



BGP Features Roadmap

First Published: May 2, 2005
Last Updated: July 15, 2009

This feature roadmap lists the Cisco IOS features documented in the *Cisco IOS IP Routing Protocols Configuration Guide* and maps them to the documents in which they appear. The roadmap is organized so that you can select your release train and see the features in that release. Find the feature name you are searching for and click on the URL in the “Where Documented” column to access the document containing that feature.

Feature and Release Support

Table 1 lists BGP feature support for the following Cisco IOS software release trains:

- [Cisco IOS Release 12.0S](#)
- [Cisco IOS Release 12.2S](#)
- [Cisco IOS Release 12.2SB](#)
- [Cisco IOS Release 12.2SR](#)
- [Cisco IOS Release 12.2SX](#)
- [Cisco IOS Releases 12.2T, 12.3, 12.3T, 12.4 and 12.4T](#)

Use Cisco Feature Navigator to find information about platform support and software image support. Cisco Feature Navigator enables you to determine which Cisco IOS and Catalyst OS software images support a specific software release, feature set, or platform. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.



Note

Table 1 lists only the Cisco IOS software release that introduced support for a given feature in a given Cisco IOS software release train. Unless noted otherwise, subsequent releases of that Cisco IOS software release train also support that feature.

Table 1 lists the most recent release of each software train first and the features in alphabetical order within the release.



Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

© 2005–2009 Cisco Systems, Inc. All rights reserved.

Table 1 **Supported BGP Features**

| Release | Feature Name | Feature Description | Where Documented |
|--------------------------------|---|--|---|
| Cisco IOS Release 12.0S | | | |
| 12.0(32)S12 | BGP Support for 4-Byte ASN | The BGP Support for 4-Byte ASN feature introduced support for 4-byte autonomous system numbers. Because of increased demand for autonomous system numbers, in January 2009 the IANA will start to allocate 4-byte autonomous system numbers in the range from 65536 to 4294967295. | Cisco BGP Overview Configuring a Basic BGP Network |
| 12.0(31)S | BGP Route-Map Continue Support for an Outbound Policy | The BGP Route-Map Continue Support for an Outbound Policy feature introduces support for continue clauses to be applied to outbound route maps. | Connecting to a Service Provider Using External BGP |
| 12.0(31)S | BGP Support for BFD | Bidirectional Forwarding Detection (BFD) is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. In addition to fast forwarding path failure detection, BFD provides a consistent failure detection method for network administrators. Because the network administrator can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different routing protocol hello mechanisms, network profiling and planning will be easier, and reconvergence time will be consistent and predictable. The main benefit of implementing BFD for BGP is a significantly faster reconvergence time. | Configuring Advanced BGP Features |
| 12.0(29)S | BGP Support for Fast Peering Session Deactivation | The BGP Support for Fast Peering Session Deactivation feature introduced an event driven notification system that allows a Border Gateway Protocol (BGP) process to monitor BGP peering sessions on a per-neighbor basis. This feature improves the response time of BGP to adjacency changes by allowing BGP to detect an adjacency change and deactivate the terminated session in between standard BGP scanning intervals. Enabling this feature improves overall BGP convergence. | Configuring BGP Neighbor Session Options |
| 12.0(29)S | BGP Support for IP Prefix Import from Global Table into a VRF Table | The BGP Support for IP Prefix Import from Global Table into a VRF Table feature introduces the capability to import IPv4 unicast prefixes from the global routing table into a Virtual Private Network (VPN) routing/forwarding (VRF) instance table using an import route map. | BGP Support for IP Prefix Import from Global Table into a VRF Table |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|-----------|---|---|--|
| 12.0(29)S | BGP Support for Next-Hop Address Tracking | The BGP Support for Next-Hop Address Tracking feature is enabled by default when a supporting Cisco IOS software image is installed. BGP next-hop address tracking is event driven. BGP prefixes are automatically tracked as peering sessions are established. Next-hop changes are rapidly reported to the BGP routing process as they are updated in the RIB. This optimization improves overall BGP convergence by reducing the response time to next-hop changes for routes installed in the RIB. When a bestpath calculation is run in between BGP scanner cycles, only next-hop changes are tracked and processed. | Configuring Advanced BGP Features |
| 12.0(27)S | BGP Cost Community Support for EIGRP MPLS VPN PE-CE | The BGP Cost Community Support for EIGRP MPLS VPN PE-CE feature provides BGP cost community support for mixed EIGRP MPLS VPN network topologies that contain back door routes. | BGP Cost Community |
| 12.0(27)S | BGP Support for Dual AS Configuration for Network AS Migrations | The BGP Support for Dual AS Configuration for Network AS Migrations feature extends the functionality of the BGP Local-AS feature by providing additional autonomous-system path customization configuration options. The configuration of this feature is transparent to customer peering sessions, allowing the provider to merge two autonomous-systems without interrupting customer peering arrangements. Customer peering sessions can later be updated during a maintenance window or during other scheduled downtime. | Configuring BGP Neighbor Session Options |
| 12.0(27)S | BGP Support for TTL Security Check | The BGP Support for TTL Security Check feature introduced a lightweight security mechanism to protect external Border Gateway Protocol (eBGP) peering sessions from CPU utilization-based attacks using forged IP packets. Enabling this feature prevents attempts to hijack the eBGP peering session by a host on a network segment that is not part of either BGP network or by a host on a network segment that is not between the eBGP peers. | Configuring BGP Neighbor Session Options |
| 12.0(26)S | BGP MIB Support Enhancements | The BGP MIB Support Enhancements feature introduced support in the CISCO-BGP4-MIB for new SNMP notifications. | Configuring Advanced BGP Features |
| 12.0(26)S | Regex Engine Performance Enhancement | The Regex Engine Performance Enhancement feature introduces a new regular expression engine that is designed to process complex regular expressions. This new regular expression engine does not replace the existing engine. The existing engine is preferred for simple regular expressions and is the default engine and in Cisco IOS software. Either engine can be selected from the command-line interface (CLI). | Regex Engine Performance Enhancement |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|-----------|---|---|---|
| 12.0(24)S | BGP Configuration Using Peer Templates | The BGP Configuration Using Peer Templates feature introduces a new mechanism that groups distinct neighbor configurations for Border Gateway Protocol (BGP) neighbors that share policies. Configuration templates provide an alternative to peer group configuration and overcome some of the limitations of peer groups. | Configuring a Basic BGP Network |
| 12.0(24)S | BGP Cost Community | The BGP Cost Community feature introduces the cost extended community attribute. The cost community is a non-transitive extended community attribute that is passed to internal BGP (iBGP) and confederation peers but not to external BGP (eBGP) peers. The cost community feature allows you to customize the local route preference and influence the best path selection process by assigning cost values to specific routes. | BGP Cost Community |
| 12.0(24)S | BGP Dynamic Update Peer Groups | The BGP Dynamic Update Peer Groups feature introduces a new algorithm that dynamically calculates and optimizes update groups of neighbors that share the same outbound policies and can share the same update messages. In previous versions of Cisco IOS software, BGP update messages were grouped based on peer group configurations. This method of grouping updates limited outbound policies and specific-session configurations. The BGP Dynamic Update Peer Group feature separates update group replication from peer group configuration, which improves convergence time and flexibility of neighbor configuration. | Configuring a Basic BGP Network |
| 12.0(24)S | BGP Link Bandwidth | The Border Gateway Protocol (BGP) Link Bandwidth feature is used to advertise the bandwidth of an autonomous system exit link as an extended community. This feature is configured for links between directly connected external BGP (eBGP) neighbors. The link bandwidth extended community attribute is propagated to iBGP peers when extended community exchange is enabled. This feature is used with BGP multipath features to configure load balancing over links with unequal bandwidth. | BGP Link Bandwidth |
| 12.0(24)S | BGP Multipath Load Sharing for eBGP and iBGP in an MPLs VPN | The BGP Multipath Load Sharing for eBGP and iBGP feature allows you to configure multipath load balancing with both external BGP (eBGP) and internal BGP (iBGP) paths in Border Gateway Protocol (BGP) networks that are configured to use Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs). This feature provides improved load balancing deployment and service offering capabilities and is useful for multi-homed autonomous systems and Provider Edge (PE) routers that import both eBGP and iBGP paths from multihomed and stub networks. | BGP Multipath Load Sharing for eBGP and iBGP in an MPLs VPN |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|-----------|---|---|---|
| 12.0(24)S | BGP Route-Map Continue | The BGP Route-Map Continue feature introduces the continue clause to BGP route map configuration. The continue clause allows for more programmable policy configuration and route filtering and introduces the capability to execute additional entries in a route map after an entry is executed with successful match and set clauses. Continue clauses allow the network operator to configure and organize more modular policy definitions so that specific policy configurations need not be repeated within the same route map. | Connecting to a Service Provider Using External BGP |
| 12.0(22)S | BGP Conditional Route Injection | The BGP Conditional Route Injection feature allows you to inject more specific prefixes into a BGP routing table over less specific prefixes that were selected through normal route aggregation. These more specific prefixes can be used to provide a finer granularity of traffic engineering or administrative control than is possible with aggregated routes. | Configuring a Basic BGP Network |
| 12.0(22)S | BGP Hybrid CLI | The BGP Hybrid CLI feature simplifies the migration of BGP networks and existing configurations from the network layer reachability information (NLRI) format to the address family identifier (AFI) format. This new functionality allows the network operator to configure commands in the AFI format and save these command configurations to existing NLRI formatted configurations. The feature provides the network operator with the capability to take advantage of new features and provides support for migration from the NLRI format to the AFI format. | Configuring a Basic BGP Network |
| 12.0(22)S | BGP Increased Support of Numbered AS-Path Access Lists to 500 | The BGP Increased Support of Numbered AS-Path Access Lists to 500 feature increases the maximum number of autonomous systems access lists that can be configured using the ip as-path access-list command from 199 to 500. | Connecting to a Service Provider Using External BGP |
| 12.0(22)S | BGP Next Hop Propagation | The BGP Next Hop Propagation feature provides additional flexibility when designing and migrating networks. The BGP Next Hop Propagation feature allows a route reflector to modify the next hop attribute for a reflected route and allows BGP to send an update to an external BGP (eBGP) multihop peer with the next hop attribute unchanged. | BGP Next Hop Propagation |

