

Migration Guide for Explaining 4-Byte Autonomous System Numbers for Cisco IOS

Motivation: Guidance on transforming a BGP network to support 4-Octet Autonomous System (AS) numbers

Last updated: January 2009

What You Will Learn

This paper will specify a procedure and considerations to be used by network operators to enable the support of 4-Octet Autonomous System Numbers for BGP routers.

Why Should You Care About 4-Byte 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-byte AS number, provides a pool of 65536 unique Autonomous System numbers. 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 Autonomous System Numbers (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-byte AS number towards a 4-byte AS number. Work for this extension was started in 2001, and has been standardized in IETF RFC 4893. It provides a theoretical 4,294,967,296 unique AS numbers.

An important aspect is the consideration regarding backward compatibility and interoperability between the 'old' 2-byte AS numbers and the 'new' 4-byte AS numbers. The IETF and the standardization of the revised 4-byte AS number has been developed in such a way that there is a high degree of backward compatibility between the 4-byte AS number and 2-byte AS number. This backward compatibility is achieved by creating a set of new BGP capabilities, attributes, communities and a well known transition Autonomous System AS_TRANS "23456".

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

- A New 4-Byte BGP Capability Advertisement
- New Update Attributes for 2-byte Transit
- New Extended Communities

These new attributes, extended BGP communities and BGP capability advertisement need to be understood when deploying routers and devices which are 4-Byte AS capable.

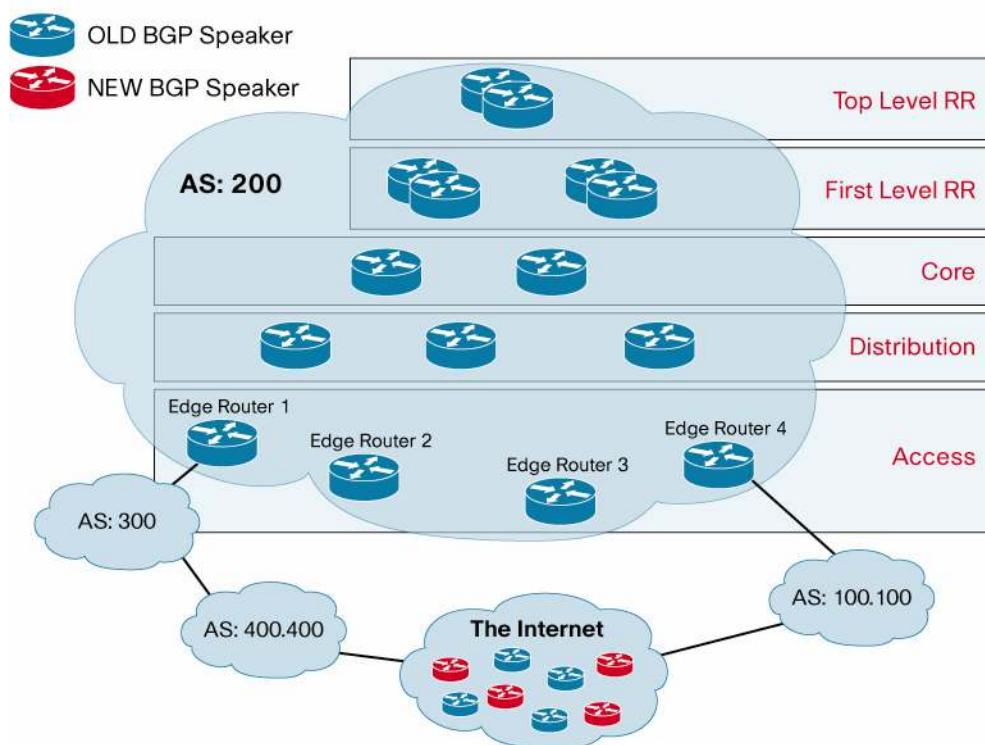
This paper will go from the assumption that the operator has an existing network implementing support for traditional 2-Byte Autonomous System Numbers (ASN). It will be required to upgrade the BGP Routers to prepare the network for the full BGP feature richness in a 4-Byte ASN situation. Note that 4-Byte ASN is backward compatible with 2-Byte ASN, however, not all BGP features regarding BGP table visibility and AS-Path filtering will be available. This paper explains the steps to upgrade the BGP speakers of a network in the most operational optimized procedure.

