Chapter 1
General Networking

Networking Basics

The International Standards Organization (ISO) developed the Open Systems Interconnection (OSI) model to enable the delineation of various functions performed by devices in the network and applications. The OSI model is a conceptual model that is composed of seven layers, each specifying a particular network function. Data being transferred from a software application on one computer system to another computer system must pass through the OSI layers. For example, if a software application in System A has information to transmit to a software application in System B, the application program in System A passes its information to the application layer (Layer 7) of System A. The application layer then passes the information to the presentation layer (Layer 6), which relays the data to the session layer (Layer 5), and so on down to the physical layer (Layer 1). At the physical layer, the information is placed on the physical network medium and is sent across the medium to System B. The physical layer of System B removes the information from the physical medium, and then its physical layer passes the information up to the data link layer (Layer 2), which passes it to the network layer (Layer 3), and so on, until it reaches the application layer (Layer 7) of System B. Finally, the application layer of System B passes the information to the recipient application program to complete the communication process. Figure 1-1 outlines the OSI model and each layer's function.

This URL provides further information on the OSI layers: http://docwiki.cisco.com/wiki/Internetworking_Basics#Open_System_Interconnection_Reference_Model.

 Please see page 131 for more details.

FIGURE 1-1 OSI Model

Application Layer	• Interface to the end user on the OSI stack • Examples: Telnet, FTP, SMTP
Presentation Layer	• Enables parity when information is transmitted between multiple systems at the application layer • Defines coding and conversion algorithms that are applied to data from the application layer • Examples: ASCII, JPEG, TIFF, MP3
Session Layer	• Manages session establishment, upkeep, and teardown between devices • Examples: H323, RTCP
Transport Layer	• Responsible for segmentation of information received from higher layers prior to network layer handoff • Also provides reliable data transport for some protocols • Fundamental entity is called a Layer 4 segment or datagram • Examples: TCP, UDP, RTP
Network Layer	• Identifies the optimal path to a specific network destination by means of routing decision • Also responsible for device identification using IP addressing • Fundamental entity is called a Layer 3 packet • Examples: IP, IPX
Data Link Layer	• Primarily performs the functions associated with transmission of data across a link reliably • Error notification, flow control, and frame sequencing are also performed by the data link layer • Consists of two sublayers: logical link control (LLC), which enables communication of devices over a single link, and MAC, which provides the means for protocols to access the physical layer media • Fundamental entity is called a Layer 2 frame • Examples: ISDN, PPP, HDLC, SDLC, Ethernet and its variants, Frame Relay
Physical Layer	• Consists of standards that define hardware specifications such as cables, connectors, NICs, electrical and mechanical specifications, bit ordering, encoding, signaling, and transmission rates • Examples: RS-232, V.35, T1, E1, 10BASE-T, 100BASE-TX, POTS, SONET, DSL, 802.11x, RJ-45

 Please see page 131 for more details.

Ethernet in a Nutshell

Ethernet uses carrier sense multiple access collision detect (CSMA/CD) to detect collisions on the Ethernet broadcast domain. Devices operating in full-duplex mode do not implement CSMA/CD.

CSMA/CD enables devices to transmit data when no other device on the broadcast domain is doing the same. In the event of contention, the contending devices implement a backoff algorithm, which forces a wait time before they attempt to access the network again to send data.

For more information about Ethernet specifications and limitations, see the Cisco Ethernet overview located at http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/ethernet.htm#wp1020792.

Bridging and Switching

Forwarding frames from one interface to the other is called *switching* or *bridging*; the forwarding decision is based on the destination MAC address.

Spanning Tree Protocol (STP) ensures loop-free topology between switches and/or bridges in a Layer 2 domain. During spanning-tree operation (which runs on all Cisco switches), a root bridge is elected based on *bridge priority* (lower priority preferred, range 0–65,535, default 32,768). If LAN switches within the domain have the same bridge priority, the lowest MAC address selects the root bridge.

MAC addresses of end stations are stored in the content-addressable memory (CAM) table on the switches. The CAM table associates MAC addresses with network interfaces. When receiving frames on a switch, the source address and incoming network interface is added to the CAM table. Frames whose destination is not identified in the CAM table are broadcast out all ports associated with the VLAN.

A VLAN consists of a group of devices (that can span across multiple switches) that function as if they were on a single broadcast domain. VLANs enable the operation of many virtual bridge domains over a shared LAN. By default, VLAN 1 is used for management purposes on all switches (*native VLAN*).

Bridges communicate using frames called bridge protocol data units (BPDU). BPDUs are sent out all ports that are not in a blocking state. A root bridge always has all of its ports in the forwarding state. To ensure a loop-free topology, nonroot bridges block any paths

to the root that are not required. BPDUs use the destination MAC address 01-08-C2-00-00-00 in Ethernet environments.

Bridge Port States

802.1D Spanning Tree Protocol defines the following port states:

- **Disabled:** The port does not participate in spanning tree.
- **Listening:** The port listens for frames, but it does not forward frames to the interface.
- **Learning:** The port does not forward frames out this port, but the source address of the end station attached to the port is added to the CAM table.
- **Forwarding:** The port forwards and receives frames on the interface.
- **Blocking:** Spanning tree has placed this port in blocking state to avoid a loop.

802.1W (Rapid Spanning Tree Protocol [RSTP]) was introduced for faster convergence. RSTP has only three operational states: discarding, learning, and forwarding. The 802.1D disabled, blocking, and listening states are merged into a unique 802.1w discarding state.

