1.jpeg)

![](_page_43_Figure_2.jpeg)

![](_page_44_Figure_1.jpeg)

![](_page_44_Figure_2.jpeg)

![](_page_45_Figure_1.jpeg)

![](_page_45_Figure_2.jpeg)

![](_page_46_Figure_1.jpeg)

![](_page_46_Picture_2.jpeg)

![](_page_47_Figure_1.jpeg)

![](_page_47_Figure_2.jpeg)

000		IP/MPLS Products from Sprint 말려	
	https 🗎 www.sprint.	t.net/index.php?p=policy_bgp C Reader O	
	WHAT YOU	U CAN CONTROL	
	AS-PATH PREPE	PENDS	
	Sprint allows of prepending with Additionally, S	customers to use AS-path prepending to adjust route preference on the network. Such vill be received and passed on properly without notifying Sprint of your change in announcements. Sprint will prepend AS1239 to eBGP sessions with certain autonomous systems depending on a multip. Currently the following ASe are properted 1468 200 2014 2010 2365 2561	
	4635, 701, 701	infumity, currently, the following Ases are supported: 1000, 209, 2914, 5500, 5550, 5549, 5501, 018, 702 and 8220.	
	String	Resulting AS Path to ASXXX	
	65000:XXX	Do not advertise to ASXXX	
	65001:XXX	1239 (default)	
	65002:XXX		
	65003:XXX	1239 1239 1239 ISP Examples: Spr	Int
	65004:XXX	1239 1239 1239 1239	
	String	Resulting AS Path to ASXXX in Asia	
	65070:XXX	Do not advertise to ASXXX	
	65071:XXX	1239 (default)	
	65072:XXX	1239 1239	
	65073:XXX	1239 1239 1239 More into at https://www.sprint.net/index.	.pnp
	65074:XXX	1239 1239 1239 1239 p=policy_bgp	
	-		
	String	Resulting AS Path to ASXXX in Europe	
	65050:XXX	Do not advertise to ASXXX	
	65051:XXX	1239 (default)	
	65052:XXX	1239 1239	
© 2010 Cisco a	65053:XXX	1239 1239 1239	ential

bar cu	stomer c	ommunities			
Custome	rs wantin	ng to alter local preference on t	heir routes.		
NTT Comm	nunication	s BGP customers may choose to af	fect our local preference on their routes by marking	their routes with	the
following o	ommunitie	98:			
Communi	tvl.ocal-n	ref Description			
(default)	120	customer			
2914:450	96	customer fallback			
2914:460	98	peer backup	Some ISP Ex	ample	s' N⊺
2914:470	100	peer			
2914:480	110	customer backup			
2914:490	120	customer default			
2914:666		blackhole			
NTT Comn communiti	nunication	s BGP customers may choose to pr	epend to all other NTT Communications BGP custo	omers with the follo	owing
Communi	ty	Description			
2914:411	prepend	is o/b to customer 1x	More info at www.us.r	ntt net/abou	t/policy
	prepend	is o/b to customer 2x	routing cfm		
2914:412	prepend	is o/b to customer 3x	routing.cim		
2914:412 2914:413	proporte				
2914:412 2914:413 Custome	rs wantin	ng to alter their route announce	ments to peers.		
2914:412 2914:413 Custome NTT Comm	rs wantin	ng to alter their route announce s BGP customers may choose to pi	ments to peers. epend to all NTT Communications peers with the fo	blowing communi	ties:
2914:412 2914:413 Custome NTT Comn Communi	rs wantin nunication	ng to alter their route announce s BGP customers may choose to pr Description	ments to peers. epend to all NTT Communications peers with the fo	llowing communi	ties:
2914:412 2914:413 Custome NTT Comm Communi 2914:421	rs wantin nunication ty prepend	ng to alter their route announce s BGP customers may choose to pr Description is o/b to peer 1x	ments to peers. epend to all NTT Communications peers with the fo	bllowing communi	ties:
2914:412 2914:413 <b>Custome</b> NTT Comm 2914:421 2914:422	rs wantin nunication ty prepend prepend	ng to alter their route announce s BGP customers may choose to pr Description is o/b to peer 1x is o/b to peer 2x	ments to peers. epend to all NTT Communications peers with the fo	blowing communi	ties:
2914:412 2914:413 <b>Custome</b> NTT Comm 2914:421 2914:422 2914:423	rs wantin nunication ty prepend prepend prepend	te to alter their route announce s BGP customers may choose to pr Description is o/b to peer 1x is o/b to peer 2x is o/b to peer 3x	ments to peers. epend to all NTT Communications peers with the fo	llowing communi	ties:

#### ISP Examples: Verizon Business Europe

aut-num:	AS702	
descr:	Verizon Bu	siness EMEA - Commercial IP service provider in Eur
remarks:	VzBi uses	the following communities with its customers:
	702:80	Set Local Pref 80 within AS702
	702:120	Set Local Pref 120 within AS702
	702:20	Announce only to VzBi AS'es and VzBi customers
	702:30	Keep within Europe, don't announce to other VzBi AS
	702:1	Prepend AS702 once at edges of VzBi to Peers
	702:2	Prepend AS702 twice at edges of VzBi to Peers
	702:3	Prepend AS702 thrice at edges of VzBi to Peers
	Advanced o	communities for customers
	702:7020	Do not announce to AS702 peers with a scope of
		National but advertise to Global Peers, European
		Peers and VzBi customers.
	702:7001	Prepend AS702 once at edges of VzBi to AS702
		peers with a scope of National.
	702:7002	Prepend AS702 twice at edges of VzBi to AS702
		peers with a scope of National.
(more)		
© 2010 Cisco and/or its affilia	tes. All rights reserved.	Cisco Confidential 99

ISP E: Verizo	xamples: n Business Europe
(more)	
	702:7003 Prepend AS702 thrice at edges of VzBi to AS702 peers with a scope of National.
	702:8020 Do not announce to AS702 peers with a scope of European but advertise to Global Peers, National Peers and VZBi customers.
	702:8001 Prepend AS702 once at edges of VzBi to AS702 peers with a scope of European.
	702:8002 Prepend AS702 twice at edges of VzBi to AS702 peers with a scope of European.
	702:8003 Prepend AS702 thrice at edges of VzBi to AS702 peers with a scope of European.
	Additional details of the VzBi communities are located at: http://www.verizonbusiness.com/uk/customer/bgp/
mnt-by:	WCOM-EMEA-RICE-MNT
Source.	
© 2010 Cisco and/or its a	Nales Al right reserved. Close Confidential 10

Some ISP	Examples	
3T Ignite		
aut-num:	AS5400	
descr:	BT Ignite European Backbone	
remarks:		
remarks:	Community to	Community to
remarks:	Not announce To peer:	AS prepend 540
remarks:		
remarks:	5400:1000 All peers & Transits	5400:2000
remarks:		
remarks:	5400:1500 All Transits	5400:2500
remarks:	5400:1501 Sprint Transit (AS1239)	5400:2501
remarks:	5400:1502 SAVVIS Transit (AS3561)	5400:2502
remarks:	5400:1503 Level 3 Transit (AS3356)	5400:2503
remarks:	5400:1504 AT&T Transit (AS7018)	5400:2504
remarks:	5400:1506 GlobalCrossing Trans(AS3549)	5400:2506
remarks:		
remarks:	5400:1001 Nexica (AS24592)	5400:2001
remarks:	5400:1002 Fujitsu (AS3324)	5400:2002
remarks:	5400:1004 C&W EU (1273)	5400:2004
<snip></snip>		
notify:	notify@eu.bt.net And mar	VI
mnt-by:	CIP-MNT	rol
source:	RIPE Many mo	re!

ome ISP E evel 3	Examples
aut-num:	AS3356
descr: <snip></snip>	Level 3 Communications
remarks:	
remarks:	customer traffic engineering communities - Suppressio
remarks:	
remarks:	64960:XXX - announce to AS XXX if 65000:0
remarks:	65000:0 - announce to customers but not to peers
remarks:	65000:XXX - do not announce at peerings to AS XXX
remarks:	
remarks:	customer traffic engineering communities - Prepending
remarks:	
remarks:	65001:0 - prepend once to all peers
remarks:	65001:XXX - prepend once at peerings to AS XXX
<snip></snip>	
remarks:	3356:70 - set local preference to 70
remarks:	3356:80 - set local preference to 80
remarks:	3356:90 - set local preference to 90
remarks:	3356:9999 - blackhole (diseard) traffic
<snip></snip>	And many
mnt-by	LEVEL3-MNT many morel
source:	RIPE Geo Confidential

![](_page_51_Picture_1.jpeg)