

Network Considerations for a Network Which Does Not Support 4-octet AS

Last updated: January 2009

Regional Registries are driving the adoption and usage of 4-octet Autonomous System (AS) numbers, however not all networks will be able to support 4-octet Autonomous System Numbers on the full network infrastructure from the first day. This paper will explain some considerations related to the full or partial lack of 4-octet AS support.

What You Will Learn

This whitepaper will document a set of considerations for a network operator, when the network does not have full 4-octet AS support.

Network Considerations for Networks Not Supporting 4-octet AS Numbers

This paper will assume that an 'OLD' BGP speaker is supporting only a 2-octet AS number, and that that a BGP speaker that supports the new 4-octet AS number is a 'NEW' BGP speaker. RFC4893 (BGP Support for Four-octet AS Number Space) describes the compatibility between OLD and NEW BGP speakers and allows a gradual transition from a 2-octet supporting router, towards 4-octet supporting router. Even though there is a high degree of compatibility, it is desirable to understand some specific considerations when operating OLD BGP speakers in a world where there may be either directly connected NEW BGP speakers or BGP prefixes containing 4-octet Autonomous System numbers.

Usage of 4-octet AS Information

It is strongly recommended that an Autonomous System will only start using the 4-octet AS information when all BGP speakers, within that Autonomous System, have been upgraded to support 4-octet AS numbers. Making sure that all BGP speakers support 4-octet AS numbers guarantees that each BGP speaker has a valid 4-octet AS number view of the BGP table. If within the Autonomous System, some nodes are OLD speakers, and other nodes are NEW BGP speakers, then the BGP table view can be dramatically different between an OLD and NEW BGP speaker. Not following this guideline may result in BGP routing loops, due to incomplete AS-Path information on the OLD BGP speakers.

Usage of a Well Known Autonomous System Number '23456' (=AS_TRANS)

To facilitate compatibility and gradual upgrade, RFC4893 defines a well known 4-octet Autonomous System number for backward compatibility. This well known AS number is '23456', also known as AS_TRANS.

This well known AS number is used by a NEW BGP speaker peering with an OLD BGP speaker in two situations:

1. If a NEW BGP speaker is using a true 4-octet AS number. A true 4-octet AS number is an AS number which can not fully be represented by a 2-octet AS number. (for example AS# "100.100" can not be represented by a 2-octet AS number, where AS# "0.100" could be represented like AS# "100" towards an OLD BGP speaker, by dropping the leading 16 zero's of the full 4-octet AS number). The OLD BGP speaker expects to peer with a BGP speaker that belongs to a 2-octet Autonomous System, and hence the NEW BGP speaker will either remove the 16 leading zero's from its AS#, or it will swap the non-mappable 4-octet AS# with well known AS# "23456"
2. In similar manner, the 4-octet AS-Path is manipulated by a NEW BGP speaker peering with an OLD BGP speaker. An OLD BGP speaker can not understand any 4-octet AS numbers, and hence the NEW BGP speaker will create from its AS-Path, consisting of all 4-octet AS hops, an AS-Path which only contains 2-octet AS hops. For each entry in the AS-Path, the NEW BGP speaker will decide if it is a mappable AS# or not. If it is a mappable AS#, then the leading 16 "zeros" are removed. If the AS# is a non-mappable AS#, then the NEW BGP speaker will swap the AS hop entry with well known AS# "23456"

A side effect from the compatibility between OLD and NEW BGP speakers is that the usage of "23456" as the BGP AS# is not possible. If an OLD BGP speaker would be configured to belong to AS "23456", then this OLD BGP speaker will result as a AS-Path loop detection drop by any BGP routes that have AS# "23456" within the AS-Path list, resulting in an incomplete view from the OLD BGP speaker perspective.

An additional side effect could be if the NEW BGP speaker belongs, for example, to AS# 100.100 and the OLD BGP speaker to AS# 200. Then, there should be an eBGP session with both BGP speakers. If the OLD BGP Speaker would be wrongly configured as belonging to AS# "23456", then instead of an eBGP peering, an iBGP peering would be established between the OLD and NEW BGP speaker, which is highly undesirable, resulting in the wrong BGP behavior.

BGP Table Memory Utilization

Through the existence of 4-octet AS numbers on the Internet, the size of the BGP table may grow gradually, and hence the memory allocation involved to store the BGP table will also grow. The growth can be accounted for by two drivers:

1. The AS-Path of each BGP route or prefix will now have a 4-octet AS-Path attached to it, instead of the earlier 2-octet AS-Path. This will be the first reason why the BGP Table Memory Utilization will grow.
2. When a NEW BGP speaker peers with an OLD BGP speaker, then the NEW BGP speaker may, for compatibility reasons, swap non 2-octet mappable AS numbers, within the AS-Path list with the well known AS_TRANS "23456" for each BGP route. In addition to this AS# swapping, the NEW BGP speaker will also send the OLD BGP speaker a new optional transitive attribute: AS4-PATH and in some cases, also an AS4-AGGREGATOR. These two attributes, created by the NEW BGP, will allow a remote NEW BGP speaker to re-create the full 4-octet AS-Path further along the BGP route propagation. While initially there will only be a few true non-mappable 4-octet AS numbers seen on the Internet, there will be an ongoing increase of these 4-octet AS#, once Regional Registries start handing out 4-octet AS numbers more aggressively.