The following URL provides a deeper understanding of RSTP: http://www.cisco.com/en/US/tech/tk389/tk621/technologies_white_paper09186a0080094cfa.shtml.

EtherChannel and Trunking

Bundling Ethernet, Fast Ethernet, or Gigabit Ethernet ports together into a single logical link is called *EtherChannel*; all ports also appear to STP as a single logical link. The ports need to be in the same VLAN or broadcast domain and have the same speed/duplex.

The maximum number of physical ports that can be bundled into an EtherChannel is eight.

The **channel-group** command is used within IOS to configure EtherChannels.

A *trunk* is a physical or logical connection between two switches that carry more than one VLAN.

Inter-Switch Link (ISL) is a Cisco proprietary VLAN tagging protocol. 802.1Q is the IEEE standard VLAN trunking protocol.

 Please see page 131 for more details.

For more information about EtherChannel load balancing, see http://www.cisco.com/warp/public/473/4.html.

IP Overview

IP is the network layer protocol in the Internet Protocol suite, and it is encapsulated in a data link layer protocol. By default, IP provides best-effort service.

IP Version 4 is the forth iteration of IP, and it is the first version of the protocol to be widely deployed. It uses 32-bit (4-byte) addresses.

IPv6 is a successor of IPv4. The main feature of IPv6 that is driving adoption today is the larger address space. Addresses in IPv6 are 128 bits long versus 32 bits in IPv4.

The *ToS bit* in the IP header identifies the priority of the packet when upper-layer protocols handle the packet. It has eight IP precedence values: 000-Routine, 001-priority, 010-immediate, 011-flash, 100-flash override, 101-critical (VoIP, real-time applications), 110-internetwork control, and 111-network control.

Differentiated services (DiffServ) is a new model in which traffic is treated by intermediate systems with relative priorities based on the Type of Service (ToS) field. Defined in RFC 2474 and RFC 2475, the DiffServ standard supersedes the original specification for defining packet priority described in RFC 791. DiffServ increases the number of definable priority levels by reallocating bits of an IP packet for priority marking. The 6 most significant bits of the DiffServ field are called the *Differentiated Services Code Point* (DSCP). The last 2 Currently Unused (CU) bits in the DiffServ field were not defined within the DiffServ field architecture; these are now used as Explicit Congestion Notification (ECN) bits. Routers at the edge of the network classify packets and mark them with either the IP Precedence or DSCP value in a DiffServ network. Other network devices in the core that support DiffServ use the DSCP value in the IP header to select a PHB behavior for the packet and provide the appropriate QoS treatment. *Flags* identify whether the packet can be fragmented (2 lower order bits of 3 total bits are used) in the IP header.

The *Protocol field* identifies the higher layer transport protocol. For a complete list of protocol numbers, see http://www.iana.org/assignments/protocol-numbers.

Figure 1-2 outlines the IP header format.

FIGURE 1-2 IPv4 Header Format

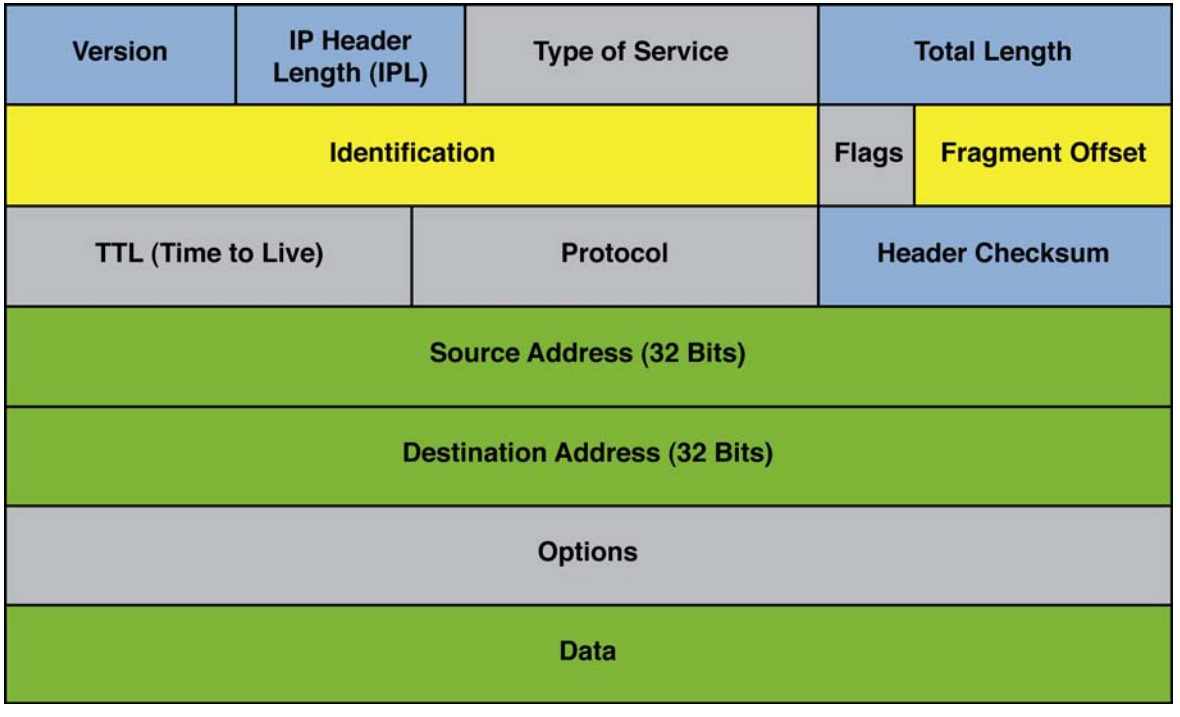

Subnetting, Variable-Length Subnet Masking, and Classless Interdomain Routing

This URL outlines the fundamentals of IP addressing, subnetting (including variable length subnet masking [VLSM]), and classless interdomain routing (CIDR): http://www.cisco.com/warp/public/701/3.html.

