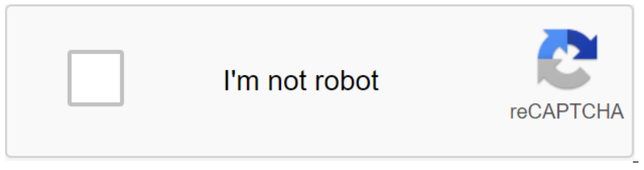

Continue

Bgp protocol configuration pdf

The protocol of routing information on the Internet IS redirects here. For other purposes, see BGP (disambigation). Internet Protocol Application Layer BGP DHCP DNS FTP http https https LDAP MGCP MZTT NNTP NTP POP PTP ONC/RPC RTP RTSP RIP SIP SMTP SNMPnet Tel TLS/SSL XMPP details... The transport layer TCP UDP DCCP SCTP RSVP more details... Internet layer IP IPv4 IPv6 ICMP ICMPv6 ECN IGMP IPsec more ... Link layer ARP NDP OSPF Tunnels L2 PPPTP MAC Ethernet Wi-Fi DSL ISDN FDDI more... vte Border Gateway Protocol (BGP) is a standardized external gateway protocol designed to share routing and accessibility information between autonomous systems (AS) on the Internet. BGP is classified as a vector route routing protocol, and it makes routing decisions based on paths, network policies, or sets of rules set up by a network administrator. BGP used for routing in an autonomous system is called the Internal Frontier Gateway Protocol, Internal BGP (iBGP). In contrast, the internet protocol application is called the External Boundary Gateway Protocol, an external BGP (eBGP). The history of the Border Gateway Protocol has been used on the Internet since 1994. IPv6 BGP was first identified in RFC 1883 in 1995, and it was upgraded to RFC 2283 in 1998. The current version of BGP is Version 4 (BGP4), which was published as RFC 4271 in 2006. RFC 4271 corrected errors, clarified ambiguities, and updated the specification with general industry practices. The main improvement was support for classless inter-domain routing (CIDR) and the use of route aggregation to reduce the size of routing tables. The new RFC allows BGP4 to carry a wide range of IPv4 and IPv6 address families. It is also called Multiprotocol Extensions, which is a Multiprotocol BGP (MP-BGP). Neighbors of the BGP operation, called peers, are set by manual configuration between routers to create a TCP session in port 179. The BGP speaker sends 19-byte keep-alive messages every 60 seconds to maintain the connection. Among routing protocols, BGP is unique in using TCP as a transport protocol. When BGP works between two peer-to-peer nodes in the same standalone system (AS), it is called The Internal BGP (i-BGP or Internal Border Gateway Protocol). When it works between different autonomous systems, it is called an external BGP (eBGP or external boundary gateway protocol). Routers on the border of one AS exchange information with another AS are called border or edge routers or simply eBGP nodes and are usually directly connected, while i-BGP nodes can be interconnected through other intermediate routers. Other topology deployments are possible, such as launching an eBGP peering into a VPN tunnel, allowing two sites share routing information in a safe and isolated way. The main difference between iBGP and eBGP is looking at how routes are derived from a single peer other peers. For example, new routes extracted from the peer-to-peer version of eBGP are usually redistributed to all iBGP colleagues, as well as all other eBGP colleagues (if transit mode is enabled on the router). However, if the new routes are learned on iBGP peering, then they are re-advertised only to all eBGP colleagues. These route distribution rules effectively require that all iBGP nodes within AS be interconnected in a full grid. How the routes are distributed can be monitored in detail with the help of the route map mechanism. This mechanism consists of a set of rules. Each rule describes for routes that meet certain criteria what action should be taken. The action may be a drop route, or it may be to change some route attributes before inserting it into the routing table. Negotiations to extend the state of the BGP machine During the handshake, when open messages are exchanged, bgP speakers can negotiate additional session capabilities, including multi-protocol extensions and various recovery modes. If multi-protocol extensions to BGP are agreed upon at the time of creation, the BGP speaker can shove information about the availability of the network layer (NLRI), which it advertises with the family address set-top box. These families include IPv4 (default), IPv6, IPv4/IPv6 Virtual Private Networks and Multicast BGP. Increasingly, BGP is being used as a generalized signaling protocol to transmit information about routes that may not be part of the global Internet, such as VPN. Connecting Active; OpenSent; OpenConfirm; and created. For each peer session, BGP supports a state variable that tracks which of the six states the session is in. BGP identifies messages that each colleague must exchange to change a session from one state to another. The first state is the state of Idleness. In Idle, BGP initiates all resources, opts out of all incoming BGP connectivity attempts, and initiates TCP's connection to peer-to-peer. The second state is Connect. In Connect, the router waits for the TCP connection to be completed and openSent to be completed if successful. If it fails, it launches the ConnectRetry timer and goes to Active after the expiration date. In active state, the active router resets the ConnectRetry timer to zero and returns to Connect state. In OpenSent, the router sends an open message and waits for it in return to go to OpenConfirm. Keepalive messages are exchanged and, after successfully receiving, the router is placed in the State Installed. In the state installed router can send/receive: Keepalive; Update; and Notification messages in/from your peer. Idle state: Refuse initiating event triggers. Initiates TCP connection with BGP-tuned peer. Listens to the TCP connection from his peer. Changes its state to connect. If an error occurs in any state of the FSM process, the BGP session stops immediately and goes back to idle