CHAPTER 10

Routing Protocols

This chapter describes the level of support that Cisco ANA provides for routing protocols, as follows:

Technology Description

BGP

The Border Gateway Protocol (BGP) routes traffic between autonomous systems. An autonomous system is a network or group of networks under common administration and with common routing policies. BGP exchanges routing information for the Internet and is the protocol used between ISPs. Customer networks, such as universities and corporations, usually employ an Interior Gateway Protocol (IGP), such as RIP or OSPF, to exchange routing information within their networks. Customers connect to ISPs, and ISPs use BGP to exchange customer and ISP routes. When BGP is used between autonomous systems, the protocol is referred to as external BGP (eBGP). If a service provider is using BGP to exchange routes within an autonomous system, the protocol is referred to as interior BGP (iBGP).

BGP is a very robust and scalable routing protocol, as evidenced by the fact that it is the routing protocol employed on the Internet. To achieve scalability at this level, BGP uses many route parameters, called attributes, to define routing policies and maintain a stable routing environment. BGP neighbors exchange full routing information when the TCP connection between neighbors is first established. When changes to the routing table are detected, the BGP routers send to their neighbors only those routes that have changed. BGP routers do not send periodic routing updates, and BGP routing updates advertise only the optimal path to a destination network.

MP-BGP

Multiprotocol BGP (MP-BGP) adds capabilities to BGP to enable multicast routing policy throughout the Internet and to connect multicast topologies within and between BGP autonomous systems. That is, MP-BGP is an enhanced BGP that carries IP multicast routes. BGP carries two sets of routes, one set for unicast routing and one set for multicast routing. The routes associated with multicast routing are used by the Protocol Independent Multicast (PIM) to build data distribution trees.

eBGP/iBGP

As noted previously, BGP is an interautonomous system routing protocol. When BGP is used between autonomous systems (AS), the protocol is referred to as external BGP (eBGP). If a service provider is using BGP to exchange routes within an AS, then the protocol is referred to as interior BGP (iBGP).

OSPF

Open Shortest Path First (OSPF) is a routing protocol developed for IP networks by the IGP working group of the Internet Engineering Task Force (IETF). It was derived from several research efforts, including a version of OSI's IS-IS routing protocol.

OSPF has two primary characteristics:

- It is an open protocol. Its specification is in the public domain (RFC 1247).
- It is based on the Shortest Path First (SPF) algorithm, sometimes known as the Dijkstra algorithm.

OSPF is a link-state routing protocol that calls for the sending of link-state advertisements (LSAs) to all other routers within the same hierarchical area. Information on attached interfaces, metrics used, and other variables are included in OSPF LSAs. As OSPF routers accumulate link-state information, they use the SPF algorithm to calculate the shortest path to each node.

EIGRP

Enhanced Interior Gateway Routing Protocol (EIGRP) is Cisco's proprietary routing protocol, based on IGRP. EIGRP is a distance-vector routing protocol, with optimizations to minimize routing instability incurred after topology changes, and the use of bandwidth and processing power in the router. Routers that support EIGRP will automatically redistribute route information to IGRP neighbors by converting the 32-bit EIGRP metric to the 24-bit IGRP metric. Most of the routing optimizations are based on the Diffusing Update Algorithm (DUAL), which guarantees loop-free operation and provides fast router convergence.

RIP

The Routing Information Protocol (RIP) is one of the oldest routing protocols still in wide use. Today's open standard version of RIP, sometimes referred to as IP RIP, is formally defined in RFC 1058 and in STD 56. RIP is a distance-vector routing protocol that uses hop count as a metric. RIP prevents routing loops by implementing a limit on the number of hops allowed in source/destination paths, and also implements split horizon, route poisoning and holddown mechanisms to prevent incorrect routing information from being propagated.

IS-IS

Intermediate system to intermediate system (IS-IS) is a link-state routing protocol. It operates by reliably flooding topology information throughout a network of routers. Each router then builds its own picture of the network's topology. Packets or datagrams are forwarded based on the best topological path through the network; IS-IS uses Dijkstra's algorithm for computing best paths. IS-IS was first defined in ISO/IEC 10589:2002 and was republished in RFC 1142 for the Internet community. IS-IS is an IGP, intended for use within one administrative domain or network only.

Information Model Objects (IMOs)

This section describes the following IMOs:

- BGP Neighbor Entry (IBgpNeighbourEntry)
- OSPF Entry (IOspfEntry)

BGP Neighbor Entry

The BGP Neighbor Entry IMO represents both the configuration and the outcome of running the Border Gateway Protocol (BGP) within a group of participating routers in a BGP neighborhood. It contains information about the connection with a remote BGP peer. It is the entry of the BGP Neighbors Table attribute of the Multi Protocol BGP Entity object (see Virtual Routing and Forwarding), representing the BGP routing service concept in the IMO.

Table 10-1 BGP Neighbor Entry (IBgpNeighbourEntry)

Attribute Name	Attribute Description	Scheme	Polling Interval
Remote Identifier	Identifier of the remote peer (*IP Address*)	Product	Configuration
Neighbor Type	Neighbor type (*Null, Client, Non Client)*	Any	Configuration
Distributing Interface	Distributing IP interface	Any	Configuration
Remote Address	Remote peer IP address	Product	Configuration
Remote Autonomous System	Remote peer autonomous system	Product	Configuration
Status	Status (*Null, Idle, Connect, Active, Open Sent, Open Confirm, Established*)	Product	Configuration
Hold Time	Established hold time in seconds	Product	Configuration
Keep Alive Time	Established keepalive time in seconds	Product	Configuration

OSPF Entry

The OSPF Entry IMO represents both the configuration and the outcome of running a single Open Shortest Path First (OSPF) protocol interface within a group of participating OSPF routers. It is aggregated by a Traffic Descriptor Container object (see Common Components).

Table 10-2 OSPF Entry (IOspfEntry)

Attribute Name	Attribute Description	Scheme	Polling Interval
Area Identifier	Area identifier (*IP address*)	Product	Configuration
IP Address	IP address	Product	Configuration
Type	OSPF type (*Null, Broadcast, NBMA, Point-to-Point, Point-to-Multipoint*)	Product	Configuration
Administrative Status	Administrative status (*Null, Enabled, Disabled*)	Product	Status
Operational Status	Operational status (*Null, Down, Loop Back, Waiting, Point-to-Point, Designated Router, Backup Designated Router, Other Designated Router*)	Product	Status

Vendor-Specific Inventory and IMOs

There are no vendor-specific inventory or IMOs for this technology.

Network Topology

Discovery of BGP neighborhood topology is done by comparing BGP router parameters on either side of potential BGP neighbors. In particular, the local and remote BGP router identification and autonomous system, as well as the connection states on both sides, are compared.

Service Alarms

The following alarms are supported for this technology:

- BGP Neighbor Loss/BGP Neighbor Found
- BGP Process Down/BGP Process Up
- BGP Link Down/BGP Link Up

For detailed information about alarms and correlation, see the *Cisco Active Network Abstraction 3.7 User Guide*.