IPv6

RFC 2460 defines IP Version 6 (IPv6) as a successor to IP Version 4 (IPv4). The motivation behind IPv6 is primarily to counter the IPv4 route exhaustion. Figure 1-3 outlines the IPv6 header format.

FIGURE 1-3 IPv6 Header Format

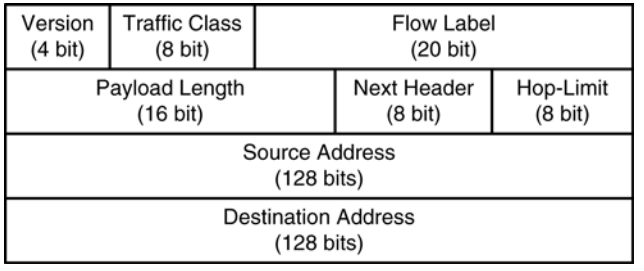

- **Version:** The first header field in an IP Packet is the 4-bit Version field. For IPv6, this has a value of 6.
- **Traffic Class:** The 8-bit Traffic Class field allows the source device to associate different priority levels to different packets generated by it. A router uses the traffic class to identify the priority level of the associated data packet.
- **Flow Label:** 20-bit field used by the source to label packet sequences that require special handling by the routers.
- **Payload Length:** 16-bit unsigned integer. Length of the IPv6 payload in octets (for example, the rest of the packet following this IPv6 header). Note that any extension headers present are considered part of the payload, meaning that they included in the length count.
- **Next Header:** 8-bit selector. Identifies the protocol type immediately following the IPv6 header. Uses the same values as the IPv4 Protocol field.
- **Hop-Limit:** 8-bit unsigned integer. Decremented by 1 by each node that forwards the packet. The packet is discarded if Hop-Limit is decremented to zero.
- **Source Address:** 128-bit address of the originator of the packet.
- **Destination Address:** 128-bit address of the intended recipient of the packet.

IPv6 supports 128-bit address space as compared to the 32-bit address space provided by IPv4. With this larger address space, IPv6 has the capability to provide unique addresses to each and every device or node attached to the Internet. An escalating demand for IP addresses was the driving force behind the development of the larger address space for IPv6.

IPv4 and IPv6 addresses are broadly classified into three categories:

 Please see page 131 for more details.

- **Unicast addresses** act as an identifier for a single interface. An IPv6 packet sent to a unicast address is delivered to the interface identified by that address.
- **Multicast addresses** act as an identifier for a group/set of interfaces that can belong to different nodes. An IPv6 packet destined to a multicast address can be delivered to multiple interfaces.
- **Anycast addresses** act as identifiers for a set of interfaces that can belong to the different nodes. An IPv6 packet destined for an anycast address is delivered to one of the interfaces identified by the anycast address.

Transmission Control Protocol

TCP is a connection-oriented Layer 4 (transport layer) protocol designed to provide reliable end-to-end transmission of data in an IP environment. TCP groups bytes into sequenced segments, and then passes them to IP for delivery. These sequenced bytes have forward acknowledgment numbers that indicate to the destination host what next byte it should see. Bytes not acknowledged to the source host within a specified time period are retransmitted, which allows devices to deal with lost, delayed, duplicate, or misread packets. This timeout mechanism allows devices to detect lost packets and request retransmission. The receiving TCP process indicates the highest sequence number it can receive without overflowing its internal buffers. TCP hosts establish a connection-oriented session with one another through a "three-way handshake" mechanism, which synchronizes both ends of a connection by allowing both sides to agree upon initial sequence numbers. Each host first randomly chooses a sequence number to use in tracking bytes within the stream it is sending and receiving.

Then, the three-way handshake proceeds in the following manner:

- The initiating host (Host A) initiates a connection by sending a packet with the initial sequence number (X) and SYN bit (or flag) set to make a connection request of the destination host (Host B).
- Host B receives the SYN bit, records the sequence number of X, and replies by acknowledging the SYN (with an ACK=X+1).
- Host B includes its own initial sequence number (Y). For example, an ACK of 20 means that Host B has received bytes 0 through 19 and expects byte 20 next. This technique is called *forward acknowledgment*.
- Host A then acknowledges all bytes from Host B with a forward acknowledgment indicating the next byte Host A expects to receive (ACK=Y+1).

 Please see page 131 for more details.

- Data transfer can now begin.

The flags used in the TCP handshaking and data transfer processes are as follows:

- **URG:** Urgent Pointer field significant.
- **ACK:** Acknowledgment field significant.
- **PSH:** Push function tells the receiver to send the data as soon as possible.
- **RST:** Reset the connection.
- **SYN:** Synchronize sequence numbers to initiate a connection.
- **FIN:** No more data from sender.

The SYN and ACK flags are used for connection establishment. The FIN and RST flags are used for connection termination.

Hot Standby Routing Protocol

Hot Standby Routing Protocol (HSRP) provides default gateway redundancy by making two or more routers/switches share a single virtual IP address that is used as a default gateway for end stations on the device connected on the segment. Routers configured to share a single virtual IP address that functions as a default gateway are called *HSRP groups*. A router functions either in an active or standby state when operating with HSRP. The router in active state performs packet-forwarding functions; a router in the standby state is ready to take over the packet-forwarding functions if the router in the active state fails.