state. Some of the reasons why the router doesn't progress out of idle state: TCP Port 179 is not open. The occasional TCP port over 1023 is not open. The peer-to-peer address is set incorrectly on any router. The AS number is set incorrectly on any router. Connection state: Waiting for successful TCP negotiations with peers. BGP doesn't spend much time in this state if the TCP session was successfully created. Sends an open message to a peer-to-peer network and changes it to OpenSent. In the event of an error, BGP goes to Active. Some reasons for the error: TCP Port 179 is not open. The occasional TCP port over 1023 is not open. The peer-to-peer address is set incorrectly on any router. The AS number is set incorrectly on any router. Active state: If the router fails to install a successful TCP session, it falls into Active state. BGP FSM tries to restart another peer-to-peer TCP session and, if successful, then sends an open message to a colleague. If it fails, the FSM is dumped into a downtime. Repeated failures can result in a cyclist riding the router between Idle and Active states. Some of the reasons for this include: TCP Port 179 is not open. The occasional TCP port over 1023 is not open. BGP configuration error. Network overload. A flapping network interface. OpenSent State: BGP FSM listens to an open message from its colleague. Once the message is received, the router checks the veracity of the open message. If there is an error, this is because one of the fields in the open message does not correspond between the nodes, for example, the non-conformity of the BGP version, the peer-to-peer router expects another My AS, etc. If there is no error, a keepalive message is sent, different timers are set and the state has been changed to OpenConfirm. OpenConfirm State: A peer listens to a Keepalive message from his colleague. If the Keepalive message is received and the timer does not expire before Keepalive is received, BGP goes into a set state. If the timer expires before the Keepalive message is received, or if there is an error, the router returns to idle state. State: In this state, nodes send updates to share information about each route advertised to a BGP colleague. If there is an error in the Update message, the peer-to-peer query is sent to a peer-to-peer message, and BGP returns to a downtime state. Router and learning routes In the simplest arrangement, all routers within the same AS and participation in the routing of BGP must set up in a full grid: Each router must be configured as a peer-to-peer for any other router. This causes scaling problems as the number of connections required grows into four parts with the number of routers involved. To alleviate the problem, BGP implements two options: route reflectors (RFC 4456) and BGP (RFC 5065). The next discussion of the BASIC UPDATE processing involves a full iBGP grid. This BGP router can accept UPDATEs network availability information (NLRI) from multiple neighbors and advertise NLRI with the same or other set of neighbors. Conceptually, BGP maintains its own master routing table, called the Local Routing Information Base (Loc-RIB), separate from the router's main routing table. For each neighbor, the BGP process supports a conceptual information base for related routing (Adj-RIB-In) containing NLRI from a neighbor, and a conceptual Adj-RIB-Out (Outgoing) for NLRI that will be sent to a neighbor. Conceptual, in the previous paragraph, means that the physical storage and structure of these different tables is decided by the BGP code executor. Their structure is not visible to other BGP routers, although they can usually be questioned with control commands on a local router. It is quite common, for example, that two Adj-RIB and Loc-RIB are stored together in the same data structure, with additional information attached to RIB records. Additional information informs the BGP process of things such as whether individual entries belong to Adj-RIBs to specific neighbors, whether the peer-to-peer route selection process made the received policies eligible for Loc-RIB, and whether Loc-RIB records can be submitted in the local router routing table management process. Under the right to be represented, BGP will present the routes it considers most appropriate for the main routing table process. Depending on the implementation of this process, the BGP route is not necessarily chosen. For example, a directly connected set-top box extracted from your router's own hardware is usually the preferred. As long as this directly connected route interface is active, the BGP route to its destination will not be placed in the routing table. Once the interface is out of the game and there are no longer preferred routes, the Loc-RIB route will be installed in the main routing table. Until recently, it was a common mistake to say that BGP was pursuing a policy. BGP actually carried out information with which the rules inside BGP-speaking routers could make policy decisions. Some of the information that is clearly intended to be used in policy decisions is community and Discrimination (IER). The BGP standard identifies a number of decision-making factors, more than those used in any other overall routing process, for selecting the NLRI to enter the Loc-RIB. Teh Teh the solution for evaluating NLRI is that its next hop attribute should be achievable (or a highlight). Another way of saying that the next transition should be achievable is that there should be an active route, already in the main router routing table, to the set-top box, in which the next hop address is achievable. Further, for each neighbor, the BGP process applies different standard and implementation-dependent criteria to decide which routes should conceptually enter Adj-RIB-In. A neighbor can send several possible routes