

Explaining 4-octet Autonomous System (AS) Numbers for Cisco IOS

- Motivation: Why should you care about 4-octet AS numbers
- Operation: What changes for 4-octet AS support
- Configuration: Configuration changes for 4-octet AS

Last updated: January 2009

What You Will Learn

This paper explains the changes the 4-octet AS has for the BGP routing infrastructure. It will also explain some technical and operational differences between the traditional 2-octet and 4-octet AS BGP technology. This paper will also explain configuration specific aspects in relationship to 4-octet AS systems.

This paper is targeted towards engineers designing, deploying or having to support a 4-octet AS BGP infrastructure.

Why Should You Care About 4-octet AS

At the time of BGP development and standardization, it was assumed that availability of a 16 bit binary number to identify the Autonomous System (AS) within BGP would have been more than sufficient. The 16 bit AS number, also known as the 2-octet AS number, provides a pool of 65536 unique Autonomous System numbers (ASN). From this pool, 1023 numbers have been reserved for local or private usage, and 3 are reserved for special usage. The remaining pool of AS's (64510) are available to support the public Internet infrastructure. The IANA manages the available BGP ASN pool, with the assignments being carried out by the Regional Registries.

The consumption rate of the publicly available ASNs is well known and can be measured by monitoring the remaining unallocated ASNs in the IANA pool. A forecast can be made based upon the current burn-rate of ASNs. These statistics indicate that the entire public 2 byte AS pool will be fully depleted by early to middle 2011 (source: <http://www.potaroo.net/tools/asn16/>). The solution to this depletion is the expansion of the existing 2-octet AS number towards a 4-octet AS number. Work for this extension was started in 2001 and is standardized in IETF RFC 4893 and provides a theoretical 4,294,967,296 unique AS numbers.

Backward Compatibility

An important aspect is the consideration regarding backward compatibility and interoperability between the 'old' 2-octet AS numbers and the 'new' 4-octet AS numbers. The IETF and the standardization of the revised 4-octet AS number is developed in such a way that there is a high degree of backward compatibility between the 4-octet and 2-octet AS number. This backward compatibility is achieved by creating a set of new BGP capabilities, attributes and communities.

Cisco IOS Software releases supporting RFC4893 (BGP Support for Four-octet AS Number Space) introduces the following new BGP aspects:

- A new 4-octet BGP capability advertisement
- New update attributes for 2-octet transit
- New extended communities
- AS_TRANS—A well known transition BGP AS Number '23456'

These new attributes, extended BGP communities and BGP capability advertisement needs understanding when deploying routers and devices which are 4-octet AS capable. The next section of this paper describes the operational behavior when working with BGP 4-octet AS numbers.

Operation

This paper will assume that an 'OLD' BGP speaker is supporting only a 2-octet AS number and that a BGP speaker that supports the new 4-octet AS number is as a 'NEW' BGP speaker. The BGP AS number can be seen within a variety of places within the BGP protocol:

In the 'OPEN' message

In the AS-PATH attribute of the 'UPDATE' message

In the AGGREGATOR attribute of the 'UPDATE' message

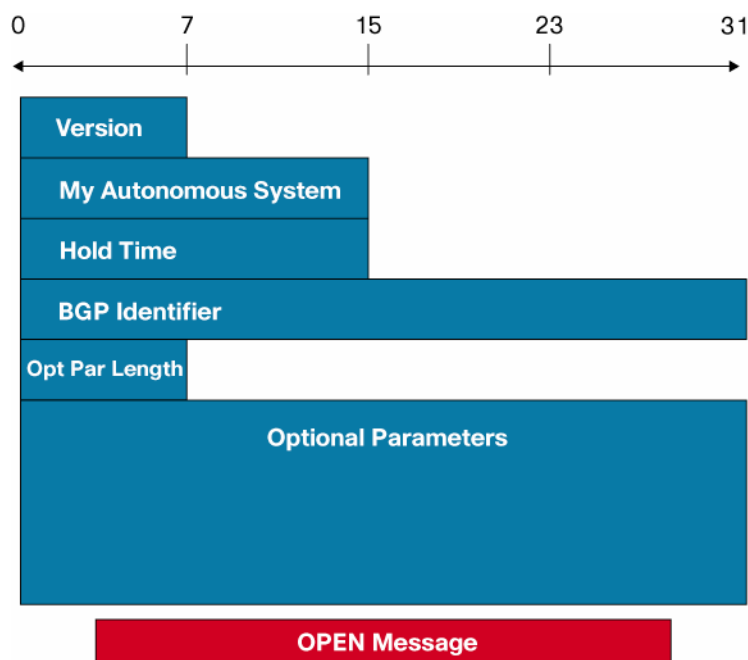
In the EXTENDED COMMUNITY for Route-Target (RT) and SITE-OF-ORIGIN (SOO)
which are used for VPN's