Figure 1-4 outlines the configuration flowchart for HSRP and shows a basic configuration for HSRP operation.

 Please see page 131 for more details.

FIGURE 1-4 HSRP Configuration Flowchart

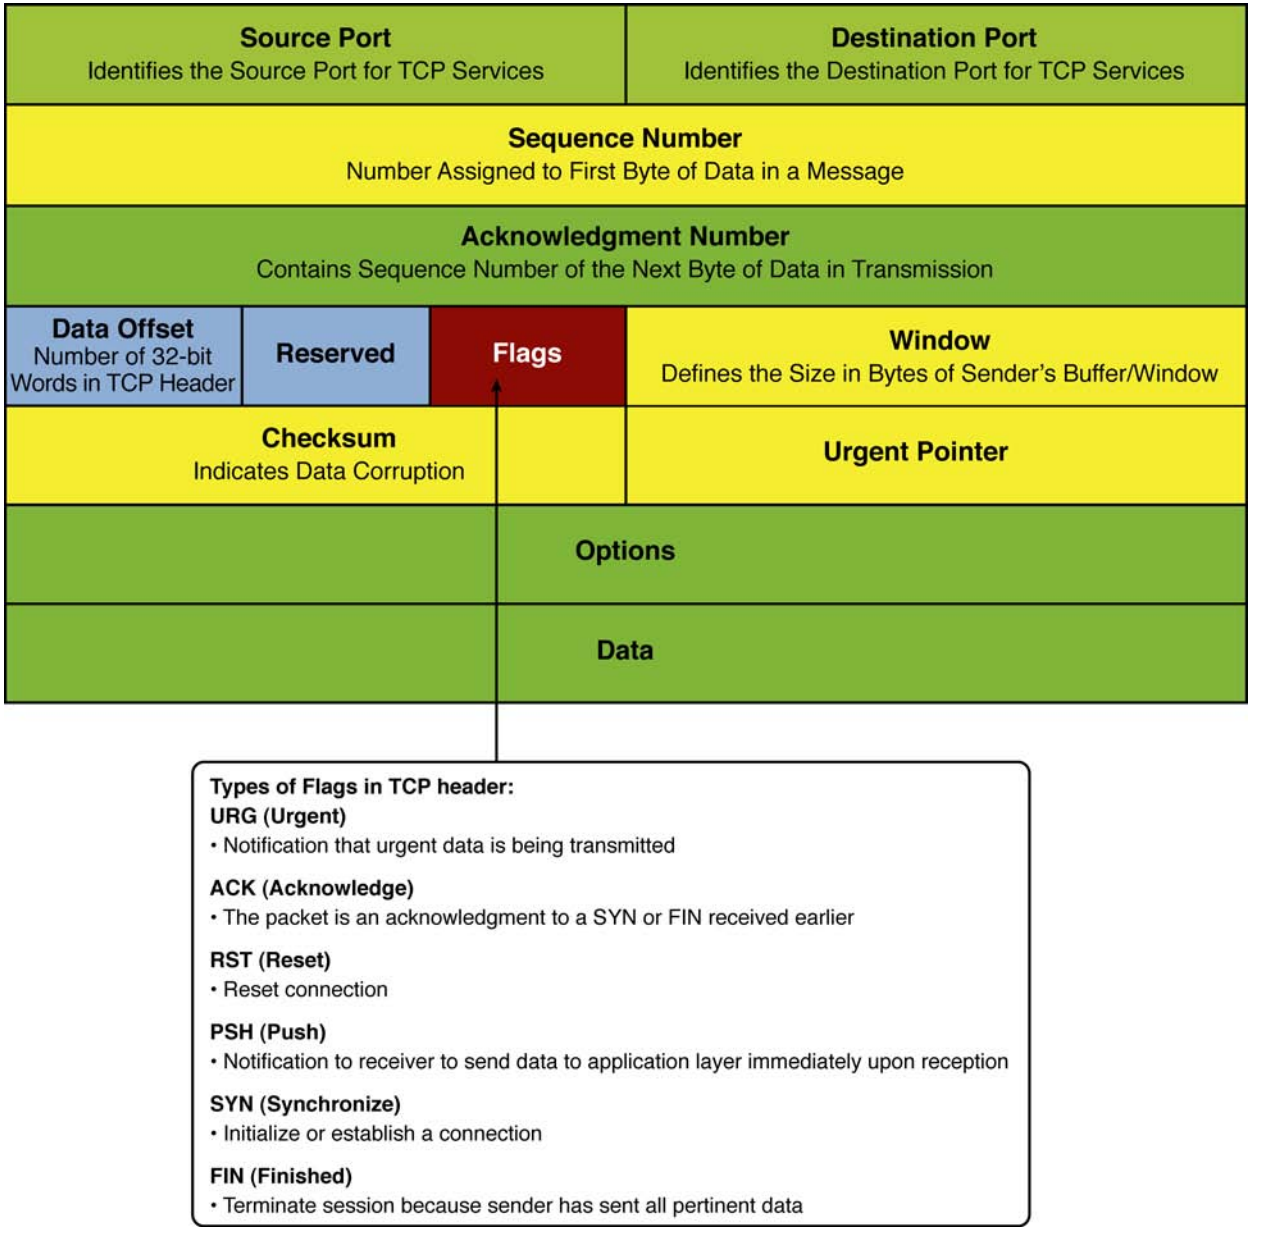

Table 1-1 lists the default values for HSRP.