These two aspects contribute to increased memory utilization for both OLD and NEW BGP speakers during the transitioning period.

Aggregation on OLD BGP Speakers

This consideration is driven due to compatibility aspects between NEW and OLD BGP speakers. Under certain conditions, it may not be possible for a NEW BGP speaker 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 oppose to a 2-octet AS number that is encoded in 4 octets). When such aggregation results in creating a route that is less specific than any of the component routes, loss of the AS path information does not create a risk of a routing loop. In all other cases, loss of the AS path information does create a risk of a routing loop.

My Autonomous System Does Not Fully Support 4-octet AS Numbers and Must Peer with a Neighbor Autonomous System That Does Support 4-octet AS

When a BGP Autonomous System has not fully transitioned yet to BGP speakers that support 4-octet AS numbers, then peering with systems that 'do' support 4-octet AS numbers must be taken into special consideration. This section has two assumptions (1) My Autonomous System Number is a 2-octet AS number and (2) the border BGP router is an OLD BGP speaker.

Most commonly the neighbor Autonomous System has an AS# that can be encoded in 2-octets (ie: 0.100). This will be for many years, the most common scenario, because Regional Registries have only recently started with allocating true 4-octet AS numbers. In this situation, is the OLD BGP speaker from My Autonomous System configured with the 2-octet AS# of the neighboring Autonomous System? The convenient aspect of this situation is that once the OLD BGP speaker gets upgraded towards an IOS supporting 4-octet AS numbers, there is nothing that changes regarding the BGP configuration.

However, if the neighbor Autonomous System has an AS# in true 4-octet format (ie: 100.100) then this has an impact on the neighbor configuration of the OLD BGP speaker. The OLD BGP speaker can not understand a 4-octet AS number, and hence the OLD BGP speaker can only peer with a neighbor AS number of 2-octet length. To work around this limitation of the OLD BGP speaker, the well known AS_TRANS "23456" is used. The OLD speaker will believe that it peers with "23456", instead of with an incompatible 4-octet AS# (ie: AS# 100.100) while the neighboring NEW BGP speaker peers with the real AS# of the OLD BGP speaker. This works great and is perfectly fine as long as the OLD BGP speaker does not understand 4-octet AS numbers. If the OLD BGP speaker gets upgraded to support 4-octet Autonomous System numbers, then the previous peering with the existing NEW BGP speaker does not reactivate automatically. To reactivate the peering between the recently upgraded BGP router and the existing NEW BGP speaker, it will need to reconfigure the neighboring AS# on the recently upgraded BGP router, from the fake well known AS number AS_TRANS "23456", towards the real 4-octet AS# (ie: 100.100).

AS-Path Filtering

AS-Path filtering has some considerations when an Autonomous System contains OLD BGP speakers. Initially there will only be a very low number of true 4-octet AS numbers found on the Internet, hence the BGP table of the OLD BGP speakers reflects very well on the Internet. However, due to the incremental increase adoption of true 4-octet AS numbers, the BGP table on an OLD BGP speaker may become gradually more and more incorrect, because the OLD BGP

speaker does not understand either 4-octet AS numbers and the new optional transitive BGP attributes AS4-PATH and AS4-AGGREGATOR. Due to the gradually growing incorrectness of the BGP table on the OLD BGP speaker, AS-Path filtering will become increasingly troublesome.

An additional aspect gradually complicating AS-Path filtering on an OLD BGP speaker as time moves forward, is the well known AS# AS_TRANS “23456”. This well known AS number can be found within the AS-Path list of a BGP route identifying an incompatible 4-octet AS#. In essence, the OLD BGP speaker does not understand the full set of information available and cannot filter based upon 4-octet Autonomous System Numbers.

Multi-Exit Discriminator

In an environment where an Autonomous System that has OLD BGP speakers peers with two or more Autonomous Systems that have NEW BGP speakers and uses AS_TRANS (rather than having a globally unique AS number), use of Multi-Exit Discriminators by the Autonomous System with the OLD speakers may result in a situation where Multi-Exit Discriminator will influence route selection among the routes that were received from different neighboring Autonomous Systems.

NetFlow Support

In a classical 2-octet Autonomous System environment, there will be a gradually growing inconsistency by the adoption of true 4-octet Autonomous System Numbers. This inconsistency gets reflected because all true AS# 4-octet data gets consolidated and the traffic matrix is lost.

A NEW BGP speaker will get support of 4-octet AS numbers through an extension in NetFlow v9 to support it. The Cisco Feature Navigator can be used to identify if the Cisco operating system being used for the NEW BGP speaker supports NetFlow for 4-octet AS numbers.

For More Information

References used to compose this paper are:

- Cisco IOS Software 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, CCDE, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IQS, 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, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, IPPhone, iQuick Study, iSeePart, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanel, ProConnect, SmartShore, SenderBase, 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. (081216)