Table 1 Supported BGP Features (continued)

| Release | Feature Name | Feature Description | Where Documented |
|-----------|---|--|---|
| 12.0(22)S | BGP Policy Accounting Output Interface Accounting | Border Gateway Protocol (BGP) policy accounting (PA) measures and classifies IP traffic that is sent to, or received from, different peers. Policy accounting was previously available on an input interface only. The BGP Policy Accounting Output Interface Accounting feature introduces several extensions to enable BGP PA on an output interface and to include accounting based on a source address for both input and output traffic on an interface. Counters based on parameters such as community list, autonomous system number, or autonomous system path are assigned to identify the IP traffic. | BGP Policy Accounting Output Interface Accounting |
| 12.0(22)S | BGP Prefix-Based Outbound Route Filtering | The BGP Prefix-Based Outbound Route Filtering feature uses BGP outbound route filtering (ORF) send and receive capabilities to minimize the number of BGP updates that are sent between BGP peers. Configuring this feature can help reduce the number of system resources required for generating and processing routing updates by filtering out unwanted routing updates at the source. For example, this feature can be used to reduce the amount of processing required on a router that is not accepting full routes from a service provider network. | Connecting to a Service Provider Using External BGP |
| 12.0(22)S | BGP Restart Session After Max-Prefix Limit | The BGP Restart Session After Max-Prefix Limit feature enhances the capabilities of the neighbor maximum-prefix command with the introduction of the restart keyword. This enhancement allows the network operator to configure the time interval at which a peering session is reestablished by a router when the number of prefixes that have been received from a peer has exceeded the maximum prefix limit. | Configuring BGP Neighbor Session Options |
| 12.0(22)S | BGP Route-Map Policy List Support | The BGP Route-Map Policy List Support feature introduces new functionality to BGP route maps. This feature adds the capability for a network operator to group route map match clauses into named lists called policy lists. A policy list functions like a macro. When a policy list is referenced in a route map, all of the match clauses are evaluated and processed as if they had been configured directly in the route map. This enhancement simplifies the configuration of BGP routing policy in medium-size and large networks because a network operator can preconfigure policy lists with groups of match clauses and then reference these policy lists within different route maps. The network operator no longer needs to manually reconfigure each recurring group of match clauses that occur in multiple route map entries. | Connecting to a Service Provider Using External BGP |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|--------------------------------|---|---|---|
| 12.0(21)S | BGP 4 MIB Support for per-Peer Received Routes | BGP 4 MIB Support for per-Peer Received Routes introduces a new table in the CISCO-BGP4-MIB that provides the capability to query (by using Simple Network Management Protocol [SNMP] commands) for routes that are learned from individual Border Gateway Protocol (BGP) peers. | BGP 4 MIB Support for per-Peer Received Routes |
| 12.0(9)S | BGP Policy Accounting | Border Gateway Protocol (BGP) policy accounting measures and classifies IP traffic that is sent to, or received from, different peers. Policy accounting is enabled on an input interface, and counters based on parameters such as community list, autonomous system number, or autonomous system path are assigned to identify the IP traffic. | BGP Policy Accounting |
| Cisco IOS Release 12.2S | | | |
| 12.2(25)S | BGP Cost Community Support for EIGRP MPLS VPN PE-CE | The BGP Cost Community Support for EIGRP MPLS VPN PE-CE feature provides BGP cost community support for mixed EIGRP MPLS VPN network topologies that contain back door routes. | BGP Cost Community |
| 12.2(25)S | BGP MIB Support Enhancements | The BGP MIB Support Enhancements feature introduced support in the CISCO-BGP4-MIB for new SNMP notifications. | Configuring Advanced BGP Features |
| 12.2(25)S | BGP Support for Dual AS Configuration for Network AS Migrations | The BGP Support for Dual AS Configuration for Network AS Migrations feature extends the functionality of the BGP Local-AS feature by providing additional autonomous-system path customization configuration options. The configuration of this feature is transparent to customer peering sessions, allowing the provider to merge two autonomous-systems without interrupting customer peering arrangements. Customer peering sessions can later be updated during a maintenance window or during other scheduled downtime. | Configuring BGP Neighbor Session Options |
| 12.2(25)S | BGP Support for IP Prefix Import from Global Table into a VRF Table | The BGP Support for IP Prefix Import from Global Table into a VRF Table feature introduces the capability to import IPv4 unicast prefixes from the global routing table into a Virtual Private Network (VPN) routing/forwarding (VRF) instance table using an import route map. | BGP Support for IP Prefix Import from Global Table into a VRF Table |
| 12.2(25)S | BGP Support for Named Extended Community Lists | The BGP Support for Named Extended Community Lists feature introduces the ability to configure extended community lists using names in addition to the existing numbered format. | Connecting to a Service Provider Using External BGP |
| 12.2(25)S | BGP Support for Sequenced Entries in Extended Community Lists | The BGP Support for Sequenced Entries in Extended Community Lists feature introduces automatic sequencing of individual entries in BGP extended community lists. This feature also introduces the ability to remove or resequence extended community list entries without deleting the entire existing extended community list. | Connecting to a Service Provider Using External BGP |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|----------------|--|---|--|
| 12.2(25)S | BGP Support for TTL Security Check | The BGP Support for TTL Security Check feature introduced a lightweight security mechanism to protect external Border Gateway Protocol (eBGP) peering sessions from CPU utilization-based attacks using forged IP packets. Enabling this feature prevents attempts to hijack the eBGP peering session by a host on a network segment that is not part of either BGP network or by a host on a network segment that is not between the eBGP peers. | Configuring BGP Neighbor Session Options |
| 12.2(25)S | Loadsharing IP Packets Over More Than Six Parallel Paths | The Loadsharing IP Packets Over More Than Six Parallel Paths feature increases the maximum number of parallel routes that can be installed to the routing table for multipath loadsharing. | Loadsharing IP Packets Over More Than Six Parallel Paths |
| 12.2(25)S | Suppress BGP Advertisement for Inactive Routes | The Suppress BGP Advertisements for Inactive Routes feature allows you to configure the suppression of advertisements for routes that are not installed in the Routing Information Base (RIB). Configuring this feature allows Border Gateway Protocol (BGP) updates to be more consistent with data used for traffic forwarding. | Configuring a Basic BGP Network |
| 12.2(22)S | BGP Policy Accounting Output Interface Accounting | Border Gateway Protocol (BGP) policy accounting (PA) measures and classifies IP traffic that is sent to, or received from, different peers. Policy accounting was previously available on an input interface only. The BGP Policy Accounting Output Interface Accounting feature introduces several extensions to enable BGP PA on an output interface and to include accounting based on a source address for both input and output traffic on an interface. Counters based on parameters such as community list, autonomous system number, or autonomous system path are assigned to identify the IP traffic. | BGP Policy Accounting Output Interface Accounting |
| 12.2(22)S | Regex Engine Performance Enhancement | The Regex Engine Performance Enhancement feature introduces a new regular expression engine that is designed to process complex regular expressions. This new regular expression engine does not replace the existing engine. The existing engine is preferred for simple regular expressions and is the default engine and in Cisco IOS software. Either engine can be selected from the command-line interface (CLI). | Regex Engine Performance Enhancement |
| 12.2(18)S | BGP Configuration Using Peer Templates | The BGP Configuration Using Peer Templates feature introduces a new mechanism that groups distinct neighbor configurations for BGP neighbors that share policies. Configuration templates provide an alternative to peer group configuration and overcome some of the limitations of peer groups. | Configuring a Basic BGP Network |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|-----------|---|--|---|
| 12.2(18)S | BGP Cost Community | The BGP Cost Community feature introduces the cost extended community attribute. The cost community is a non-transitive extended community attribute that is passed to internal BGP (iBGP) and confederation peers but not to external BGP (eBGP) peers. The cost community feature allows you to customize the local route preference and influence the best path selection process by assigning cost values to specific routes. | BGP Cost Community |