Table 1-1 HSRP Default Values

HSRP Parameters	HSRP Values
Standby group number	0–255.
Standby MAC address	System assigned as 0000.0c07.ac*XX*, where *XX* is the HSRP group number.
Standby priority	Default is 100, range is 0–255. (Higher priority is preferred as active in HSRP group.)
Standby delay	Default is 0 delay, both minimum and reload delays can be set in a range of 0–10,000 seconds.
Standby track interface priority	Default is 10.
Standby hello time	3 seconds. (When configured for **msec** option, range is 15–999 milliseconds.)
Standby hold time	10 seconds. (When configured for **msec** option, range is 50–3000 milliseconds.)

Virtual Router Redundancy Protocol

Virtual Router Redundancy Protocol (VRRP) is the standard implementation of HSRP in the industry that is not Cisco proprietary. VRRP process and operation is similar to that of HSRP. For more information on VRRP and its configuration, visit http://www.cisco.com/en/US/docs/ios/12_0st/12_0st18/feature/guide/st_vrrpx.html.

Generic Routing Encapsulation

Generic routing encapsulation (GRE) is an IP-tunneling mechanism that encapsulates a multitude of protocol packet types inside an IP tunnel, thereby creating logical point-to-point links to remote routers over an IP network. RFC 2784 highlights the various characteristics of this tunneling protocol.

Figure 1-5 highlights the composition of the GRE header. GRE introduces additional 24 bytes of overhead to the original IP packet. For further information about GRE, refer to http://www.cisco.com/en/US/tech/tk827/tk369/tk287/tsd_technology_support_sub-protocol_home.html and RFC 2784.

 Please see page 131 for more details.

FIGURE 1-5 GRE Header

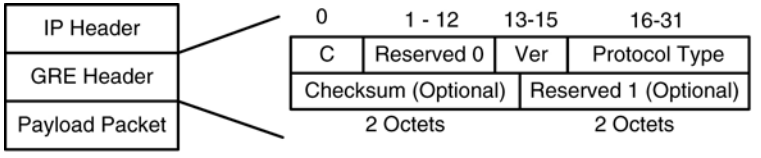

1. The Protocol Type field contains the protocol type of the payload packet.
2. When IPv4 is being carried as the GRE payload, the Protocol Type field MUST be set to 0x800.
3. The IPv4 protocol 47 [RFC1700] is used when GRE packets are encapsulated in IPv4.

Next Hop Resolution Protocol

RFC 2332 defines Next Hop Resolution Protocol (NHRP) as a protocol that allows a host or router to communicate over a non-broadcast multiple access (NMBA) network by learning the NBMA address of the another host or remote router. NHRP is a client server protocol where hubs are consider as next-hop servers (NHS) and spokes are considered next-hop clients (NHC). The hub maintains an NHRP database of the public interface addresses of each spoke. Each spoke registers its real address when it boots and queries the NHRP database for real addresses of the destination spokes to build direct tunnels. With NHRP, systems attached to an NBMA network dynamically learn the NBMA address of the other systems that are part of that network, allowing these systems to directly communicate without requiring traffic to use an intermediate hop.

NHRP allows two functions to help support these NBMA networks:

- **NHRP registration:** NHRP allows NHCs to dynamically register with NHSs. This registration function allows the NHCs to join the NBMA network without configuration changes on the NHSs.
- **NHRP resolution:** NHRP allows one NHC (spoke) to dynamically discover the logical VPN IP to physical NBMA IP mapping for another NHC (spoke) within the same NBMA network. Without this discovery, IP packets traversing from hosts behind one spoke to hosts behind another spoke would have to traverse by way of the NHS (hub) router. This process increases the utilization of the hub's physical bandwidth and CPU to process these packets that enter and exit the hub on the multipoint interface. With NHRP, systems attached to an NBMA network dynamically learn the NBMA address of the other systems that are part of that network, which allows these systems

 Please see page 131 for more details.

to directly communicate without requiring traffic to use an intermediate hop. This function alleviates the load on the intermediate hop (NHS) and can increase the overall bandwidth of the NBMA network to be greater than the bandwidth of the hub router.

Dynamic Multipoint VPN

Dynamic Multipoint VPN (DMVPN) allows large and small IPsec VPNs to scale by combining GRE tunnels, IPsec encryption, and NHRP routing. DMVPN also simplifies configuration via crypto profiles, which override the requirement for defining static crypto maps, and enables dynamic discovery of tunnel endpoints. This feature requires the following two key technologies:

- **NHRP:** A client and server protocol where the hub is the server and the spokes are the clients. The hub maintains an NHRP database of the public interface addresses of the each spoke. Each spoke registers its real address when it boots and queries the NHRP database for real addresses of the destination spokes to build direct tunnels.
- **mGRE tunnel interface:** Allows a single GRE interface to support multiple IPsec tunnels and simplifies the size and complexity of the configuration.

The key benefit of implementing DMVPN is to reduce hub router configuration. This is accomplished by grouping multiple spoke sites into a single multipoint GRE interface. The feature allows IPsec to be immediately triggered for point-to-point tunneling or when the GRE peer address is resolved via NHRP for the mGRE tunnel. With DMPVN, there is no longer a need for hub routers to have knowledge of the physical interface IP address of the spoke routers. DMVPN allows spoke routers to have dynamic physical interface IP addresses, and on coming online sends registration packets containing its physical interface IP address to the hub router. DMVPN eliminates the need for spoke-to-spoke configuration for directs tunnels and allows spoke routers to dynamically create an IPsec tunnel between them so that data can be sent.