A NEW Speaker will have to support 4-octet AS values for each of the above identified messages. For a NEW BGP speaker, the values regarding the BGP Autonomous System with the OPEN message, the AS-PATH and AGGREGATOR attribute and the extended communities are based upon 4-octet values, while for an OLD BGP speaker, this is based upon 2-octet values.

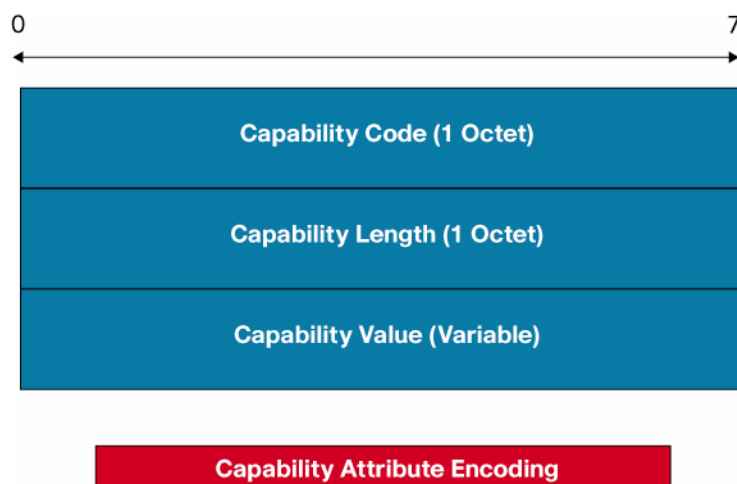
To support a high degree of backward compatibility between an OLD BGP speaker and a NEW BGP speaker two additional optional transitive BGP attributes have been defined within RFC 4893: AS4_PATH and AS4_AGGREGATOR attributes. These attributes are created and attached to each BGP route by a NEW BGP speaker that must communicate with an OLD BGP speaker. These are 'optional' attributes because these attributes may not be understood by an OLD BGP 2-octet AS BGP implementation. The 'transitive' character of these optional attributes means that the attributes are forwarded by an OLD BGP speaker to all its BGP neighbors, even if the actual 2-octet AS BGP router does not understand the meaning of the attributes. More details on the why and how this fully works will be discussed later in this paper.

One of the basic steps of the BGP protocol is to establish adjacencies between BGP peers. Without successful completion of this step, no exchange of updates between the peers will be possible. Successful neighbor negotiation is based upon successful completion of a TCP transport connection, the successful processing of the BGP OPEN message and the periodic detection of BGP keepalive messages. An important task during the neighbor establishment is the negotiation of the BGP capabilities by each peer during the "OPEN" message exchanges.

The "OPEN" BGP message includes a set of parameters and has to be agreed upon before a full BGP adjacency can be established. The "OPEN" BGP message is encoded as follows:



Among other capabilities, a 'NEW' BGP speaker uses BGP Capability Advertisements [RFC2842] to advertise to its neighbors (either internal or external) about the support of 4-octet AS extensions. These BGP Capability Advertisements are encoded within the 'Optional Parameters' of the OPEN message. The Capability parameter contains one or more triples <Capability Code, Capability Length, Capability Value>, where each triple is encoded as shown below:



The use and meaning of these fields are as follows:

Capability Code:

Capability Code is a one octet field that unambiguously identifies individual capabilities. IANA has assigned code '65' to identify the 4-octet AS capability.