| 12.2(18)S | BGP Dynamic Update Peer Groups | The BGP Dynamic Update Peer Groups feature introduces a new algorithm that dynamically calculates and optimizes update groups of neighbors that share outbound policies and can share update messages. In previous versions of Cisco IOS software, BGP update messages were grouped based on peer-group configurations. This method of grouping updates limited outbound policies and specific-session configurations. The BGP Dynamic Update Peer Groups feature separates update group replication from peer group configuration, which improves convergence time and flexibility of neighbor configuration. | Configuring a Basic BGP Network |
| 12.2(18)S | BGP Increased Support of Numbered AS-Path Access Lists to 500 | The BGP Increased Support of Numbered AS-Path Access Lists to 500 feature increases the maximum number of autonomous systems access lists that can be configured using the ip as-path access-list command from 199 to 500. | Connecting to a Service Provider Using External BGP |
| 12.2(18)S | BGP Restart Session After Max-Prefix Limit | The BGP Restart Session After Max-Prefix Limit feature enhances the capabilities of the neighbor maximum-prefix command with the introduction of the restart keyword. This enhancement allows the network operator to configure the time interval at which a peering session is reestablished by a router when the number of prefixes that have been received from a peer has exceeded the maximum prefix limit. | Configuring BGP Neighbor Session Options |
| 12.2(18)S | BGP Route-Map Continue | The BGP Route-Map Continue feature introduces the continue clause to BGP route map configuration. The continue clause allows for more programmable policy configuration and route filtering and introduces the capability to execute additional entries in a route map after an entry is executed with successful match and set clauses. Continue clauses allow the network operator to configure and organize more modular policy definitions so that specific policy configurations need not be repeated within the same route map. Continue clauses are supported in outbound route maps only in Cisco IOS Release 12.0(31)S and later releases. | Connecting to a Service Provider Using External BGP |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|----------------|--|--|---|
| 12.2(18)S | BGP Route-Map Policy List Support | The BGP Route-Map Policy List Support feature introduces new functionality to BGP route maps. This feature adds the capability for a network operator to group route map match clauses into named lists called policy lists. A policy list functions like a macro. When a policy list is referenced in a route map, all of the match clauses are evaluated and processed as if they had been configured directly in the route map. This enhancement simplifies the configuration of BGP routing policy in medium-size and large networks because a network operator can preconfigure policy lists with groups of match clauses and then reference these policy lists within different route maps. The network operator no longer needs to manually reconfigure each recurring group of match clauses that occur in multiple route map entries. | Connecting to a Service Provider Using External BGP |
| 12.2(14)S | BGP 4 MIB Support for per-Peer Received Routes | BGP 4 MIB Support for per-Peer Received Routes introduces a new table in the CISCO-BGP4-MIB that provides the capability to query (by using Simple Network Management Protocol [SNMP] commands) for routes that are learned from individual Border Gateway Protocol (BGP) peers. | BGP 4 MIB Support for per-Peer Received Routes |
| 12.2(14)S | BGP Conditional Route Injection | The BGP Conditional Route Injection feature allows you to inject more specific prefixes into a BGP routing table over less specific prefixes that were selected through normal route aggregation. These more specific prefixes can be used to provide a finer granularity of traffic engineering or administrative control than is possible with aggregated routes. | Configuring a Basic BGP Network |
| 12.2(14)S | BGP Link Bandwidth | The Border Gateway Protocol (BGP) Link Bandwidth feature is used to advertise the bandwidth of an autonomous system exit link as an extended community. This feature is configured for links between directly connected external BGP (eBGP) neighbors. The link bandwidth extended community attribute is propagated to iBGP peers when extended community exchange is enabled. This feature is used with BGP multipath features to configure load balancing over links with unequal bandwidth. | BGP Link Bandwidth |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|-----------|---|--|---|
| 12.2(14)S | BGP Multipath Load Sharing for eBGP and iBGP in an MPLs VPN | The BGP Multipath Load Sharing for eBGP and iBGP feature allows you to configure multipath load balancing with both external BGP (eBGP) and internal BGP (iBGP) paths in Border Gateway Protocol (BGP) networks that are configured to use Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs). This feature provides improved load balancing deployment and service offering capabilities and is useful for multi-homed autonomous systems and Provider Edge (PE) routers that import both eBGP and iBGP paths from multihomed and stub networks. | BGP Multipath Load Sharing for eBGP and iBGP in an MPLs VPN |
| 12.2(14)S | BGP Named Community Lists | The BGP Named Community Lists feature introduces a new type of community list called the named community list. The BGP Named Community Lists feature allows the network operator to assign meaningful names to community lists and increases the number of community lists that can be configured. A named community list can be configured with regular expressions and with numbered community lists. All rules of numbered communities apply to named community lists except that there is no limitation on the number of community attributes that can be configured for a named community list. | Connecting to a Service Provider Using External BGP |
| 12.2(14)S | BGP Next Hop Propagation | The BGP Next Hop Propagation feature provides additional flexibility when designing and migrating networks. The BGP Next Hop Propagation feature allows a route reflector to modify the next hop attribute for a reflected route and allows BGP to send an update to an external BGP (eBGP) multihop peer with the next hop attribute unchanged. | BGP Next Hop Propagation |
| 12.2(14)S | BGP Prefix-Based Outbound Route Filtering | The BGP Prefix-Based Outbound Route Filtering feature uses BGP ORF send and receive capabilities to minimize the number of BGP updates that are sent between BGP peers. Configuring this feature can help reduce the number of system resources required for generating and processing routing updates by filtering out unwanted routing updates at the source. For example, this feature can be used to reduce the amount of processing required on a router that is not accepting full routes from a service provider network. | Connecting to a Service Provider Using External BGP |
| 12.2(14)S | iBGP Multipath Load Sharing | The iBGP Multipath Load Sharing feature enables the BGP speaking router to select multiple iBGP paths as the best paths to a destination. The best paths or multipaths are then installed in the IP routing table of the router. | iBGP Multipath Load Sharing |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|---------------------------------|---|--|---|
| Cisco IOS Release 12.2SB | | | |
| 12.2(33)SB | BGP Graceful Restart per Neighbor | The BGP Graceful Restart per Neighbor feature enables or disables the BGP graceful restart capability for an individual BGP neighbor, including using peer session templates and BGP peer groups. | Configuring Advanced BGP Features |
| 12.2(33)SB | BGP Neighbor Policy | The BGP Neighbor Policy feature introduces new keywords to two existing commands to display information about local and inherited policies. When BGP neighbors use multiple levels of peer templates, it can be difficult to determine which policies are applied to the neighbor. Inherited policies are policies that the neighbor inherits from a peer-group or a peer-policy template. | Configuring a Basic BGP Network |
| 12.2(33)SB | BGP Route-Map Continue Support for an Outbound Policy | The BGP Route-Map Continue Support for an Outbound Policy feature introduces support for continue clauses to be applied to outbound route maps. | Connecting to a Service Provider Using External BGP |
| 12.2(33)SB | BGP Support for BFD | Bidirectional Forwarding Detection (BFD) is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. In addition to fast forwarding path failure detection, BFD provides a consistent failure detection method for network administrators. Because the network administrator can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different routing protocol hello mechanisms, network profiling and planning will be easier, and reconvergence time will be consistent and predictable. The main benefit of implementing BFD for BGP is a significantly faster reconvergence time. | Configuring Advanced BGP Features |