Figure 1-6 shows a high-level DMVPN network where R1is the hub router; R2, R3, and R4 being spoke routers.

 Please see page 131 for more details.

FIGURE 1-6 DMVPN

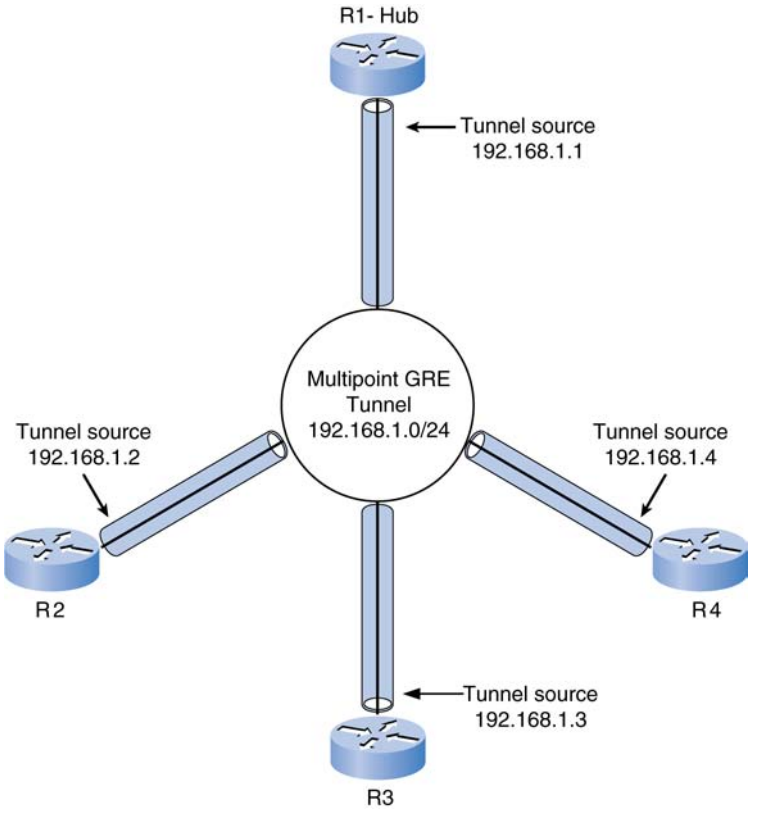

Here is the DMVPN configuration:

```
R1-Hub configuration
interface FastEthernet0/0
 ip address 172.16.1.1 255.255.255.252
!
interface Tunnel0
 ip address 192.168.1.1 255.255.255.0
 ip nhrp map multicast dynamic
 ip nhrp network-id 1
 tunnel source 172.16.1.1
 tunnel mode gre multipoint
```

 Please see page 131 for more details.

As shown in this configuration, tunnel destination is not specified, because multipoint tunnels are built dynamically from the DMVPN spokes to the hub router; the hub router doesn't need to be preconfigured with spoke addresses. Note that the tunnel mode has been designated as multipoint GRE. **ip nhrp network-id 1** uniquely identifies the DMVPN network; tunnels will not form between routers with different network IDs. **ip nhrp multicast dynamic** enables the forwarding of multicast traffic across the tunnel to dynamic spokes.

Here is the configuration of spoke router R2:

```
interface FastEthernet0/0
 ip address 172.16.2.1 255.255.255.252
!
interface Tunnel0
 ip address 192.168.1.2 255.255.255.0
 ip nhrp map 192.168.1.1 172.16.1.1
 ip nhrp map multicast 172.16.1.1
 ip nhrp network-id 1
 ip nhrp nhs 192.168.1.1
 tunnel source 172.16.2.1
 tunnel mode gre multipoint
```

ip nhrp nhs 192.168.1.1 designates R1 as the NHS, and **ip nhrp map 192.168.1.1 172.16.1.1** statically maps the NHS address to R1's physical address. The **ip nhrp multicast** command also differs slightly from its application on the hub in that multicast traffic is only being allowed from spokes to the hub, not from spoke to spoke.

Routing Protocols

Routing Information Protocol (and Routing Information Protocol Version 2)

Routing Information Protocol (RIP) is a distance-vector protocol. RIPv1 is classful, RIPv2 is classless, metric is hop count, and the maximum hop count = 15 hops.

In a classless routing protocol, the netmask is always propagated with the route being advertised; in a classful routing protocol, the netmask cannot be propagated.

 Please see page 131 for more details.

RIPv2 supports authentication for sessions and equal-cost load balancing.

Timers are Update (30Sec), Invalid (180Sec), HoldDown (unused), and Flush (120Sec).

RIPv2 uses multicast addresses to send updates in the network; 224.0.0.9 is the address used to send updates (triggered and normal) to all RIP routers in network.

Configuring RIP

To configure RIP, follow these steps:

Step 1. Enable the RIP routing process by using the command **router rip**.

Step 2. Configure the version number of the RIP process using the **version** command under the RIP routing process.

Step 3. Configure the networks to be enabled for RIP routing using the **network** *network-number* command under the RIP routing process.

Step 4. (Optional) Configure passive interfaces for the RIP routing process to process only inbound RIP updates using the **passive-interface** command. Thus, they do not discover neighbors or form an adjacency out that interface.

Step 5. Authentication is configured under the **interface** configuration using the commands shown in Table 1-2.

