Software Engineering 4C03

Border Gateway Protocol

By: Kenneth Tsui (9701592)
Date: March 19, 2001

What is BGP?

Computer scientists use the term Exterior Gateway Protocol (EGP) to denote any protocol used to pass routing information between two autonomous systems.
Border gateway protocol (BGP) is replacing EGP as an exterior protocol of choice. BGP exchanges reachability information between autonomous systems, but provides more capabilities then EGP.

BGP is designed to efficiently manage a large, multi-organization routing table, such as the global Internet routing table. This provides the level of control to manage a heterogeneous environment where there is no single organization managing all parts of the network (it's a network of autonomous systems). None of the interior routing protocols were designed to perform this task nor can they efficiently manage the large number of routes on the Internet. RIP (and other interior routing protocols) will send out updates listing all the networks it knows about every few minutes (depending on protocol). This consumes valuable bandwidth. BGP using the Transmission Control Protocol (TCP) and send updated router table information only when one host has detected a change.

When do we need BGP?

BGP is an Exterior Gateway Protocol. As in, your routers speak BGP to routers in another Autonomous System (AS), generally because of policy and control issues. The other AS might be your Internet Services Provider (ISP), but it also might be another company.

Another use of BGP is to break a large network up into smaller pieces. For instance, a company with a presence on each continent might consider using BGP between continents.

What is autonomous system?

A group of networks and routers controlled by a single administrative authority is called an autonomous system.

eBGP and iBGP

BGP uses a formally registered Autonomous System number, at least with entities that will advertise the AS number to the Internet. If your router and its BGP neighbor have different AS numbers, they're speaking external BGP, eBGP. When two BGP neighbors have the same AS number, they're speaking internal BGP, iBGP.

Functionally, there is no difference between eBGP and iBGP, iBGP is just used to exchange route information among multiple routers in the same AS. The range of private BGP AS numbers is 64512-65535.

BGP characteristics

Next-Hop Paradigm: Like distance-vector routing protocols, BGP supplies new hop information for each destination.

Policy Support: BGP can implement policies that the local administrator chooses. A router running BGP can be configured to distinguish between the set of destinations reachable by computers inside its autonomous system and the set of destinations advertised to other autonomous system.

Reliable Transport: BGP uses TCP as a reliable transport medium and so it needs only to send out updates when necessary rather than continuously.

Path information: in addition to specifying destinations that can be reached and a next hop for each, BGP advertisements include path information that allows a receiver to learn a series of autonomous systems along a path to the destination.

Classless Addressing: BGP-4 makes it easy to use Classless Inter-Domain Routing (CIDR), which is a way to have more addresses within the network than with the current Internet Protocol address assignment scheme.

Route Aggregation: BGP conserves network bandwidth by allowing a sender to aggregate route information and send a single entry to represent multiple, related destinations.

Authentication: BGP allows a receiver to authenticate messages.

Path attributes

Path attributes may include, for example, administrative preferences based on political, organizational, or security considerations in the routing decision.

On connection start, BGP peers exchange complete copies of their routing tables, which can be quite large. However, only changes (deltas) are then exchanged, which makes long running BGP sessions more efficient than shorter ones.

The routing table contains a list of known routers, the addresses they can reach, and a cost metric associated with the path to each router so that the best available route is chosen.

BGP uses a predefined set of criteria when selecting routes to enter in the routing table. A simplified description of the path selection criteria used in BGP routers can be summarized as follows:

- Select the path with the highest administrative weight.
- If weights are equal, prefer the route with the highest local preference.
- If the local preferences are the same, prefer the route originated by this router.
- If none of the routes originated from this router, prefer the route that passes through the fewest AS.
- Beyond this, there are some BGP attributes that are examined as final tiebreakers, ultimately ending with preference being given to the route originating from the router with the lowest IP address.

Reference:

http://www.ccci.com/learn/welcher/papers/bgp1.htm
http://www.networksorcery.com/enp/protocol/bgp.htm
http://www.tech-nic.net/html/bgp.html
http://www.freesoft.org/CIE/Topics/88.htm
http://info.connect.com.au/docs/routing/general/multi-faq.shtml#q1
http://www.networkcomputing.com/813/813ws2.html
http://vrienden.pandora.be/thor/info/bgp.html