Steps to Upgrade a BGP Network to Support 4-Byte ASN

To upgrade a network to enjoy the full 4-Byte ASN BGP feature support, all BGP speakers within the network will have to understand 4-Byte Autonomous System Numbers. Obviously upgrading the BGP speakers is not a task that can be done without careful planning and operational downtime studies.

This migration guide provides guidelines for a successful migration, with direct feedback mechanisms to allow a secure fallback safety procedure if a router upgrade would result in undesired routing behavior. It is recommended to upgrade the BGP speakers to support 4-Byte ASN one-by-one, especially for the first set of BGP routers. This allows the ability to quickly reverse the upgrade, if necessary.

A reference network to illustrate the implementation of 4-Byte ASN support can be found below.



The diagram shows a traditional network running BGP and has Autonomous System Number 200 (a traditional 2-Byte ASN). The network is designed based upon a variety of layers and external connectivity options. For illustrational purposes, it is assumed that 'Edge router 1' peers with a traditional 2-Byte AS network, while 'Edge Router 4' peers with a 4-Byte ASN system. This difference will have considerations when upgrading the IOS towards a version supporting 4-Byte ASN. In the diagram, is also a hierarchical setup of Route-Reflectors (RR) seen. In this example, a two level hierarchy is used, however, the procedure will work for a single layer of RR or for more layers of RR. When the Edge Routers are also used as a VPN concentrator, then the same considerations for upgrading the IOS version to support 4-Byte ASN must be taken as for the above Interconnectivity model.

The High-Level Approach to Integrate 4-Byte ASN

An assumption is that the devices will not use the 4-Byte ASN information as long as all BGP routers have gained 4-Byte ASN support.

There are two major ways to add in a structured approach 4-Byte ASN technology to the network components. One can implement 4-Byte ASN Top-Down or either Bottom-Up. Either from Top Level RR towards the Edge Routers, or one can start from the Edge Routers and work the way up towards the Top Level RR. Both will result in a network fully supporting 4-Byte ASN, and in both models, the 4 byte ASN implementation is quite operationally structured and secure.

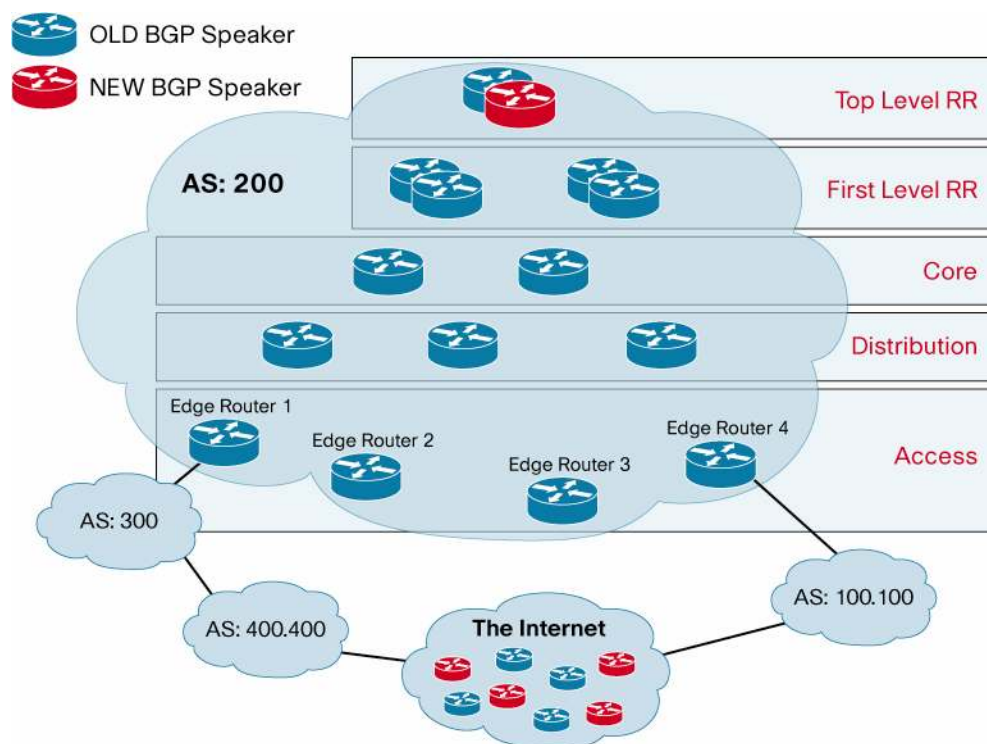
There is a small advantage to start with the Top Level RR:

- It allows to see on each router software upgrade the direct successful or failure effect
- The Top Level Route-Reflector has all information from the full 4-Byte AS-Path
- Each BGP speaker can be upgraded one-by-one and understand the impact immediately

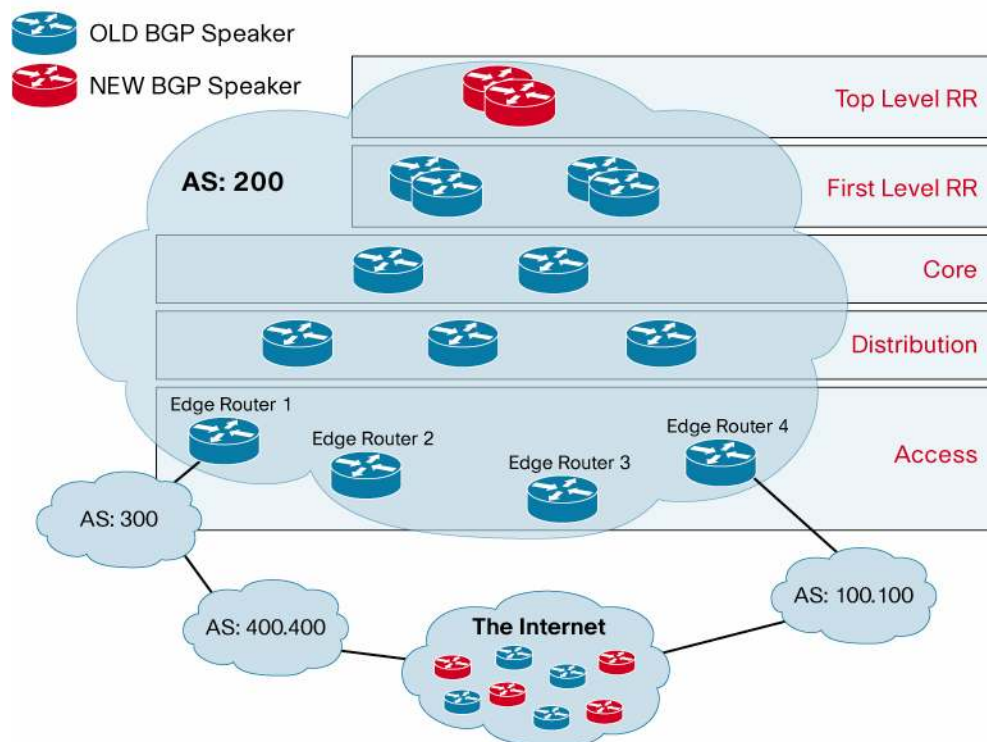
Another important aspect is that there is no configuration required to enable 4-Byte AS on the BGP router. By default, a Cisco IOS router has 4-Byte ASN enabled if the version of software supports the technology. This will make the inclusion of 4-Byte AS logistics very simple, because a majority of the BGP routers will get 4-Byte AS support silently, due to a BGP router upgrade.