to its destination, but the first level of preference is at the neighbor's level. Only one route to each destination will be set in the conceptual Adj-RIB-In. This process will also remove from Adj-RIB-In any routes that are removed by the neighbor. Whenever the conceptual Adj-RIB-In changes, the basic PROCESS of BGP decides if any of the new neighbor routes are preferable to routes already in Loc-RIB. If so, he replaces them. If this route is recalled by a neighbor and there is no other route to that destination, the route is removed from Loc-RIB and no longer sent to BGP by the general manager of the routing table. If the router does not have a route to that destination from any source that is not BGP, the recalled route will be removed from the main routing table. After verifying that the next transition is achievable if the route comes from an internal (i.e. iBGP) peer, the first rule that applies, according to the standard, is to study the LOCAL_PREFERENCE attribute. If there are several iBGP routes from your neighbor, you will choose one with the highest LOCAL_PREFERENCE if there are no multiple routes with the same LOCAL_PREFERENCE. In the latter case, the route selection process moves on to the next tie-break. While LOCAL_PREFERENCE is the first rule in the standard once the availability of NEXT_HOP is verified, Cisco and a number of other providers first consider a decision factor called WEIGHT, which is local to the router (i.e. not transmitted by BGP). The route with the highest WEIGHT is preferable. You can LOCAL_PREFERENCE, WEIGHT, and other criteria that can be manipulated with local configuration and software capabilities. Such manipulations go beyond the standard, but are widely used. For example, the COMMUNITY attribute (see below) is not directly used by the BGP selection process. However, in the neighbor's process, BGP may have a rule that LOCAL_PREFERENCE or other factor based on a manually programmed rule to set up an attribute if the COMMUNITY value corresponds to a pattern matching criterion. If the route has been extracted from the external peer-to-peer BGP process, it calculates the value of the LOCAL_PREFERENCE from local policy rules and then compares the LOCAL_PREFERENCE of all From a neighbor. At the per-neighbor level - ignoring policy modifiers specific to implementation - the order of communication violations: I prefer the route route shortest AS_PATH. The AS_PATH set as numbers that must be passed to get to the advertised destination. AS1-AS2-AS3 shorter as AS4-AS5-AS6-AS7. Prefer routes with the lowest ORIGIN attribute. Prefer routes with MULTI_EXIT_DISC (multi-profile discriminator or MED). Until the last edition of bgP, if UPDATE didn't matter MULTI_EXIT_DISC, several implementations created a MED with the highest possible value. The current standard, however, indicates that missing MED should be treated as the lowest possible value. Because the current rule may cause a different behavior than the vendor's interpretations, BGP implementations that used a non-standard default have a configuration function that allows you to choose an old or standard rule. After receiving routes from neighbors, Loc-RIB software applies additional tie-breaks to routes to the same destination. If at least one route has been explored from an external neighbor (i.e. the route was extracted from eBGP), discard all routes recovered from iBGP. Prefer the route with the lowest internal cost to NEXT_HOP, according to the main routing table. If two neighbors advertised the same route, but one neighbor can be reached by a low bitrate link and the other by a high bitrate link, and the interior routing protocol calculates the lowest cost based on the highest bitreit, the route through a high-bitrate link will be preferable, and other routes will be removed. If there are still multiple routes connected at the moment, several BGP implementations offer a customizable option to download between routes, taking everything (or all to a certain number). Prefer a route extracted from the BGP speaker with numerically low BGP ID Preferred route extracted from BGP speakers with the lowest contact IP address of the BGP Community Community attribute tags that can be applied to incoming or outgoing prefixes to achieve some common goal (RFC 1997). While it can often be said that BGP allows an administrator to set policies on how prefixes are handled by providers, this is generally not possible, strictly speaking. For example, BGP at home does not have the concept to allow one AS to tell another AS to limit prefix advertising to only North American peering customers. Instead, the provider typically publishes a list of known or nonfree communities with a description for each of them, which essentially becomes an agreement on how prefixes should be considered. RFC 1997 also identifies three well-known communities that are of global importance; NO_EXPORT, NO_ADVERTISE and NO_EXPORT_SUBCONFED. RFC 7611 defines ACCEPT_OWN. Examples of common communities include local preference adjustments, geographic or type restrictions, DoS avoidance (black holing) and AS pre-retirement options. The provider can say that any routes received from customers with the community XXX:500 XXX:500 will be advertised for all peers (by default), while the XXX:501 community will restrict advertising in North America. The customer simply adjusts its configuration to include the right community or community for each route, and the provider is responsible for controlling who the prefix is advertised to. The end user does not have the technical ability to ensure that the provider takes the right action, although problems in this area are usually rare and accidental. This is a common tactic for end customers to use BGP communities (usually ASN:70,90,90,100) to control local