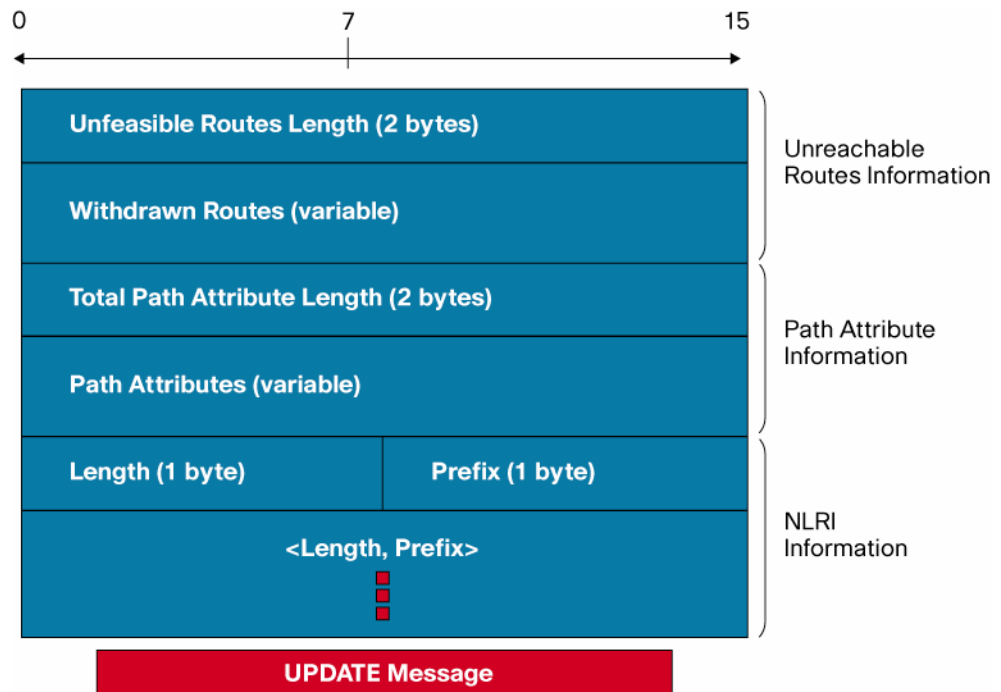
Capability Length:

Capability Length is a one octet field that contains the length of the Capability Value field in octets. It will always be '0' for the 4-octet AS capability.

Capability Value:

Capability Value is a variable length field that is interpreted according to the value of the Capability Code field, and because there is a Capability length of '0' specified, there is no Capability Value specified or required.

Once both BGP peers have agreed upon mutual capabilities, they can start exchanging routing information by means of BGP UPDATE messages. If we assume that both peers are 'NEW' BGP speakers, then they will carry AS path information expressed in terms of 4-octet Autonomous Systems numbers. In addition, NEW BGP speakers use the same 4-octet structure for the AGGREGATOR attribute.