| 12.2(31)SB2 | Per-VRF Assignment of BGP Router ID | The Per-VRF Assignment of BGP Router ID feature introduces the ability to have VRF-to-VRF peering in Border Gateway Protocol (BGP) on the same router. BGP is designed to refuse a session with itself because of the router ID check. The per-VRF assignment feature allows a separate router ID per VRF using a new keyword in the existing bgp router-id command. The router ID can be manually configured for each VRF or can be assigned automatically either globally under address family configuration mode or for each VRF. | Per-VRF Assignment of BGP Router ID |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|----------------|--|---|---|
| 12.2(31)SB | BGP Support for Fast Peering Session Deactivation | The BGP Support for Fast Peering Session Deactivation feature introduced an event driven notification system that allows a Border Gateway Protocol (BGP) process to monitor BGP peering sessions on a per-neighbor basis. This feature improves the response time of BGP to adjacency changes by allowing BGP to detect an adjacency change and deactivate the terminated session in between standard BGP scanning intervals. Enabling this feature improves overall BGP convergence. | Configuring BGP Neighbor Session Options |
| 12.2(31)SB | BGP Selective Address Tracking | The BGP Selective Address Tracking feature introduced the use of a route map for next-hop route filtering and fast session deactivation. Selective next-hop filtering uses a route map to selectively define routes to help resolve the BGP next hop, or a route map can be used to determine if a peering session with a BGP neighbor should be reset when a route to the BGP peer changes. | Configuring Advanced BGP Features Configuring BGP Neighbor Session Options |
| 12.2(31)SB | BGP Support for TCP Path MTU Discovery per Session | Border Gateway Protocol (BGP) support for Transmission Control Protocol (TCP) path maximum transmission unit (MTU) discovery introduced the ability for BGP to automatically discover the best TCP path MTU for each BGP session. The TCP path MTU is enabled by default for all BGP neighbor sessions, but you can disable, and subsequently enable, the TCP path MTU globally for all BGP sessions or for an individual BGP neighbor session. | Configuring BGP Neighbor Session Options |
| 12.2(28)SB | BGP 4 MIB Support for per-Peer Received Routes | BGP 4 MIB Support for per-Peer Received Routes introduces a new table in the CISCO-BGP4-MIB that provides the capability to query (by using Simple Network Management Protocol [SNMP] commands) for routes that are learned from individual Border Gateway Protocol (BGP) peers. | BGP 4 MIB Support for per-Peer Received Routes |
| 12.2(28)SB | BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO) | The BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO) feature enables provider edge (PE) routers to maintain Border Gateway Protocol (BGP) state with customer edge (CE) routers and ensure continuous packet forwarding during a Route Processor (RP) switchover or during a planned In-Service Software Upgrade (ISSU) for a PE router. CE routers do not need to be Nonstop Forwarding (NSF)-capable or NSF-aware to benefit from BGP NSR capabilities on PE routers. Only PE routers need to be upgraded to support BGP NSR—no CE router upgrades are required. BGP NSR with SSO, thus, enables service providers to provide the benefits NSF with the additional benefits of NSR without requiring CE routers to be upgraded to support BGP graceful restart. | BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO) |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|-------------|---|--|---|
| 12.2(27)SBC | BGP Configuration Using Peer Templates | The BGP Configuration Using Peer Templates feature introduces a new mechanism that groups distinct neighbor configurations for BGP neighbors that share policies. Configuration templates provide an alternative to peer group configuration and overcome some of the limitations of peer groups. | Configuring a Basic BGP Network |
| 12.2(27)SBC | BGP Cost Community | The BGP Cost Community feature introduces the cost extended community attribute. The cost community is a non-transitive extended community attribute that is passed to internal BGP (iBGP) and confederation peers but not to external BGP (eBGP) peers. The cost community feature allows you to customize the local route preference and influence the best path selection process by assigning cost values to specific routes. | BGP Cost Community |
| 12.2(27)SBC | BGP Dynamic Update Peer Groups | The BGP Dynamic Update Peer Groups feature introduces a new algorithm that dynamically calculates and optimizes update groups of neighbors that share outbound policies and can share update messages. In previous versions of Cisco IOS software, BGP update messages were grouped based on peer group configurations. This method of grouping updates limited outbound policies and specific-session configurations. The BGP Dynamic Update Peer Groups feature separates update group replication from peer group configuration, which improves convergence time and flexibility of neighbor configuration. | Configuring a Basic BGP Network |
| 12.2(27)SBC | BGP Increased Support of Numbered AS-Path Access Lists to 500 | The BGP Increased Support of Numbered AS-Path Access Lists to 500 feature increases the maximum number of autonomous systems access lists that can be configured using the ip as-path access-list command from 199 to 500. | Connecting to a Service Provider Using External BGP |
| 12.2(27)SBC | BGP Route-Map Continue | The BGP Route-Map Continue feature introduces the continue clause to BGP route map configuration. The continue clause allows for more programmable policy configuration and route filtering and introduces the capability to execute additional entries in a route map after an entry is executed with successful match and set clauses. Continue clauses allow the network operator to configure and organize more modular policy definitions so that specific policy configurations need not be repeated within the same route map. Continue clauses are supported in outbound route maps only in Cisco IOS Release 12.0(31)S and later releases. | Connecting to a Service Provider Using External BGP |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|---------------------------------|---|--|---|
| 12.2(27)SBC | BGP Route-Map Policy List Support | The BGP Route-Map Policy List Support feature introduces new functionality to BGP route maps. This feature adds the capability for a network operator to group route map match clauses into named lists called policy lists. A policy list functions like a macro. When a policy list is referenced in a route map, all of the match clauses are evaluated and processed as if they had been configured directly in the route map. This enhancement simplifies the configuration of BGP routing policy in medium-size and large networks because a network operator can preconfigure policy lists with groups of match clauses and then reference these policy lists within different route maps. The network operator no longer needs to manually reconfigure each recurring group of match clauses that occur in multiple route map entries. | Connecting to a Service Provider Using External BGP |
| 12.2(27)SBC | BGP Support for IP Prefix Import from Global Table into a VRF Table | The BGP Support for IP Prefix Import from Global Table into a VRF Table feature introduces the capability to import IPv4 unicast prefixes from the global routing table into a Virtual Private Network (VPN) routing/forwarding (VRF) instance table using an import route map. | BGP Support for IP Prefix Import from Global Table into a VRF Table |
| 12.2(27)SBC | BGP Support for Named Extended Community Lists | The BGP Support for Named Extended Community Lists feature introduces the ability to configure extended community lists using names in addition to the existing numbered format. | Connecting to a Service Provider Using External BGP |
| 12.2(27)SBC | BGP Support for Sequenced Entries in Extended Community Lists | The BGP Support for Sequenced Entries in Extended Community Lists feature introduces automatic sequencing of individual entries in BGP extended community lists. This feature also introduces the ability to remove or resequence extended community list entries without deleting the entire existing extended community list. | Connecting to a Service Provider Using External BGP |
| Cisco IOS Release 12.2SR | | | |
| 12.2(33)SRC | BGP Graceful Restart per Neighbor | The BGP Graceful Restart per Neighbor feature enables or disables the BGP graceful restart capability for an individual BGP neighbor, including using peer session templates and BGP peer groups. | Configuring Advanced BGP Features |
| 12.2(33)SRC | BGP MIB Support Enhancements | The BGP MIB Support Enhancements feature introduced support in the CISCO-BGP4-MIB for new SNMP notifications. | Configuring Advanced BGP Features |
| 12.2(33)SRB | BGP Neighbor Policy | The BGP Neighbor Policy feature introduces new keywords to two existing commands to display information about local and inherited policies. When BGP neighbors use multiple levels of peer templates, it can be difficult to determine which policies are applied to the neighbor. Inherited policies are policies that the neighbor inherits from a peer-group or a peer-policy template. | Configuring a Basic BGP Network |