provider preferences assigning advertised routes rather than using MED (the effect is similar). The community attribute is transit, but the community applied by the client is very rarely distributed outside of the next AS hop. Not all providers give out their communities to the public, while some others do. The extended format consists of one or two octets for a type field followed by seven or six octets for the relevant content of the community attribute. The definition of this extended community attribute is documented in RFC 4360. IANA manages the BGP Extended Community Registry. Extended community attribute is itself a transit supplement of BGP. However, a little in the type field in the attribute decides whether the coded extended community is transitive or non-useful. Therefore, the IANA registry provides different number ranges for attribute types. Because of the extended range of attributes, its use can be diverse. RFC 4360 exemplary defines The October 2 as a specific extended community, IPv4 Address of a specific extended community, opaque extended community, Route Target Community, and Community Origin Route. A number of BGP oS projects also use this extended structure of community attributes for inter-domain signaling. Note: With RFC 7153, extended communities are compatible with 32-bit ASNs. With the introduction of 32-bit AS numbers, some problems were immediately obvious with the community attribute, which identifies only 16 bits of ASN field, which prevents a comparison between this field and the real value of ASN. That's why RFC 8092 and RFC 8195 introduce a Large Community attribute of 12 bytes, divided into three fields of 4 bytes each (AS:function:parameter). Multi-exit discriminators MEDs, defined mainly by bgP standard, were originally intended to show another AS neighbor a preference for AS advertising as to which of the few links is preferable to incoming traffic. Another application of MEDs is advertising values, usually based on delay, several AS that have a presence on IXP that they impose to send traffic Destination. The headline message format Is the next headline format of the BGP Version 4 message: bit is compensated 0-15 16-23 24-31 0 Marker 32 64 96 128 Length Type Marker: Included for compatibility, must be installed for all of them. Length: Total length of message in octets, including title. Type: BGP message type. The following values are defined: Open (1) Update (2) Notice (3) KeepAlive (4) Route-Refresh (5) BGP's Internal Scalability is the most scalable of all routing protocols. An autonomous system with an internal BGP (iBGP) should have all of its iBGP nodes connected to each other in a full grid (where everyone speaks to everyone directly). This full grid configuration requires that each router maintain a session for any other router. In large networks, this number of sessions can impair router performance due to lack of memory or high processor requirements. Route reflectors reduce the number of connections required in AS. One router (or two for redundancy) can be made by a route reflector: other routers in AS should be configured only as peer-to-peer for them. The route reflector offers an alternative to the logical requirement for a full network of internal border gateway protocol (IBGP). RR serves as a focal point for IBGP sessions. The goal of RR is concentration. Multiple BGP routers can work with a central point, RR - acting as a route reflector server - instead of working with each other router in a full grid. All other IBGP routers become customers of the route reflector. This approach, similar to the DR/BDR OSPF, provides large networks with additional IBGP scalability. In a fully networked IBGP of 10 routers, 90 individual CLI operators (distributed across all routers in topology) are only needed to identify the remote AS of each peer: it is fast becoming a management headache. Topology RR can reduce these 90 applications to 18, offering a viable solution for larger networks managed by providers. The route reflector is one point of failure, so at least a second route reflector can be configured to ensure redundancy. Since it's an extra peer-to-peer for the other 10 routers, it comes with an additional statement account to double minus 2 of one Route Reflector installation. Additional statements are 11'2-2'20 in this case due to the addition of an additional router. In addition, in a multipathic BGP environment, this can also benefit by adding local-switching/bandwidth if RRs act as traditional routers, rather than simply as a dedicated Route Reflector Server role. Rules BGP Route Reflector Deployment configuration proposed by Section 6, RFC 4456. Rr servers distribute routes within AS based on the following rules: If the route is received from a non-customer, reflect only EBGP customers and colleagues. If the route is received customer colleague, reflects on all non-clients, as well as on fellow clients, with the exception of the route maker and affects EBGP colleagues. The RR cluster and its customers form a cluster. Cluster-ID is then attached to each route advertised by RR to its customer or non-client colleagues. Cluster-ID is a cumulative, impassable BGP attribute, and each RR must pre-use CLUSTER_ID on CLUSTER_LIST to avoid routing cycles. Route and conferlier reflectors reduce the number of iBGP nodes for each router and thus reduce processing overheads. Route reflectors are pure performance-enhancing techniques, while confederations can also be used to implement more fine-grained policies. BGP Confederations are sets of autonomous systems. In normal practice, only one of the AS numbers of the confederation is seen by the Internet as a whole. The Confederacys are used in very large networks where large AS can be configured to cover smaller, more manageable internal