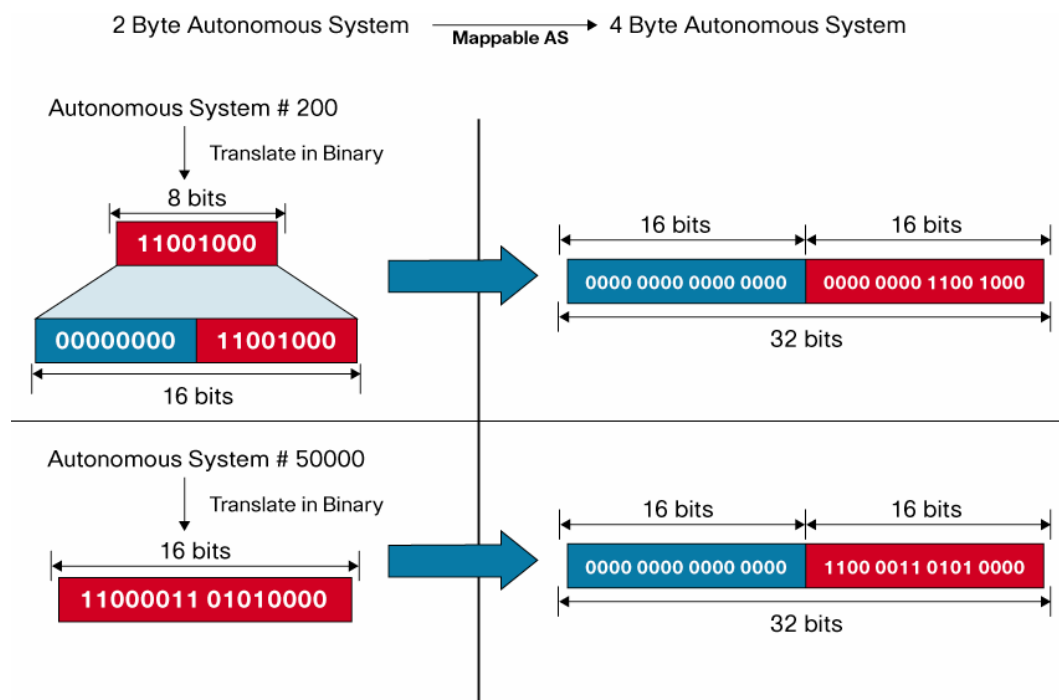
The difference between a full 4-octet AS system and a 2-octet AS system is negligible because the BGP operation does not change at all. The only difference is that in one case the AS is 2 Bytes, it will be in the other case, 4 Bytes. The major difference is driven by the operation principle during transition and compatibility between 2 and 4 Byte systems. An assumed 2-to-4 Byte AS transition constant has been that the existing 2-octet AS infrastructure must be unchanged due to potential historical BGP implementations, and that the 4-octet BGP technology must accommodate for that assumption. RFC4893 has succeeded for most of the deployments.

Transition

When an OLD BGP speaker is peering with a NEW BGP speaker, special care needs to be taken regarding both the BGP adjacency establishment and BGP routing loop detection. An OLD BGP speaker has no way of understanding a 4-Octet Autonomous System number. A NEW BGP speaker has for that purpose, four tricks available that may be used simultaneously for the sending or receiving of BGP updates between an NEW and OLD BGP speaker. These four tricks explain the BGP behavior during a transition scenario.

Trick #1

If the NEW BGP speaker is part of an Autonomous System which is 2-octet mappable, then it will just send the low order 16 bits of the Autonomous System towards the peering OLD BGP speaker within the OPEN message. A 2-octet mappable Autonomous System number is an AS number for which the two higher order octets of the 4-octet AS number is set to zero. The 2-octet mappable addresses are all found with the Autonomous System numbers “0 to 65536 ($=2^{16}$)”. If this Autonomous System would be translated towards a 4-Octet binary value, then the higher order 2-octet will be zero, while the lower two octets carry the full information on the Autonomous System number, hence no information is lost.



However, if the NEW BGP speaker is part of an Autonomous System which is not 2-octet mappable (ie: 100.200), then the NEW BGP speaker will swap the actual 4-octet Autonomous System number with the well know Autonomous System '23456', also known as AS_TRANS within the BGP OPEN message. The impact this will have on the OLD BGP speaker is that the OLD BGP speaker believes it is peering with Autonomous System '23456' (a two byte AS) instead of the real 4-octet AS (ie: Autonomous System 100.200)