Table 1 Supported BGP Features (continued)

| Release | Feature Name | Feature Description | Where Documented |
|-------------|---|--|---|
| 12.2(33)SRB | BGP per Neighbor SoO Configuration | The BGP per neighbor SOO configuration feature simplifies the configuration of the site-of-origin (SoO) parameter. In previous releases, the SoO parameter is configured using an inbound route map that sets the SoO value during the update process. The per neighbor SoO configuration introduces two new commands that can be used under router configuration mode to set the SoO value. | BGP per Neighbor SoO Configuration |
| 12.2(33)SRB | BGP Route-Map Continue Support for an Outbound Policy | The BGP Route-Map Continue Support for an Outbound Policy feature introduces support for continue clauses to be applied to outbound route maps. | Connecting to a Service Provider Using External BGP |
| 12.2(33)SRB | BGP Selective Address Tracking | The BGP Selective Address Tracking feature introduces the use of a route map for next-hop route filtering and fast session deactivation. Selective next-hop filtering uses a route map to selectively define routes to help resolve the BGP next hop, or a route map can be used to determine if a peering session with a BGP neighbor should be reset when a route to the BGP peer changes. | Configuring Advanced BGP Features Configuring BGP Neighbor Session Options |
| 12.2(33)SRB | BGP Support for MTR | BGP support for MTR introduces a new configuration hierarchy and command-line interface (CLI) commands to support multi-topology routing (MTR) topologies. The new configuration hierarchy, or scope, can be implemented by BGP independently of MTR. MTR allows the configuration of service differentiation through class-based forwarding. MTR supports multiple unicast topologies and a separate multicast topology. A topology is a subset of the underlying network (or base topology) characterized by an independent set of Network Layer Reachability Information (NLRI). In 12.2(33)SRB, this feature was introduced on the Cisco 7600 series routers. | Configuring Advanced BGP Features |
| 12.2(33)SRB | BGP Support for the L2VPN Address Family | BGP support for the L2VPN address family introduces a BGP-based autodiscovery mechanism to distribute Layer 2 Virtual Private Network (L2VPN) endpoint provisioning information. BGP uses a separate L2VPN routing information base (RIB) to store endpoint provisioning information, which is updated each time any Layer 2 virtual forwarding instance (VFI) is configured. When BGP distributes the endpoint provisioning information in an update message to all its BGP neighbors, the endpoint information is used to set up a pseudowire mesh to support L2VPN-based services. | BGP Support for the L2VPN Address Family |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|----------------|---|--|---|
| 12.2(33)SRB | Multiprotocol BGP (MP-BGP) Support for the CLNS | The Multiprotocol BGP (MP-BGP) Support for CLNS feature provides the ability to scale Connectionless Network Service (CLNS) networks. The multiprotocol extensions of Border Gateway Protocol (BGP) add the ability to interconnect separate Open System Interconnection (OSI) routing domains without merging the routing domains, thus providing the capability to build very large OSI networks. | Multiprotocol BGP (MP-BGP) Support for the CLNS |
| 12.2(33)SRA | BGP MIB Support Enhancements | The BGP MIB Support Enhancements feature introduced support in the CISCO-BGP4-MIB for new SNMP notifications. | Configuring Advanced BGP Features |
| 12.2(33)SRA | BGP Support for BFD | Bidirectional Forwarding Detection (BFD) is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. In addition to fast forwarding path failure detection, BFD provides a consistent failure detection method for network administrators. Because the network administrator can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different routing protocol hello mechanisms, network profiling and planning will be easier, and reconvergence time will be consistent and predictable. The main benefit of implementing BFD for BGP is a significantly faster reconvergence time. | Configuring Advanced BGP Features |
| 12.2(33)SRA | BGP Support for Dual AS Configuration for Network AS Migrations | The BGP Support for Dual AS Configuration for Network AS Migrations feature extends the functionality of the BGP Local-AS feature by providing additional autonomous-system path customization configuration options. The configuration of this feature is transparent to customer peering sessions, allowing the provider to merge two autonomous-systems without interrupting customer peering arrangements. Customer peering sessions can later be updated during a maintenance window or during other scheduled downtime. | Configuring BGP Neighbor Session Options |
| 12.2(33)SRA | BGP Support for Fast Peering Session Deactivation | The BGP Support for Fast Peering Session Deactivation feature introduced an event driven notification system that allows a Border Gateway Protocol (BGP) process to monitor BGP peering sessions on a per-neighbor basis. This feature improves the response time of BGP to adjacency changes by allowing BGP to detect an adjacency change and deactivate the terminated session in between standard BGP scanning intervals. Enabling this feature improves overall BGP convergence. | Configuring BGP Neighbor Session Options |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|---------------------------------|---|--|---|
| 12.2(33)SRA | BGP Support for IP Prefix Import from Global Table into a VRF Table | The BGP Support for IP Prefix Import from Global Table into a VRF Table feature introduces the capability to import IPv4 unicast prefixes from the global routing table into a Virtual Private Network (VPN) routing/forwarding (VRF) instance table using an import route map. | BGP Support for IP Prefix Import from Global Table into a VRF Table |
| 12.2(33)SRA | BGP Support for Named Extended Community Lists | The BGP Support for Named Extended Community Lists feature introduces the ability to configure extended community lists using names in addition to the existing numbered format. | Connecting to a Service Provider Using External BGP |
| 12.2(33)SRA | BGP Support for Sequenced Entries in Extended Community Lists | The BGP Support for Sequenced Entries in Extended Community Lists feature introduces automatic sequencing of individual entries in BGP extended community lists. This feature also introduces the ability to remove or resequence extended community list entries without deleting the entire existing extended community list. | Connecting to a Service Provider Using External BGP |
| 12.2(33)SRA | BGP Support for TCP Path MTU Discovery per Session | Border Gateway Protocol (BGP) support for Transmission Control Protocol (TCP) path maximum transmission unit (MTU) discovery introduced the ability for BGP to automatically discover the best TCP path MTU for each BGP session. The TCP path MTU is enabled by default for all BGP neighbor sessions, but you can disable, and subsequently enable, the TCP path MTU globally for all BGP sessions or for an individual BGP neighbor session. | Configuring BGP Neighbor Session Options |
| 12.2(33)SRA | Per-VRF Assignment of BGP Router ID | The Per-VRF Assignment of BGP Router ID feature introduces the ability to have VRF-to-VRF peering in Border Gateway Protocol (BGP) on the same router. BGP is designed to refuse a session with itself because of the router ID check. The per-VRF assignment feature allows a separate router ID per VRF using a new keyword in the existing bgp router-id command. The router ID can be manually configured for each VRF or can be assigned automatically either globally under address family configuration mode or for each VRF. | Per-VRF Assignment of BGP Router ID |
| Cisco IOS Release 12.2SX | | | |
| 12.2(33)SXII | BGP Support for 4-Byte ASN | The BGP Support for 4-Byte ASN feature introduced support for 4-byte autonomous system numbers. Because of increased demand for autonomous system numbers, in January 2009 the IANA will start to allocate 4-byte autonomous system numbers in the range from 65536 to 4294967295. | Cisco BGP Overview Configuring a Basic BGP Network |
| 12.2(33)SXI | BGP Route-Map Continue Support for an Outbound Policy | The BGP Route-Map Continue Support for an Outbound Policy feature introduces support for continue clauses to be applied to outbound route maps. | Connecting to a Service Provider Using External BGP |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|-------------|---|--|--|
| 12.2(33)SXH | BGP Dynamic Neighbors | BGP dynamic neighbor support allows BGP peering to a group of remote neighbors that are defined by a range of IP addresses. Each range can be configured as a subnet IP address. BGP dynamic neighbors are configured using a range of IP addresses and BGP peer groups. After a subnet range is configured for a BGP peer group and a TCP session is initiated for an IP address in the subnet range, a new BGP neighbor is dynamically created as a member of that group. The new BGP neighbor will inherit any configuration for the peer group. The output for three show commands has been updated to display information about dynamic neighbors. | Configuring BGP Neighbor Session Options |
| 12.2(33)SXH | BGP MIB Support Enhancements | The BGP MIB Support Enhancements feature introduced support in the CISCO-BGP4-MIB for new SNMP notifications. | Configuring Advanced BGP Features |
| 12.2(33)SXH | BGP Support for BFD | Bidirectional Forwarding Detection (BFD) is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. In addition to fast forwarding path failure detection, BFD provides a consistent failure detection method for network administrators. Because the network administrator can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different routing protocol hello mechanisms, network profiling and planning will be easier, and reconvergence time will be consistent and predictable. The main benefit of implementing BFD for BGP is a significantly faster reconvergence time. | Configuring Advanced BGP Features |
| 12.2(33)SXH | BGP Support for Dual AS Configuration for Network AS Migrations | The BGP Support for Dual AS Configuration for Network AS Migrations feature extends the functionality of the BGP Local-AS feature by providing additional autonomous-system path customization configuration options. The configuration of this feature is transparent to customer peering sessions, allowing the provider to merge two autonomous-systems without interrupting customer peering arrangements. Customer peering sessions can later be updated during a maintenance window or during other scheduled downtime. | Configuring BGP Neighbor Session Options |
| 12.2(33)SXH | BGP Support for Fast Peering Session Deactivation | The BGP Support for Fast Peering Session Deactivation feature introduced an event driven notification system that allows a Border Gateway Protocol (BGP) process to monitor BGP peering sessions on a per-neighbor basis. This feature improves the response time of BGP to adjacency changes by allowing BGP to detect an adjacency change and deactivate the terminated session in between standard BGP scanning intervals. Enabling this feature improves overall BGP convergence. | Configuring BGP Neighbor Session Options |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|-------------|---|---|---|
| 12.2(33)SXH | BGP Support for IP Prefix Import from Global Table into a VRF Table | The BGP Support for IP Prefix Import from Global Table into a VRF Table feature introduces the capability to import IPv4 unicast prefixes from the global routing table into a Virtual Private Network (VPN) routing/forwarding (VRF) instance table using an import route map. | BGP Support for IP Prefix Import from Global Table into a VRF Table |
| 12.2(33)SXH | BGP Support for Named Extended Community Lists | The BGP Support for Named Extended Community Lists feature introduces the ability to configure extended community lists using names in addition to the existing numbered format. | Connecting to a Service Provider Using External BGP |
| 12.2(33)SXH | BGP Support for Next-Hop Address Tracking | The BGP Support for Next-Hop Address Tracking feature is enabled by default when a supporting Cisco IOS software image is installed. BGP next-hop address tracking is event driven. BGP prefixes are automatically tracked as peering sessions are established. Next-hop changes are rapidly reported to the BGP routing process as they are updated in the RIB. This optimization improves overall BGP convergence by reducing the response time to next-hop changes for routes installed in the RIB. When a bestpath calculation is run in between BGP scanner cycles, only next-hop changes are tracked and processed. | Configuring Advanced BGP Features |
| 12.2(33)SXH | BGP Support for Sequenced Entries in Extended Community Lists | The BGP Support for Sequenced Entries in Extended Community Lists feature introduces automatic sequencing of individual entries in BGP extended community lists. This feature also introduces the ability to remove or resequence extended community list entries without deleting the entire existing extended community list. | Connecting to a Service Provider Using External BGP |
| 12.2(33)SXH | BGP Support for TCP Path MTU Discovery per Session | Border Gateway Protocol (BGP) support for Transmission Control Protocol (TCP) path maximum transmission unit (MTU) discovery introduced the ability for BGP to automatically discover the best TCP path MTU for each BGP session. The TCP path MTU is enabled by default for all BGP neighbor sessions, but you can disable, and subsequently enable, the TCP path MTU globally for all BGP sessions or for an individual BGP neighbor session. | Configuring BGP Neighbor Session Options |
| 12.2(33)SXH | Per-VRF Assignment of BGP Router ID | The Per-VRF Assignment of BGP Router ID feature introduces the ability to have VRF-to-VRF peering in Border Gateway Protocol (BGP) on the same router. BGP is designed to refuse a session with itself because of the router ID check. The per-VRF assignment feature allows a separate router ID per VRF using a new keyword in the existing bgp router-id command. The router ID can be manually configured for each VRF or can be assigned automatically either globally under address family configuration mode or for each VRF. | Per-VRF Assignment of BGP Router ID |