Table 1-2 Configuring RIP Authentication

Command	Function
ip rip authentication key-chain *name-of-chain*	Enable RIP authentication on the interface in interface configuration mode
ip rip authentication mode {**text** | **md5**}	Configure authentication mode on interface in interface configuration mode

In addition, key management needs to be configured by defining a key chain. You must also identify the keys that belong to the key chain and specify how long each key is valid. Each key has its own key identifier (specified with the **key** *number* command), which is stored locally. The combination of the key identifier and the interface associated with the message uniquely identifies the authentication algorithm and message digest algorithm 5 (MD5) authentication key in use. Table 1-3 identifies the commands used to configure key management.

 Please see page 131 for more details.

Table 1-3 Configuring Key Management

Command	Function
key chain *name-of-chain*	Defines the name of the key chain
key *number*	Configures a key number
key-string *text*	Configures a key string that is used for authentication
accept-lifetime *start-time* {**infinite** | *end-time* | **duration** *seconds*}	Defines the time period when the key can be received
send-lifetime *start-time* {**infinite** | *end-time* | **duration** *seconds*}	Defines the time period when the key can be sent

Interior Gateway Routing Protocol

Interior Gateway Routing Protocol (IGRP) is a distance-vector protocol, classful in nature. IGRP uses a composite metric that factors in internetwork delay, bandwidth, reliability, and load. It enables unequal-cost load balancing using the **variance** command. IGRP accepts up to four paths to the same destination.

Timers are Update (90Sec), Invalid (270Sec=3xUpdateTimer), HoldDown (280sec=(3xUpdateTimer+10sec)), and Flush (630Sec=7xUpdateTimer).

IGRP metric = [K1 * Bandwidth + (K2 * Bandwidth) / (256 – Load) + K3 * Delay] * [K5 / (Reliability + K4)], where the default constant values are K1 = K3 = 1 and K2 = K4 = K5 = 0.

Configuring IGRP

The steps to configure IGRP are as follows.

Step 1. Enable the IGRP routing process using the **router irgp** *autonomous-system-number* command.

Step 2. Associate networks with an IGRP routing process using the **network** *network-number* command.

Step 3. (Optional) Adjust the IGRP metric weights using the command **metric weights** *tos k1 k2 k3 k4 k5*.

Step 4. Define the variance associated with a particular path to enable unequal-cost load balancing using the command **variance** *multiplier*.

Step 5. Distribute traffic proportionately to the ratios of metrics, or by the minimum-cost route using the **traffic-share** {**balanced** | **min**} command.

For further details on IGRP, refer to http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008009405c.shtml, http://www.cisco.com/en/US/docs/ios/12_0/np1/configuration/guide/1cigrp.html, and http://www.ciscopress.com/articles/article.asp?p=102174&seqNum=.

Open Shortest Path First Protocol.

OSPF is a link-state routing protocol that uses Dijkstra's shortest path first (SPF) algorithm. OSPF is an open standard, and there are several OSPF advantages, including fast convergence, classless routing, VLSM (variable-length subnet mask) support, authentication support, support for much larger internetworks, the use of areas to minimize routing protocol traffic, and a hierarchical design. The backbone area, called Area 0, is the heart of the OSPF domain. All other nonbackbone areas need to be connected to Area 0. In the event they are not, virtual links must be configured to make the area appear like it is connected to Area 0. The virtual links are defined within Area 0 but transit another nonzero area.

Designated router (DR) and backup designated router (BDR) election happens on multiaccess networks. Updates are sent either to AllSPFRouters (224.0.0.5) or to AllDRouters (224.0.0.6), which includes the DR and the BDR.

A router running the OSPF protocol floods link-state advertisements (LSA) over all its adjacencies. The LSAs describe the entire router's links, or interfaces; the router's neighbors; and the state of the links wherein the links might connect to stub networks (other OSPF routers either in the same area or different areas or routers that are not part of the OSPF domain). Because of the varying types of link-state information, OSPF defines multiple LSA types:

- **Type 1 (Router LSA):** Contains information on router and directly connected links; flooded within the area.
- **Type 2 (Network LSA):** Contains information on networks and routers connected to the same; generated by DR, flooded within the area.
- **Type 3 (Summary LSA):** Identifies networks reachable outside the area; generated by the Area Border Router (ABR).

 Please see page 131 for more details.

- **Type 4 (ASBR Summary LSA):** Identifies network reachability to an Autonomous System Boundary Router (ASBR) from an ABR; generated by an ABR.
- **Type 5 (External LSA):** Generated by the ASBR and identifies networks reachable by ASBR; flooded thorough the OSPF domain.
- **Type 7 (External LSA):** To provide the ability of carrying external routes originated in the NSSA, this LSA type was defined. It has the structure and semantics of a Type 5 (External) LSA, with two differences:

  Type 7 LSAs can be originated and propagated within the NSSA; they do not cross area borders like Type 5 LSAs. Type 5 LSAs are not supported in NSSA; they can be neither originated nor propagated in NSSA.
  Translation of Type 7 to Type 5 LSA at the ABR

You can find more information on Type 7 LSA at http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a00801ec9ef.shtml.

For more information about OSPF and configuring OSPF, see the Cisco OSPF design guide at http://www.cisco.com/warp/public/104/1.html#t20 (recommended).

To configure authentication in OSPF, three modes are supported: null, plain text, and MD5. By default, null authentication is used. Table 1-4 identifies the commands required to enable OSPF authentication.

Table 1-4 Configuring OSPF Authentication

Command	Function
ip ospf authentication	Enables OSPF authentication under the interface configuration
ip ospf authentication-key *key*	Configures a plain-text authentication key on the interface
ip ospf message-digest-key *key*	Configures an MD5 authentication key under the interface configuration
area *area-number* **authentication**	Enables all interfaces in an area for plain-text authentication (under OSPF process configuration)
area *area-number* **authentication message-digest**	Enables all interfaces in an area for MD5 authentication (under OSPF process configuration)

Enhanced Interior Gateway Routing Protocol

Enhanced Interior Gateway Routing Protocol (EIGRP) is a hybrid routing protocol; it's classless in nature with a metric calculated using the same formula as IGRP.