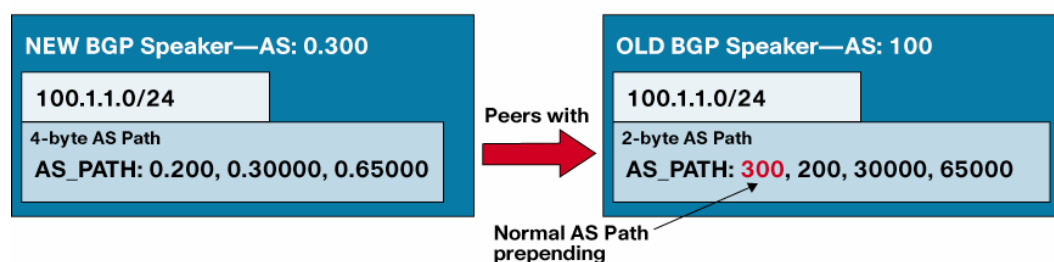
Trick #2

If a NEW BGP speaker needs to send a BGP route update to an OLD BGP speaker, then the NEW BGP speaker must make sure that his old peer can understand the information he is sending. A crucial piece of information is the AS_PATH information attached to a BGP route. This information is very important for AS_PATH based loop detection within the OLD BGP speaker world. What the NEW BGP speaker will have to achieve is to translate his 4-octet AS_PATH towards a 2-octet AS_PATH.

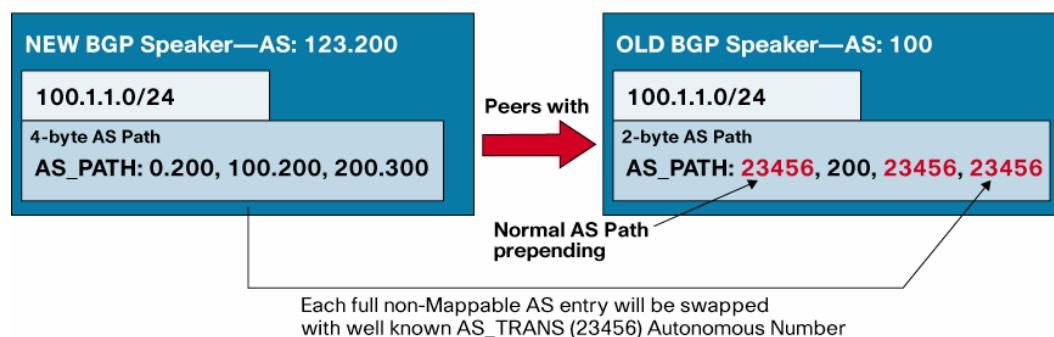
If a NEW BGP speaker is peering with an OLD BGP speaker, he will manipulate the 4-octet AS_PATH to make it 2-octet compatible, which the OLD BGP speaker will understand. For each BGP route, the new BGP speaker will either:

1. Remove the first two octet's of the Autonomous System number, if it was a 4-octet mappable Autonomous System number, or
2. The NEW BGP speaker will swap each non-mappable Autonomous System number with the well know Autonomous System '23456', also known as AS_TRANS.