ASs. Confederate AS consists of several ASs. Each Confederate AS only has an iBGP completely netted and has ties with other ASs within the Confederation. Although these ASS have eBGP peers ASs within the confederation, ASs sharing routing as if they were using iBGP. Thus, the confederation retains the following information about the transition, metric and local preferences. For the outside world, the confederation seems to be one AS. With this solution, the iBGP AS transit problem can be solved because iBGP requires a complete grid between all BGP routers: a large number of TCP sessions and unnecessary duplication of routing traffic. The Confederacy can be used in conjunction with route reflectors. Both confederations and route reflectors may be subject to constant fluctuations unless specific design rules affecting both BGP and interior routing protocol are followed. However, these alternatives may introduce their own problems, including the following: sub-optimal growth of the route routing route time of the BGP convergence In addition, route reflectors and BGP confederation were not designed to facilitate the configuration of the BGP router. However, these are common tools for experienced architects of the BGP network. These tools can be combined, for example, as a hierarchy of route reflectors. The Stability routing tables managed by BGP are constantly adjusted to account for actual network changes, such as breaking and restoring links or moving routers and restoring them. In the network as a whole it is normal for these changes to occur almost continuously, but for any particular router or link, changes should be relatively rare. If the router is wrong or incorrectly configured, it can get into a fast loop between states down and up. This re-inference and re-announcement model, known as route slamming, can lead to excessive over-re-introduction in all other routers that are aware of the broken link, as the same route is constantly entered and removed from the routing tables. The design of BGP is such that traffic delivery may not function during route upgrades. On the Internet, changing BGP routing can cause disruptions within minutes. The feature, known as route valve damping (RFC 2439), is built into many BGP implementations in an attempt to mitigate the effects of the route slamming. Without damping, excessive activity can put a heavy strain on routers, which in turn can delay upgrades on other routes and thus affect the overall stability of routing. With the damping, the flapping route decomposes exponentially. First of all, when the route becomes inaccessible and appears quickly, the damping is not in effect in order to maintain the normal time of the BGP failure. In the second case, BGP avoids this prefix for a certain period of time; subsequent cases of sukm exponentially. Once the anomalies have stopped and the time has passed for the offensive route, the prefixes can be restored and the slate wiped. Damping can also mitigate denial-of-service attacks; Damping timing is very customizable. It is also proposed in RFC 2439 (according to Design Choice - Stability Sensitive Suppression Route Advertising) that route flap damping is a function more desirable if implemented for external boundary gateway protocol sessions (eBGP sessions or simply called external peers) rather than at internal border gateway protocol sessions (iGPB sessions or simply called internal peers); With this approach, when the route flaps inside the autonomous system, it does not extend to external ASs - the slapping route to eBGP will have a chain flapping for a specific route across the spine. This method also successfully avoids the overhead cost of damping the route valve for iBGP sessions. However, subsequent studies have shown that sealing the flap can actually lengthen the convergence time in some cases, and can cause communication disruptions even if the links do not slam. Moreover, as backbone links and router processors became faster, some network architects have suggested that sealing flaps may not be as important as it used to be, as changes in the routing table can be handled by routers much faster. This led to the RIPE Routing working group writing that in the current implementations of BGP valve damping, the use of sealing of flaps in ISP networks is NOT recommended. ... If the flap damping is implemented, the provider working on the network will cause side effects for its and internet users of the content and services of their customers..... These side effects are likely to be worse than the impact caused simply by not running a flap of damping at all. Improving stability without valve problems is the subject of ongoing research. The growth of the BGP routing table on the Internet as an AS

internet number versus the number of registered AS One of the biggest challenges facing BGP, and the Internet infrastructure as a whole, is the growth of the Internet routing table. If the global routing table grows to the point where some older, less capable routers can't handle the memory requirements or processor load to maintain the table, these routers will no longer be effective gateways between the parts of the Internet they connect. In addition, and perhaps more importantly, it takes longer (see above) to stabilize larger routing tables after a major change in connectivity, making network maintenance unreliable or even inaccessible in the interim. By the end of 2001, the global routing table was growing exponentially, eventually threatening widespread connectivity gaps. In an effort to prevent this, ISPs collaborated to maintain the global routing table as little as possible, using class-free inter-domain routing (CIDR) and route aggregation. Although this slowed the growth of the routing table to linear process for several years, with