Updates do not occur at regular intervals, but only during a network or topology change (triggered). In addition, the updates are partial, so that only route changes are propagated versus the entire routing table and only sent to routers where the change affects routing decisions.

EIGRP can route IP, Internetwork Packet Exchange (IPX), and AppleTalk. It uses the DUAL algorithm for faster convergence. For information on DUAL concepts and data structures, see http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080093f07.shtml#eigrp_concepts.

EIGRP uses multicast to send updates by sending messages to 224.0.0.10, which enables the message/update to be sent to all EIGRP speakers in the domain.

Configuring EIGRP

Steps to configure EIGRP are as follows

Step 1. Enable the EIGRP routing process using the **router eigrp** *autonomous-system-number* command in global configuration mode.

Step 2. Configure networks to be enabled for EIGRP routing using the command **network** *network* command.

Step 3. Disable automatic summarization using the command **no auto-summary**.

For more information about EIGRP and its configuration, see the EIGRP design guide at http://www.cisco.com/warp/public/103/eigrp-toc.html (recommended).

Authentication is configured on EIGRP similar to RIPv2: by configuring the authentication modes on the interface and associating an authentication key chain instance (see Table 1-5).

Table 1-5 Configuring EIGRP Authentication

Command	Function
ip authentication key-chain eigrp *autonomous-system name-of-chain*	Associates an EIGRP autonomous system and key chain per interface in the interface configuration mode
ip authentication mode eigrp *autonomous-system* **md5**	Configures authentication mode as MD5 on the interface in the interface configuration mode

In addition, if used, the key chain will have to be configured as previously defined in the section "Routing Information Protocol (and Routing Information Protocol Version 2)."

Border Gateway Protocol

Border Gateway Protocol (BGP) is an exterior gateway protocol that's used for interdomain routing within the Internet today. BGP is considered a *path-vector protocol* because routing information exchange also propagates information on the path of autonomous systems via which the route was learned. BGP uses TCP port 179 (transport layer) for information exchange. In addition, BGP maintains a BGP table that contains information about all available paths to reach a specific destination. Only the best path is imported into the routing table. For complete coverage of BGP attributes and their operation in route selection, go to http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm. In addition, it is recommended that you visit the BGP Cisco FAQ (http://www.cisco.com/warp/public/459/bgpfaq_5816.shtml).

Configuring BGP (Basics Only)

Steps to configure BGP are as follows

Step 1. Enable BGP on the router using the command **router bgp** *autonomous-system-number*.

Step 2. Configure explicit neighbors using the **neighbor** *ip-address* **remote-as** *remote-as-number* command.

Step 3. (Optional) Configure networks to be advertised into the BGP process using the **network** *network-number* **mask** *subnet-mask* command.

Step 4. For interior BGP (iBGP) sessions, change the source of BGP updates to a specific interface using the command **neighbor** *ip-address* **update-source** *interface-type interface-number*.

For further configurations and in-depth coverage of BGP, see the Cisco BGP case studies at http://www.cisco.com/warp/public/459/bgp-toc.html (recommended).

Authentication (MD5) can be enabled per neighbor using the command **neighbor** *ip-address* **password** *string*.

IP Multicast Overview

Multicast is a subset of broadcast wherein just a specific subset of hosts receive the packet (versus all hosts on a broadcast domain). The host chooses whether to become a member of a certain multicast group address, thus enabling the host to receive packets destined for that group. Multicast addresses are Class D addresses ranging from 224.0.0.0 to 239.255.255.255. A large number of multicast protocols are in use today. You can find detailed coverage about these protocols at http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/ipmulti.htm and http://www.cisco.com/univercd/cc/td/doc/cisintwk/intsolns/mcst_sol/mcst_ovr.htm. In addition, special multicast addresses from the 224.0.0.0/8 block are used by control protocols (for example, 224.0.0.1 [all systems] and 224.0.0.2 [all routers on a subnet]).

Questions and Answers

1. IP routing occurs at what layer of the OSI model?

   a. Network layer
   b. Data link layer
   c. Transport layer
   d. Application layer

2. IP RIP runs over ___, port number ___.

   a. UDP, 21
   b. TCP, 24
   c. TCP, 520
   d. UDP, 520

 Please see page 131 for more details.

3. In what field(s) does the IP checksum calculate the checksum value?
   a. Data payload only
   b. Header and data payload
   c. Header only
   d. Not used in an IP packet

4. Which of the following routing protocols support authentication mechanisms? (Choose all that apply.)
   a. OSPFv2
   b. BGP
   c. RIPv1
   d. EIGRP
   e. IGRP

5. The default value for HSRP priority is ___.
   a. 100
   b. 110
   c. 150

6. The default values for BGP local preference, MED, and weight are ___.
   a. 100, 100, 100
   b. 100, 32,768, 100
   c. 32,768, 100, 32,768
   d. 100, 32,768, 0

7. The number of unique multicast IP addresses that map to a single Layer 2 multicast address is ___.

   a. 16
   b. 8
   c. 4
   d. 32
   e. 64

8. The process of configuring a multicast sparse mode network to provide for fault tolerance and load sharing within a single multicast domain is called _____.

   a. Source-based trees
   b. Shared trees
   c. Anycast RP
   d. MBGP