Example 1: Only Mappable Autonomous System Numbers

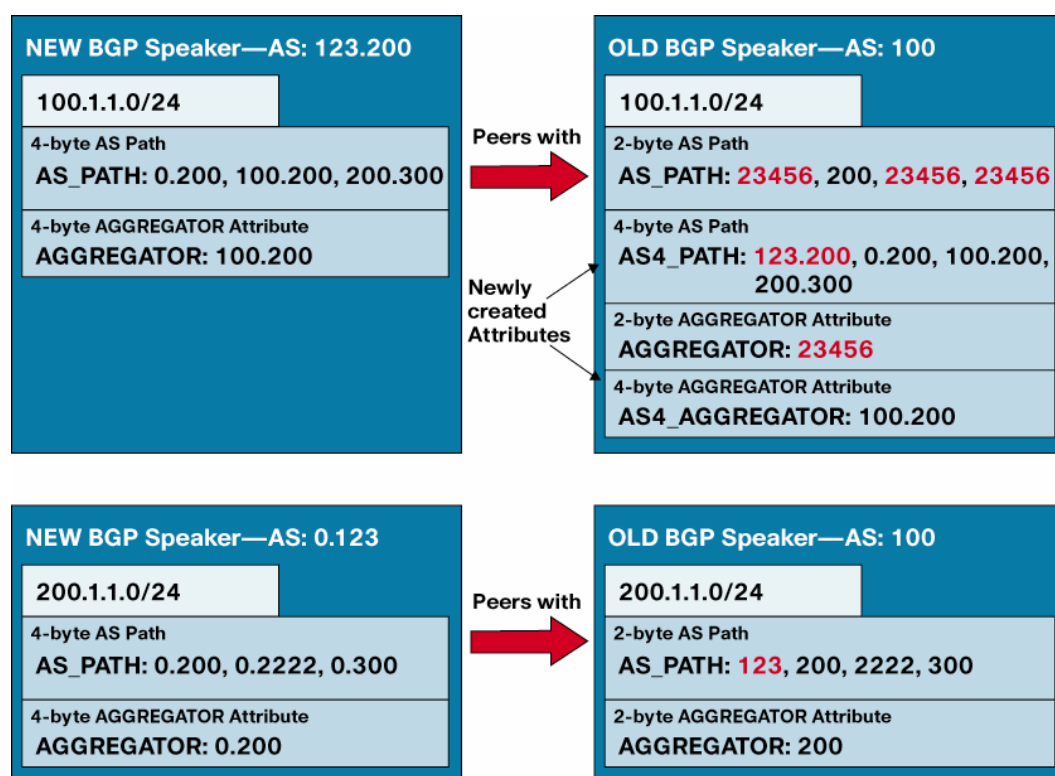


Example 2: Mix of Mappable & non-Mappable Autonomous System Numbers



Trick #3

If the NEW BGP speaker has to deal with a non-mappable Autonomous System number, then in addition to swapping this non-mappable Autonomous System number with '23456', it also adds two 'optional transitive' attributes AS4_PATH (Attribute number 17) and AS4_AGGREGATOR (Attribute number 18) to the BGP route which is sent to the OLD BGP speaker. The AS4_PATH contains the full content of the original 4-octet AS_PATH, whereas the newly created AS_PATH has lost some information due to the conversion from 4-octet AS's into 2-octet AS's. The AS4_AGGREGATOR has same context as the traditional AGGREGATOR attribute. The AGGREGATOR attribute provides information about the last AS number that formed the aggregated BGP route and is followed by the IP address of the BGP speaker that formed the aggregated BGP route.

**Trick #4**

When a NEW BGP speaker receives an update from an OLD one, it should be prepared to receive the AS4_PATH attribute along with the existing 2-octet AS_PATH attribute. If the AS4_PATH attribute is also received, both the attributes will be used by the NEW BGP speaker to construct the exact AS path information, and therefore the information carried by both the attributes will be considered for BGP AS path loop detection within the 4-octet AS world.

It may be that a route has traversed a series of autonomous systems with 2-octet AS numbers and OLD BGP speakers only. In that case, if the route carries the AS4_PATH attribute, this attribute must have remained unmodified since the route left the last NEW BGP speaker. The trailing AS path information (representing autonomous systems with 2-octet AS numbers and OLD BGP


speakers only) is contained only in the current AS_PATH attribute (encoded in the leading part of the AS_PATH attribute).

Under certain conditions, it may not be possible to reconstruct the entire AS path information from the AS_PATH and the AS4_PATH attributes of a route. This occurs when two or more routes that carry the AS4_PATH attribute are aggregated by an OLD BGP speaker, and the AS4_PATH attribute of at least one of these routes carries at least one 4-octet AS number (as opposed to a 2-octet AS number that is encoded in 4 bytes). Depending on the implementation, either the AS4_PATH attribute would be lost during route aggregation, or both the AS_PATH attribute and the AS4_PATH attribute would contain valid, yet partial information that cannot be combined seamlessly, resulting in incomplete AS path information in these cases.

Note that the NEW BGP speaker must also expect to receive the AS4_AGGREGATOR attribute to make sure that the full set of information attached to a BGP route is completely valid and contributes to the global BGP mechanism to aid BGP loop detection.




BGP Loop detection during a hybrid (OLD and NEW) BGP environment:

During a transition scenario, there will be a hybrid BGP environment due to the existence of both NEW and OLD BGP speakers within a single network. It is crucial that BGP can perform loop detection correctly in all situations, including a hybrid BGP environment that exists out of NEW and OLD BGP speakers and is demonstrated by the picture below. At any time and at any location a BGP router has sufficient information available for doing a successful AS-PATH loop detection procedure based on the information available within the AS4-PATH and AS-PATH attributes.