Detailed 4-Byte ASN Implementation

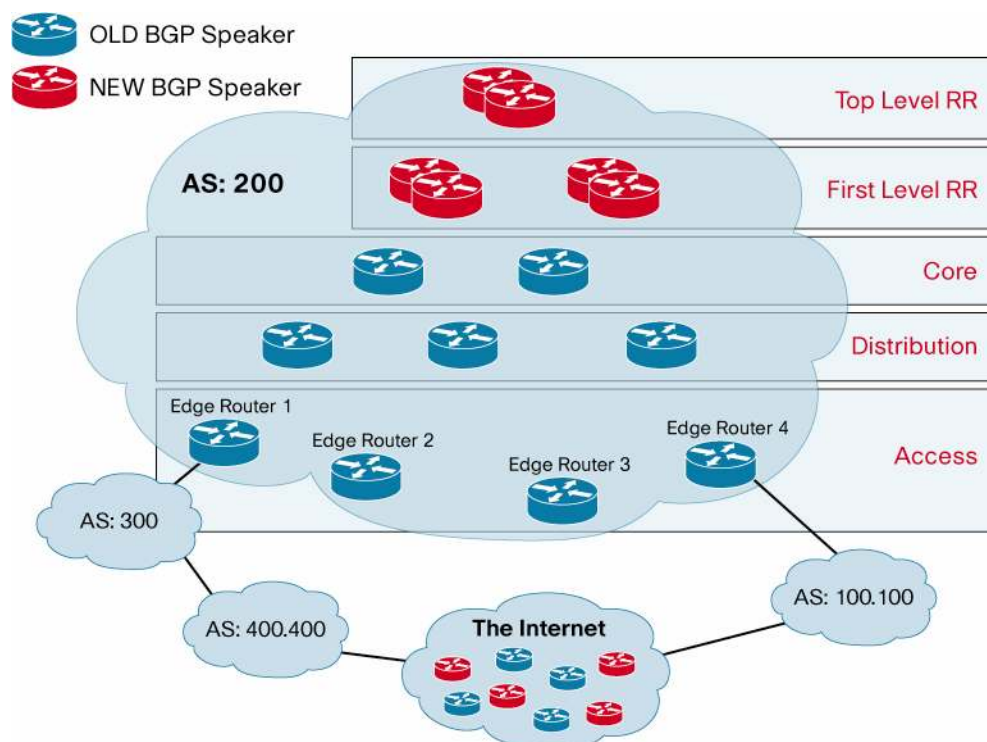
The preferred procedure described is to start with the Top Level Route-Reflector. In many cases, the Route-Reflector setup is highly redundant and exists out of a set of two RR, which work in a high-available setup. During a first software upgrade, only a single Top Level RR should be upgraded, followed by a period to monitor the Route-Reflector behavior for stability and understanding potentially unexpected behavior.



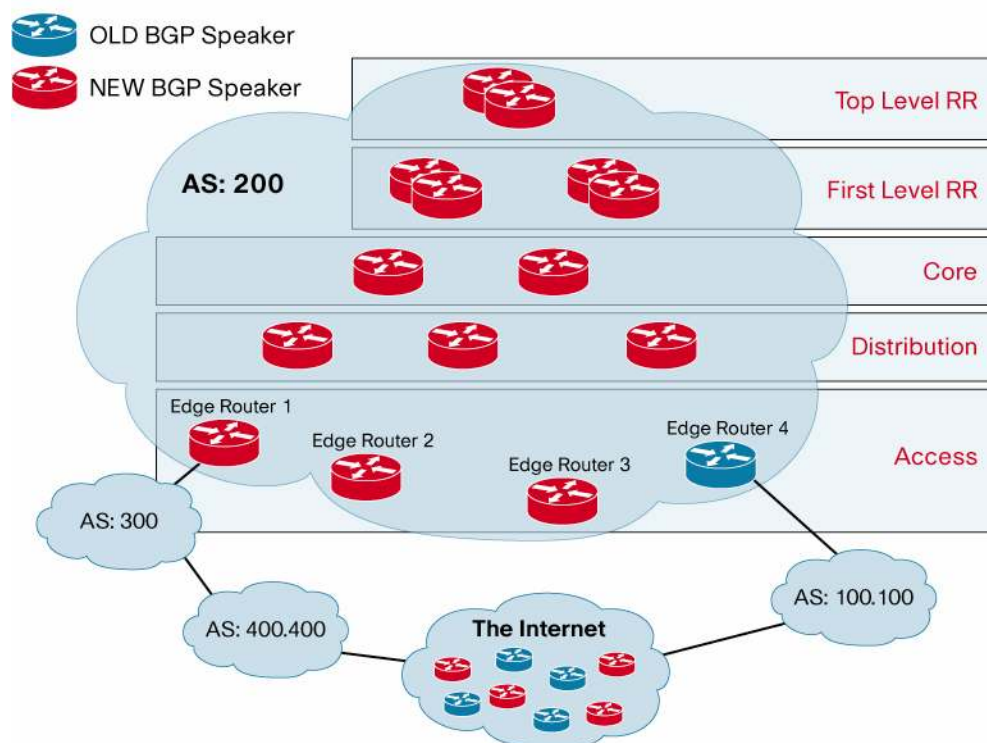
Once the situation is found stable and the operational behavior is found satisfying, then the second Top Level Route-Reflector should be upgraded.