Table 1 Supported BGP Features (continued)

| Release | Feature Name | Feature Description | Where Documented |
|--|---|---|---|
| 12.2(33)SXH | Suppress BGP Advertisement for Inactive Routes | The Suppress BGP Advertisements for Inactive Routes feature allows you to configure the suppression of advertisements for routes that are not installed in the Routing Information Base (RIB). Configuring this feature allows Border Gateway Protocol (BGP) updates to be more consistent with data used for traffic forwarding. | Configuring a Basic BGP Network |
| 12.2(18)SXE | BGP Multipath Load Sharing for eBGP and iBGP in an MPLs VPN | The BGP Multipath Load Sharing for eBGP and iBGP feature allows you to configure multipath load balancing with both external BGP (eBGP) and internal BGP (iBGP) paths in Border Gateway Protocol (BGP) networks that are configured to use Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs). This feature provides improved load balancing deployment and service offering capabilities and is useful for multi-homed autonomous systems and Provider Edge (PE) routers that import both eBGP and iBGP paths from multihomed and stub networks. | BGP Multipath Load Sharing for eBGP and iBGP in an MPLs VPN |
| 12.2(18)SXE | BGP Support for TTL Security Check | The BGP Support for TTL Security Check feature introduced a lightweight security mechanism to protect external Border Gateway Protocol (eBGP) peering sessions from CPU utilization-based attacks using forged IP packets. Enabling this feature prevents attempts to hijack the eBGP peering session by a host on a network segment that is not part of either BGP network or by a host on a network segment that is not between the eBGP peers. | Configuring BGP Neighbor Session Options |
| Cisco IOS Releases 12.2T, 12.3, 12.3T, 12.4 and 12.4T | | | |
| 12.4(24)T | BGP Support for 4-Byte ASN | The BGP Support for 4-Byte ASN feature introduced support for 4-byte autonomous system numbers. Because of increased demand for autonomous system numbers, in January 2009 the IANA will start to allocate 4-byte autonomous system numbers in the range from 65536 to 4294967295. | Cisco BGP Overview Configuring a Basic BGP Network |
| 12.4(20)T | BGP Support for TCP Path MTU Discovery per Session | Border Gateway Protocol (BGP) support for Transmission Control Protocol (TCP) path maximum transmission unit (MTU) discovery introduced the ability for BGP to automatically discover the best TCP path MTU for each BGP session. The TCP path MTU is enabled by default for all BGP neighbor sessions, but you can disable, and subsequently enable, the TCP path MTU globally for all BGP sessions or for an individual BGP neighbor session. | Configuring BGP Neighbor Session Options |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|----------------|---|---|---|
| 12.4(20)T | Per-VRF Assignment of BGP Router ID | The Per-VRF Assignment of BGP Router ID feature introduces the ability to have VRF-to-VRF peering in Border Gateway Protocol (BGP) on the same router. BGP is designed to refuse a session with itself because of the router ID check. The per-VRF assignment feature allows a separate router ID per VRF using a new keyword in the existing <code>bgp router-id</code> command. The router ID can be manually configured for each VRF or can be assigned automatically either globally under address family configuration mode or for each VRF. | Per-VRF Assignment of BGP Router ID |
| 12.4(11)T | BGP Neighbor Policy | The BGP Neighbor Policy feature introduces new keywords to two existing commands to display information about local and inherited policies. When BGP neighbors use multiple levels of peer templates, it can be difficult to determine which policies are applied to the neighbor. Inherited policies are policies that the neighbor inherits from a peer-group or a peer-policy template. | Configuring a Basic BGP Network |
| 12.4(11)T | BGP per Neighbor SoO Configuration | The BGP per neighbor SOO configuration feature simplifies the configuration of the site-of-origin (SoO) parameter. In Cisco IOS Release 12.4(9)T, 12.2(33)SRA, and previous releases, the SoO parameter is configured using an inbound route map that sets the SoO value during the update process. The per neighbor SoO configuration introduces two new commands that can be used under router configuration mode to set the SoO value. | BGP per Neighbor SoO Configuration |
| 12.4(4)T | BGP Route-Map Continue Support for an Outbound Policy | The BGP Route-Map Continue Support for an Outbound Policy feature introduces support for continue clauses to be applied to outbound route maps. | Connecting to a Service Provider Using External BGP |
| 12.4(4)T | BGP Selective Address Tracking | The BGP Selective Address Tracking feature introduces the use of a route map for next-hop route filtering and fast session deactivation. Selective next-hop filtering uses a route map to selectively define routes to help resolve the BGP next hop, or a route map can be used to determine if a peering session with a BGP neighbor should be reset when a route to the BGP peer changes. | Configuring Advanced BGP Features Configuring BGP Neighbor Session Options |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|----------------|---|--|---|
| 12.4(4)T | BGP Support for BFD | Bidirectional Forwarding Detection (BFD) is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. In addition to fast forwarding path failure detection, BFD provides a consistent failure detection method for network administrators. Because the network administrator can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different routing protocol hello mechanisms, network profiling and planning will be easier, and reconvergence time will be consistent and predictable. The main benefit of implementing BFD for BGP is a significantly faster reconvergence time. | Configuring Advanced BGP Features |
| 12.3(14)T | BGP Support for Fast Peering Session Deactivation | The BGP Support for Fast Peering Session Deactivation feature introduced an event driven notification system that allows a Border Gateway Protocol (BGP) process to monitor BGP peering sessions on a per-neighbor basis. This feature improves the response time of BGP to adjacency changes by allowing BGP to detect an adjacency change and deactivate the terminated session in between standard BGP scanning intervals. Enabling this feature improves overall BGP convergence. | Configuring BGP Neighbor Session Options |
| 12.3(14)T | BGP Support for IP Prefix Import from Global Table into a VRF Table | The BGP Support for IP Prefix Import from Global Table into a VRF Table feature introduces the capability to import IPv4 unicast prefixes from the global routing table into a Virtual Private Network (VPN) routing/forwarding (VRF) instance table using an import route map. | BGP Support for IP Prefix Import from Global Table into a VRF Table |
| 12.3(14)T | BGP Support for Next-Hop Address Tracking | The BGP Support for Next-Hop Address Tracking feature is enabled by default when a supporting Cisco IOS software image is installed. BGP next-hop address tracking is event driven. BGP prefixes are automatically tracked as peering sessions are established. Next-hop changes are rapidly reported to the BGP routing process as they are updated in the RIB. This optimization improves overall BGP convergence by reducing the response time to next-hop changes for routes installed in the RIB. When a bestpath calculation is run in between BGP scanner cycles, only next-hop changes are tracked and processed. | Configuring Advanced BGP Features |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|----------------|---|---|---|