Extended Community Support for 4-octet Autonomous System Numbers

When a router is operating in a VPN modus, this router will be using Extended BGP Communities to identify the VPN and to aid with VPN based routing loop detection. The two communities used in this case are the Route-Target (RT) and Site-of-Origin (SOO) community. Originally they are constructed as follows:



Four-octet AS Specific BCP Extended Community
(ietf-13vpn-as4octet-ext-community-02.txt)



As can be seen, the available data-field for the Autonomous System number is a 16 bit value. When deploying a system supporting 4-octet Autonomous Systems numbers, this is extended towards a 32 bit data-field. The BGP Extended Community is identified for either 2-octet or a 4-octet Autonomous System by the setting of the Extended Community Type fields (Type high and Type Low). The following values for the type fields have been standardized by IETF for Route-Target (RT) and Site-of-Origin (SoO):

Type High	Type Low	Type of Extended Community
00	02	2-byte AS RT
00	03	2-byte AS SoO
02	02	4-byte AS RT
02	03	4-byte AS SoO

IOS Configuration

Once a Cisco router is running an IOS that has 4-octet AS support, then there is no additional configuration required to enable 4-octet AS support for the Cisco IOS router. A few minimal changes have been implemented for regular expressions and the representation of BGP Autonomous System numbers. The Cisco default configuration of the 4-octet AS number is ASPLAIN as specified in RFC5396—"Textual Representation of Autonomous System (AS) Numbers". Changing notation between ASPLAIN and ASDOT can be made according with the following configuration command:

```
router bgp 1.1
  bgp asnotation dot
```

The following can be used to configure a 4-octet AS number on a Cisco router:

ASDOT notation:

```
router bgp 4.4
  bgp log-neighbor-changes
  neighbor 134.0.0.3 remote-as 3.3
```

ASPLAIN notation:

```
router bgp 262148
  bgp log-neighbor-changes
  neighbor 134.0.0.3 remote-as 196611
```

Also the Cisco Command Line Interface (CLI) has been enhanced for BGP CLIs to display both 2-octet and 4-octet AS path value. In case of 4-octet AS value, the value is displayed in XX.YY format.

ASDOT notation:

```
R4#sh ip bgp 1.1.1.0
BGP routing table entry for 1.1.1.0/24, version 2
```

Paths: (1 available, best #1, table default)

Flag: 0x820

Not advertised to any peer

3.3 2 1.1

134.0.0.3 from 134.0.0.3 (134.0.0.3)

Origin IGP, localpref 100, valid, external, best

R4#**sh ip bgp sum**

BGP router identifier 134.0.0.4, local AS number **4.4**

BGP table version is 2, main routing table version 2

1 network entries using 124 bytes of memory

1 path entries using 52 bytes of memory

2/1 BGP path/bestpath attribute entries using 184 bytes of
memory

1 BGP AS-PATH entries using 40 bytes of memory

0 BGP route-map cache entries using 0 bytes of memory

0 BGP filter-list cache entries using 0 bytes of memory

BGP using 400 total bytes of memory

BGP activity 1/0 prefixes, 1/0 paths, scan interval 60 secs

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
134.0.0.3	4	3.3	28	27	2	0	0	00:25:33	1

ASPLAIN notation:

R4#**sh ip bgp 1.1.1.0**

BGP routing table entry for 1.1.1.0/24, version 2

Paths: (1 available, best #1, table default)

Flag: 0x820

Not advertised to any peer

196611 2 65537

134.0.0.3 from 134.0.0.3 (134.0.0.3)

Origin IGP, localpref 100, valid, external, best

R4#**sh ip bgp sum**

BGP router identifier 134.0.0.4, local AS number **262148**

BGP table version is 2, main routing table version 2

1 network entries using 124 bytes of memory

1 path entries using 52 bytes of memory

2/1 BGP path/bestpath attribute entries using 184 bytes of
memory

```

1 BGP AS-PATH entries using 40 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 400 total bytes of memory
BGP activity 1/0 prefixes, 1/0 paths, scan interval 60 secs

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
134.0.0.3	4	196611	28	27	2	0	0	00:25:33	1

Redistribution between routing protocols is done for a 4-octet AS system identically the same as for the existing 2-octet AS:

ASDOT notation:

```

R1(config)#router ospf 1
R1(config-router)#redistribute bgp 3.3

```

```

R3#sh ip rout | include B

```

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP

```

B          2.2.2.0 [20/0] via 123.0.0.2, 00:11:01

```

```

B      192.0.0.0/24 [20/0] via 123.0.0.2, 00:11:01

```

```

R3#sh ip route 192.0.0.0

```

Routing entry for 192.0.0.0/24

Known via "bgp 3.3", distance 20, metric 0

Tag 2, type external

Redistributing via ospf 1

Advertised by ospf 1

Last update from 123.0.0.2 00:12:14 ago

Routing Descriptor Blocks:

* 123.0.0.2, from 123.0.0.2, 00:11:09 ago

Route metric is 0, traffic share count is 1

AS Hops 1

Route tag 2

ASPLAIN notation:

```

R1(config)#router ospf 1
R1(config-router)#redistribute bgp 196611

```

```

R3#sh ip rout | include B

```

```

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B -
BGP

B          2.2.2.0 [20/0] via 123.0.0.2, 00:11:01
B        192.0.0.0/24 [20/0] via 123.0.0.2, 00:11:01

R3#sh ip route 192.0.0.0
Routing entry for 192.0.0.0/24
    Known via "bgp 196611", distance 20, metric 0
    Tag 2, type external
    Redistributing via ospf 1
    Advertised by ospf 1
    Last update from 123.0.0.2 00:12:14 ago
    Routing Descriptor Blocks:
    * 123.0.0.2, from 123.0.0.2, 00:11:09 ago
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 2

```

It is possible to configure BGP policies regarding AS_PATH filtering. A first example demonstrates a filter for a router if you would like to receive only the routes originated from AS 1.4 (and no Internet routes):

```

ip as-path access-list 1 permit ^1\.4$

router bgp 1
  neighbor 4.4.4.4 remote-as 1.4
  neighbor 4.4.4.4 route-map foo in

route-map foo permit 10
  match as-path 1

```

This ensures only networks originated from AS 1.4 are allowed into the configured router. Important to note is the dot notation "X.Y" contains a "." which is a special character in Cisco regular expression. To circumvent the special meaning of this, the "." must be taken away when using X.Y notation patterns. This is done by preceding a backslash (\) to remove the special meaning of the "." character.

Conclusion

This paper provides a technological overview to explain why the Internet will migrate to using the 4-octet AS system; the paper also explains the working principles and provides insight on configuration aspects of a 2/4-octet system. The Cisco IOS 4-octet Autonomous System support is

fully backward compatible with the traditional 2-octet Autonomous System and provides a reliable and proven path for your next generation Internet infrastructure.

For More Information

References used to compose this paper are:

- Cisco IOS BGP Configuration Guide
 - http://www.cisco.com/en/US/tech/tk365/tk80/tsd_technology_support_sub-protocol_home.html
- RFC4893—"BGP Support for Four-octet AS Number Space"
- RFC5396—"Textual Representation of Autonomous System (AS) Numbers "
- RFC2842—"Capabilities Advertisement with BGP-4 "
- 16-bit AS Number Report
 - <http://www.potaroo.net/tools/asn16/>
- ARIN, AS Number Change on 1 January 2009
 - <http://www.arin.net/announcements/07242008.html>
- RIPE NCC, AS Number change could affect Internet routing from 1 January 2009
 - <http://www.ripe.net/news/asn-32-pr2008.html>
- APNIC, AS number change could affect Internet routing from 1 January 2009
 - <http://www.apnic.net/news/2008/0725.html>



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, COBNT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumina, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks, and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco ICS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FrameShare, GigaDrive, HomeLink, Internet Gateway, IOS, iPhones, iQuik Study, iSeePart, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanel, ProConnect, SmartShare, SenderBase, SMI, SmartNet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (081215)