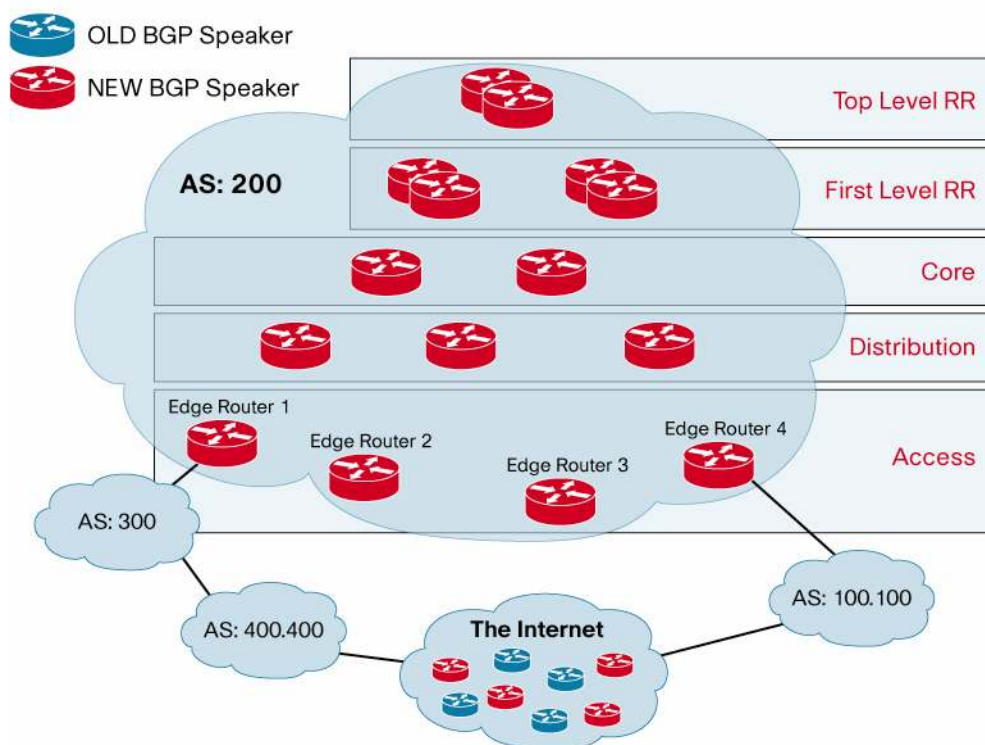
The next step to migrate the network towards 4-Byte ASN support is to upgrade the next Level of BGP Route-Reflectors one-by-one, until all Route-Reflectors have been upgraded.



Then, the best approach is to upgrade the routers within the different layers of the network one-by-one, until the access layer has been reached. Special care has to be taken if an Edge Router has an eBGP peering with an Autonomous System with a true 4-Byte ASN (ie: Edge Router 4). For the other Edge Routers, no configuration changes are required.



For those Edge Routers (ie: Edge Router 4) that have an eBGP peering with a true 4-Byte AS Number, there is a need for a configuration change. Because the Edge Router does not understand a 4-Byte AS Number, it will, for compatibility reasons, have an eBGP session using the special purpose BGP AS_TRANS Autonomous System Number “23456”. This Autonomous System is used by the neighborhood between OLD (traditional 2-Byte BGP router) and NEW BGP (a BGP router supporting 4-Byte AS Numbers) speakers. In the above diagram, Edge Router 4 will need to be reconfigured for the neighboring AS from AS_TRANS “23456” (the special AS number used for backward compatibility) into its true 4-Byte AS Number, which is in the example case, AS Number 100.100. Only when this configuration change has been done, the eBGP neighborhood will go “UP” once the Edge Router has been upgraded to support 4-Byte AS Numbers.



Conclusion

Migrating from a BGO infrastructure supporting 2-Byte Autonomous System Numbers, to an infrastructure understanding 4-Byte Autonomous System Numbers can be achieved in a structured and operationally secure procedure, without major network redesign and operational enhancements.

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, CCENT, 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 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, FormShare, GigaDrive, HomeLink, Internet QuikStart, IOS, iPhone, iQuikStart, iStartPart, the iStartPart logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanel, ProConnect, SmartShare, SenderBase, SNAKit, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet QuikStart, 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)