the expansion of demand for multi-hired end-user networks, growth was once again superlinean by mid-2004. The 512k day Y2K-like overflow is triggered in 2014 for those models that have not been properly updated. While the full IPv4 BGP table from August 2014 (512k a day) is more than 512,000 prefixes, many older routers have a limit of 512k (512,000-524,288). On August 12, 2014, disruptions caused by full tables hit eBay, LastPass, and Microsoft Azure, among others. A number of Cisco routers commonly used to store advertised BGP routes have TCAM, a form of high-speed memory addressed by content. On affected routers, TCAM by default allocates 512k records for IPv4 routes, and 512k entries for IPv6 routes. While the number of advertised IPv6 routes advertised was only about 20k, the number of advertised IPv4 routes reached the default limit, causing a side effect as routers tried to compensate for the problem by slowing the software routing (as opposed to quickly routing equipment via TCAM). The main method of solving this problem involves operators changing the distribution of TCAM to allow more IPv4 records, by redistributing some of the TCAM reserved for IPv6 routes. This requires a reboot on most routers. The 512k problem was predicted in advance by a number of IT professionals. The actual appropriations that have pushed the number of above 512k was the announcement of 15,000 new routes in short order starting at 07:48 UTC. Almost all of these routes were Verizon Autonomous Systems 701 and 705, created as a result of deaggregating large blocks, introducing thousands of new /24 routes, and making do reach 515,000 entries. The new routes appeared to have been re-aggregated within 5 minutes, but the instability on the Internet appeared to have continued for a number of hours. Even if Verizon didn't cause the routing table to exceed 512k inputs in the short spike, it would happen soon anyway through natural growth. Routeization is often used to improve the aggregation of the global BGP routing table, thereby reducing the required table size in AS routers. Consider AS1 was allocated a large address space 172.16.0.0/16, this will be counted as one route in the table, but because of customer needs or traffic engineering goals, AS1 wants to announce smaller, more specific routes 172.16.0.0/18, 172.16.64.0/18, and 172.16.128.0/18. The console 172.16.192.0/18 has no hosts, so AS1 does not announce a specific route 172.16.192.0/18. All this is considered AS1 announces four routes. AS2 will see four routes from AS1 (172.16.0.0/16, 172.16.0.0/18, 172.16.64.0/18, and 172.16.128.0/18), and it's up to routing the AS2 policy to decide whether to take a copy of the four routes or how 172.16.0/16 overlaps all other specific routes to just keep a resume, 172.16.0/16. If AS2 wants to send data to the console on 172.16.192.0/18, it will be sent to AS1 routers on route 172.16.0.0/16. On the AS1 router, it will either be deleted or the ICMP message will be sent back, depending on the configuration of the AS1 routers. If AS1 later decides to abandon route 172.16.0/16, leaving 172.16.0.0/18, 172.16.64.0/18, and 172.16.128.0/18, as1 will drop the number of routes it announces to three. AS2 will see three routes, and depending on the AS2 routing policy, it will store a copy of the three routes, or aggregate the set-top box 172.16.0.0/18 and 172.16.64.0/18 to 172.16.0/17, thereby reducing the number of as2 store routes to two: 172.16.0.0/17 and 172.16.128.0/18. If AS2 wants to send data to the set-top box 172.16.192.0/18, it will be deleted or the destination ICMP message will be sent back to AS2 routers (not AS1 as before) because 172.16.192.0/18 will not be in the routing table. AS depletion numbers and 32-bit ASNs RFC 1771 (Border Gateway Protocol 4 (BGP-4)) are planned to encode AS numbers at 16 bits, for 64,510 possible public AS, since ASN 64512 to 65534 have been reserved for private use (0 and 65535 prohibited). In 2011, only 15,000 AS numbers were available, and projections were for the complete depletion of existing AS numbers in September 2013. RFC 6793 extends as coding from 16 to 32 bits (keeping the range of 16 bits as 0 to 65535, and its reserved AS numbers), which now allows up to 4 billion AC available. An additional private range of AS is also defined in RFC 6996 (with up to 4294967294, 4294967295 prohibited A new OT attribute is used AS4_PATH use to prevent router groups from managing these new ASNs. ASN's 32-bit assignments began in 2007. Balancing load Another factor causing this growth of the routing table is the need to balance the load of multi-family networks. It is not a trivial task to balance incoming traffic in a multi-stage network through several incoming paths, due to the restriction of the BGP route selection process. For a multi-home network, if it announces the same network blocks in all of its BGP colleagues, the result may be that one or more of its incoming links become overloaded, while other links remain underutilized because external networks have all chosen this set of congested paths as optimal. Like most other routing protocols, BGP does not detect congestion. To get around this problem, BGP administrators on this multi-configured network can divide a large adjacent IP address block into smaller blocks and set up a route announcement to make the different blocks look optimal on different paths, so that external networks will choose a different path to reach the different blocks of this multi-home network. Such cases will increase the number of routes, as outlined in the global BGP table. One way to increase popularity to solve the load balancing problem is to deploy BGP/LISP (Locator/ID separation protocol) at the Point of Internet Exchange to allow traffic engineering on multiple links. This method does not