| 12.3(11)T | BGP Support for Dual AS Configuration for Network AS Migrations | The BGP Support for Dual AS Configuration for Network AS Migrations feature extends the functionality of the BGP Local-AS feature by providing additional autonomous-system path customization configuration options. The configuration of this feature is transparent to customer peering sessions, allowing the provider to merge two autonomous-systems without interrupting customer peering arrangements. Customer peering sessions can later be updated during a maintenance window or during other scheduled downtime. | Configuring BGP Neighbor Session Options |
| 12.3(11)T | BGP Support for Named Extended Community Lists | The BGP Support for Named Extended Community Lists feature introduces the ability to configure extended community lists using names in addition to the existing numbered format. | Connecting to a Service Provider Using External BGP |
| 12.3(11)T | BGP Support for Sequenced Entries in Extended Community Lists | The BGP Support for Sequenced Entries in Extended Community Lists feature introduces automatic sequencing of individual entries in BGP extended community lists. This feature also introduces the ability to remove or resequence extended community list entries without deleting the entire existing extended community list. | Connecting to a Service Provider Using External BGP |
| 12.3(8)T | BGP Cost Community Support for EIGRP MPLS VPN PE-CE | The BGP Cost Community Support for EIGRP MPLS VPN PE-CE feature provides BGP cost community support for mixed EIGRP MPLS VPN network topologies that contain back door routes. | BGP Cost Community |
| 12.3(7)T | BGP MIB Support Enhancements | The BGP MIB Support Enhancements feature introduced support in the CISCO-BGP4-MIB for new SNMP notifications. | Configuring Advanced BGP Features |
| 12.3(7)T | BGP Support for TTL Security Check | The BGP Support for TTL Security Check feature introduced a lightweight security mechanism to protect external Border Gateway Protocol (eBGP) peering sessions from CPU utilization-based attacks using forged IP packets. Enabling this feature prevents attempts to hijack the eBGP peering session by a host on a network segment that is not part of either BGP network or by a host on a network segment that is not between the eBGP peers. | Configuring BGP Neighbor Session Options |
| 12.3(4)T | BGP Configuration Using Peer Templates | The BGP Configuration Using Peer Templates feature introduces a new mechanism that groups distinct neighbor configurations for BGP neighbors that share policies. Configuration templates provide an alternative to peer group configuration and overcome some of the limitations of peer groups. | Configuring a Basic BGP Network |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|----------|---|---|---|
| 12.3(4)T | BGP Dynamic Update Peer Groups | The BGP Dynamic Update Peer Groups feature introduces a new algorithm that dynamically calculates and optimizes update groups of neighbors that share outbound policies and can share update messages. In previous versions of Cisco IOS software, BGP update messages were grouped based on peer group configurations. This method of grouping updates limited outbound policies and specific-session configurations. The BGP Dynamic Update Peer Groups feature separates update group replication from peer group configuration, which improves convergence time and flexibility of neighbor configuration. | Configuring a Basic BGP Network |
| 12.3(4)T | BGP Policy Accounting Output Interface Accounting | Border Gateway Protocol (BGP) policy accounting (PA) measures and classifies IP traffic that is sent to, or received from, different peers. Policy accounting was previously available on an input interface only. The BGP Policy Accounting Output Interface Accounting feature introduces several extensions to enable BGP PA on an output interface and to include accounting based on a source address for both input and output traffic on an interface. Counters based on parameters such as community list, autonomous system number, or autonomous system path are assigned to identify the IP traffic. | BGP Policy Accounting Output Interface Accounting |
| 12.3(4)T | Regex Engine Performance Enhancement | The Regex Engine Performance Enhancement feature introduces a new regular expression engine that is designed to process complex regular expressions. This new regular expression engine does not replace the existing engine. The existing engine is preferred for simple regular expressions and is the default engine and in Cisco IOS software. Either engine can be selected from the command-line interface (CLI). | Regex Engine Performance Enhancement |
| 12.3(2)T | BGP Cost Community | The BGP Cost Community feature introduces the cost extended community attribute. The cost community is a non-transitive extended community attribute that is passed to internal BGP (iBGP) and confederation peers but not to external BGP (eBGP) peers. The cost community feature allows you to customize the local route preference and influence the best path selection process by assigning cost values to specific routes. | BGP Cost Community |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|-----------|---|---|--|
| 12.3(2)T | BGP Route-Map Continue | The BGP Route-Map Continue feature introduces the continue clause to BGP route map configuration. The continue clause allows for more programmable policy configuration and route filtering and introduces the capability to execute additional entries in a route map after an entry is executed with successful match and set clauses. Continue clauses allow the network operator to configure and organize more modular policy definitions so that specific policy configurations need not be repeated within the same route map. | Connecting to a Service Provider Using External BGP |
| 12.3(2)T | Loadsharing IP Packets Over More Than Six Parallel Paths | The Loadsharing IP Packets Over More Than Six Parallel Paths feature increases the maximum number of parallel routes that can be installed to the routing table for multipath loadsharing. | Loadsharing IP Packets Over More Than Six Parallel Paths |
| 12.2(15)T | BGP Hybrid CLI | The BGP Hybrid CLI feature simplifies the migration of BGP networks and existing configurations from the NLRI format to the AFI format. This new functionality allows the network operator to configure commands in the AFI format and save these command configurations to existing NLRI formatted configurations. The feature provides the network operator with the capability to take advantage of new features and provides support for migration from the NLRI format to the AFI format. | Configuring a Basic BGP Network |
| 12.2(15)T | BGP Increased Support of Numbered AS-Path Access Lists to 500 | The BGP Increased Support of Numbered AS-Path Access Lists to 500 feature increases the maximum number of autonomous systems access lists that can be configured using the ip as-path access-list command from 199 to 500. | Connecting to a Service Provider Using External BGP |
| 12.2(15)T | BGP Nonstop Forwarding (NSF) Awareness | Nonstop Forwarding (NSF) awareness allows a router to assist NSF-capable neighbors to continue forwarding packets during a Stateful Switchover (SSO) operation. The BGP Nonstop Forwarding Awareness feature allows an NSF-aware router that is running BGP to forward packets along routes that are already known for a router that is performing an SSO operation. This capability allows the BGP peers of the failing router to retain the routing information that is advertised by the failing router and continue to use this information until the failed router has returned to normal operating behavior and is able to exchange routing information. The peering session is maintained throughout the entire NSF operation. | Configuring Advanced BGP Features |