increase the number of routes seen in the global BGP table. Security By design, BGP-operated routers accept advertised routes from other BGP routers by default. This allows for automatic and decentralized routing of internet traffic, but also makes the Internet potentially vulnerable to accidental or malicious failures known as BGP hijacking. Because of the extent to which BGP is embedded in the main Internet systems, as well as the number of different networks managed by the various organizations that collectively make up the Internet, fixing this vulnerability (for example, by implementing the use of cryptographic keys to verify the identity of BGP routers) is a technically and economically complex issue. BgP Expansion Extension is the use of multipating - it usually requires identical MED, weight, origin, and AS-way, although some implementations provide the ability to relax AS-way checks only to expect equal lengths of track, rather than the actual AS numbers on the way are expected to match too. This can be expanded further with features such as dmzlink-bw Cisco, which allows traffic-based traffic ratios to be configured on References. The multiprotocol extension for BGP (MBGP), sometimes referred to as Multiprotocol BGP or Multicast BGP and identified in IETF RFC 4760, is a continuation (BGP) that that different types of addresses (known as address families) that will be distributed in parallel. While standard BGP only supports IPv4 unicast addresses, Multiprotocol BGP supports IPv4 and IPv6 addresses and supports single-point and multi-cast versions of each. Multiprotocol BGP allows you to share information about the topology of IP routers with multicast routers separate from the topology of conventional single-station IPv4 routers. Thus, it allows multi-tissue topology routing differs from single-tissue topology routing. While MBGP allows cross-domain multichannel routing information to be shared, other protocols, such as the Protocol Independent Multicast family, are needed to build trees and promote multi-track traffic. Multiprotocol BGP is also widely deployed in the case of MPLS L3 VPN, for the exchange of VPN labels learned for routes from customers sites on the MPLS network in order to distinguish different customer sites when traffic from other customer sites comes to the router Provider Edge (PE router) for routing. The use of BGP4 is the standard for routing the Internet and requires most Internet service providers (PROVIDER) to route between it. Very large private IP networks use BGP internally. An example would be the merging of a number of large Open Shortest Path First (OSPF) networks when the OSPF itself does not scale to the required size. Another reason for using BGP is to multi-print the network for better redundancy, either for multiple access points of one provider or for multiple providers. Implementation routers, especially small ones for use in Small Office/Home Office (SOHO), may not include BGP software. Some SOHO routers simply can't work BGP/using BGP routing tables of any size. Other commercial routers may need certain software, an image that contains BGP, or a license that allows it to do so. Open source packages that run BGP include GNU zebra, kuagga, OpenBGPD, BIRD, XORP, and Vyatta. Devices such as layer 3 switches are less likely to support BGP than devices that will be sold as routers, but high-quality Layer 3 switches can usually work under BGP. Products such as switches may or may not have restrictions on the size of BGP tables, such as 20,000 routes, much smaller than the full Internet table plus internal routes. These devices, however, can be quite reasonable and useful when used to route a smaller part of the network, such as the Confederation-AS, which represents one of several small businesses that a bgp, or a small company that announces routes to a provider but only accepts the default route and possibly a small number of aggregated routes. The BGP router, used only for a single Internet entry point, can have a much smaller routing table size (and therefore a RAM and processor requirement) than a multi-screen network. Even simple multi-breeding can have a modest routing routing Size. See RFC 4098 for providers of independent performance settings for a single BGP router convergence in plane control. The actual amount of memory required in the BGP router depends on the amount of BGP information exchanged with other BGP speakers and how a particular router stores BGP information. The router may have to store more than one copy of the route so that it can manage different policies to advertise routes and adopt in a specific neighboring AS. The term view is often used for these different policy relationships on a running router. If one router implementation takes more memory per route than another implementation, it may be a legitimate design choice, the speed of processing trade versus memory. The full IPv4 BGP table for August 2015 exceeds 590,000 prefixes. Large providers can add another 50% for domestic and customer routes. Again, depending on the implementation, individual tables can be stored for each view of different peer AS. Notable free and open source BGP implementations include: BIRD Internet routing Daemon, GPL routing package for Unix-like system. FRRouting, kuaggi fork for systems similar to Unix. GNU zebra, GPL routing kit that supports BGP4. (decommissioned) OpenBGPD, BSD-licensed implementation by the OpenBSD team. Kwagga, the GNU zebra fork for the Unix-like system. XORP, eXtensible Open Router Platform, licensed BSD routing protocol. Systems to verify BGP compliance, load or stress performance come from suppliers, such as: Agilent Technologies GNS3 open source network simulator Ixia Spirent Standards Communications documents selective route Update for BGP, IETF project RFC 1772, Application of Border Gateway Protocol in Internet Protocol (BGP-4) using SMIv2 RFC 2439, BGP Flap Route Damping RFC 2918, Route Capability for BGP-4 RF , NoPEER Community for Border Gateway Protocol (BGP) Route Management RFC 4271, Border Gateway Protocol 4 (BGP-4) RFC 4272, BGP Vulnerability Analysis RFC 4273, Definitions of Managed Objects for BGP-4 RFC 4274, BGP-4 Protocol Analysis RFC 4275, BGP-4 MIB Implementation Review RFC 4276, BGP-4 Implementation Report RFC 4277, Experience with BGP-4 RfC Protocol 4278, Maturation Time Standards Option TCP MD5 (RFC 2385) and BGP-4 Specification RFC 4456 , BGP Route Reflection - Alternative to the full grid of internal BGP (iBGP) RFC 4724, fine reboot mechanism for BGP RFC 4760, Multiprotocol extension for BGP-4 RFC 4893, BGP Support four-octet AS space number RFC 5065, Autonomous Confederation System for BGP RFC 5492, Advertising Capabilities with BGP-4 RFC 5575, RFC Flow Distribution Rules RFC 7752, Distribution link-state and road engineering information using BGP RFC 7911 , Multi-Road Advertising in BGP BGP BGP Customs Decision Process, February 3, 2017 RFC 3392, Outdated - Advertising Opportunities with BGP-4 RFC 2796, Outdated - BGP Route Reflection - Alternative to the complete grid iBGP RFC 3065, Outdated - Autonomous Confederate System for BGP RFC 1965, Outdated - Autonomous Confederation System for BGP RFC 1771, Outdated - Border Gateway Protocol 4 (BGP-4) RFC 1657, Outdated - Border Gateway Protocol 4 (BGP-4) RFC 1657, Outdated - Border Gateway Protocol 4 (BGP-4) RFC 1657, Outdated - Border Gateway Protocol 4 (BGP-4) RFC 1657, Outdated - Border Gateway Protocol 4 (BGP-4) RFC 1657, Outdated - Definitions of managed objects for the fourth version of the border gateway RFC 1655, Outdated - Application of the border gateway protocol on the Internet RFC 1654, Outdated - Border Gateway Protocol 4 (BGP-4) RFC 1105, Outdated - Border Gateway Protocol (Border Gateway) BGP) RFC 2858, Outdated - Multiprotocol Extensions for BGP-4 See also AS 7007 Incident Internet Assigned Numbers Organ Package Shipping Private IP SPPB Regional Online Registry Route Routing Asset Routing Database Links : Border Gateway Explained Protocol. Orbital-computer Solutions.Com. Archive from the original 2013-09-28. Received 2013-10-08. Sobrinho, Jose Luis (2003). Routing a network with path vector protocols: theory and application (PDF). Received on March 16, 2018. History of the border gateway protocol. blog.datapath.io. Border Gateway Protocol 4 (BGP-4). RFC 4271. BGP Keepalive Messages. INetDaemon IT lessons. RFC 4274 - Advertising Opportunities with BGP-4, RFC 2842, R. Chandra and J. Scudder, May 2000 - Multi-protocol expansion for BGP-4, RFC 2858, T. Bates et al., June 2000 - BGP/MPLS VPNs., RFC 2547, E. Rosen and Y. Rekhter, April 2004 - BGP Community. Received on April 13, 2015. IANA Registry for BGP Advanced Community Types, IANA,2008 - IETF Projects at BGP Signaled zoS Archive 2009-02-23 on Wayback Machine, Thomas Knoll,2008 - Border Gateway Protocol (BGP). Cisco.com. - BGP Route Reflection: Alternative to the Full Net Internal BGP (iBGP), RFC 4456, T. Bates et al., April 2006 - Info. www.ietf.org. received 2019-12-17. Information. www.ietf.org. received 2019-12-17. Information. www.ietf.org. received 2019-12-17. The route flap damping exacerbates internet routing convergence (PDF). November 1998. Chang, Eychuan; Pei Dan; Daniel Massey; Lisa Chang (June 2005). The interaction of the timer in the flap Damping route (PDF). IEEE 25th International Conference on Distributed Computing Systems. Received 2006-09-26. We show that the current damping design only leads to the purpose of behavior when the route is constantly clapped. When the number of flaps is small, the global routing dynamics deviate significantly from the expected behavior with a longer convergence delay. BGP Route Flap Damping. Tools.ietf.org. No. 10 May 2006 (2006-05-10). RIPE Route Working Group Recommendations on the route of flap damping. RIPE Network Coordination Center. Received 2013-12-04. To use-02 - Creating a Flap route You can eat it in use. Tools.ietf.org. Received 2013-12-04. As of August 12, 2014, several Internet routers manufactured by Cisco and other vendors had experienced a default software limit of 512K (512,000 - 524,288) of Cisco switch problem. Renesys 512k global routes. b BGP Reports. potaroo.net. - CAT 6500 and 7600 Series routers and TCAM distribution adjustment switches. Cisco. March 9, 2015. Jim Cowie. The Internet deals with half a million routes: outages are possible next week. Dean research. Garside, Juliet; Samuel Gibbs (August 14, 2014). The internet infrastructure needs updating or more outages will happen. Keeper. Received on August 15, 2014. BoF report (PDF). www.nanog.org. received 2019-12-17. Greg Ferro. TCAM - A deeper look and influence of IPv6. EtherealMind. The site of IPv4 depletion. ipv4depletion.com. - What caused today's hiccups on the Internet. bgpmon.net. - 16-bit report on the autonomy system, Jeff Houston 2011 (originally archived on //www.potaroo.net/tools/asn16/) - Craig Timberg (2015-05-31). A quick solution to the early problem of the Internet lives on a quarter of a century later. The Washington Post. 2015-06-01. GNU zebra. Further reading of the Chapter Of the Boundary Gateway Protocol (BGP) in Cisco iOS Technology Handbook External links BGP Routing Resources (includes a special section on BGP and provider Core Security) BGP statistics tables extracted from bgp routing protocol configuration. bgp routing protocol configuration in packet tracer. bgp routing protocol configuration example. bgp multi-protocol configuration. bgp protocol configuration pdf. bgp protocol configuration cisco. bgp protocol configuration commands. bgp protocol configuration cisco packet tracer