Table 1 Supported BGP Features (continued)

| Release | Feature Name | Feature Description | Where Documented |
|-----------|--|--|---|
| 12.2(15)T | BGP Restart Session After Max-Prefix Limit | The BGP Restart Session After Max-Prefix Limit feature enhances the capabilities of the neighbor maximum-prefix command with the introduction of the restart keyword. This enhancement allows the network operator to configure the time interval at which a peering session is reestablished by a router when the number of prefixes that have been received from a peer has exceeded the maximum prefix limit. | Configuring BGP Neighbor Session Options |
| 12.2(15)T | BGP Route-Map Policy List Support | The BGP Route-Map Policy List Support feature introduces new functionality to BGP route maps. This feature adds the capability for a network operator to group route map match clauses into named lists called policy lists. A policy list functions like a macro. When a policy list is referenced in a route map, all of the match clauses are evaluated and processed as if they had been configured directly in the route map. This enhancement simplifies the configuration of BGP routing policy in medium-size and large networks because a network operator can preconfigure policy lists with groups of match clauses and then reference these policy lists within different route maps. The network operator no longer needs to manually reconfigure each recurring group of match clauses that occur in multiple route map entries. | Connecting to a Service Provider Using External BGP |
| 12.2(13)T | BGP Policy Accounting | Border Gateway Protocol (BGP) policy accounting measures and classifies IP traffic that is sent to, or received from, different peers. Policy accounting is enabled on an input interface, and counters based on parameters such as community list, autonomous system number, or autonomous system path are assigned to identify the IP traffic. | BGP Policy Accounting |
| 12.2(8)T | BGP Named Community Lists | The BGP Named Community Lists feature introduces a new type of community list called the named community list. The BGP Named Community Lists feature allows the network operator to assign meaningful names to community lists and increases the number of community lists that can be configured. A named community list can be configured with regular expressions and with numbered community lists. All rules of numbered communities apply to named community lists except that there is no limitation on the number of community attributes that can be configured for a named community list. | Connecting to a Service Provider Using External BGP |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|----------------|---|---|---|
| 12.2(8)T | Multiprotocol BGP (MP-BGP) Support for the CLNS | The Multiprotocol BGP (MP-BGP) Support for CLNS feature provides the ability to scale Connectionless Network Service (CLNS) networks. The multiprotocol extensions of Border Gateway Protocol (BGP) add the ability to interconnect separate Open System Interconnection (OSI) routing domains without merging the routing domains, thus providing the capability to build very large OSI networks. | Multiprotocol BGP (MP-BGP) Support for the CLNS |
| 12.2(4)T | BGP Conditional Route Injection | The BGP Conditional Route Injection feature allows you to inject more specific prefixes into a BGP routing table over less specific prefixes that were selected through normal route aggregation. These more specific prefixes can be used to provide a finer granularity of traffic engineering or administrative control than is possible with aggregated routes. | Configuring a Basic BGP Network |
| 12.2(4)T | BGP Multipath Load Sharing for eBGP and iBGP in an MPLs VPN | The BGP Multipath Load Sharing for eBGP and iBGP feature allows you to configure multipath load balancing with both external BGP (eBGP) and internal BGP (iBGP) paths in Border Gateway Protocol (BGP) networks that are configured to use Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs). This feature provides improved load balancing deployment and service offering capabilities and is useful for multi-homed autonomous systems and Provider Edge (PE) routers that import both eBGP and iBGP paths from multihomed and stub networks. | BGP Multipath Load Sharing for eBGP and iBGP in an MPLs VPN |
| 12.2(4)T | BGP Prefix-Based Outbound Route Filtering | The BGP Prefix-Based Outbound Route Filtering feature uses BGP ORF send and receive capabilities to minimize the number of BGP updates that are sent between BGP peers. Configuring this feature can help reduce the number of system resources required for generating and processing routing updates by filtering out unwanted routing updates at the source. For example, this feature can be used to reduce the amount of processing required on a router that is not accepting full routes from a service provider network. | Connecting to a Service Provider Using External BGP |

Table 1 **Supported BGP Features (continued)**

| Release | Feature Name | Feature Description | Where Documented |
|----------|-----------------------------|---|---|
| 12.2(2)T | BGP Link Bandwidth | The Border Gateway Protocol (BGP) Link Bandwidth feature is used to advertise the bandwidth of an autonomous system exit link as an extended community. This feature is configured for links between directly connected external BGP (eBGP) neighbors. The link bandwidth extended community attribute is propagated to iBGP peers when extended community exchange is enabled. This feature is used with BGP multipath features to configure load balancing over links with unequal bandwidth. | BGP Link Bandwidth |
| 12.2(2)T | iBGP Multipath Load Sharing | The iBGP Multipath Load Sharing feature enables the BGP speaking router to select multiple iBGP paths as the best paths to a destination. The best paths or multipaths are then installed in the IP routing table of the router. | iBGP Multipath Load Sharing |

CCDE, CCENT, CCSI, Cisco Eos, Cisco HealthPresence, Cisco Ironport, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Stackpower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flip Video, Flip Video (Design), Flipshare (Design), Flip Ultra, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Store, and Flip Gift Card 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 IOS, 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, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, 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. (0907R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2007–2009 Cisco Systems, Inc. All rights reserved.

