GLOSSARY

A

AAA — authentication, authorization, and accounting (pronounced “triple a”).

AAL5 — See ATM adaptation layer 5.

ABR — See available bit rate.

ADSL — See asymmetric digital subscriber line.

anonymous link — A link that does not provide a PPP username or endpoint. Multilink PPP (MLP) combines all of the anonymous links into a single bundle called an anonymous bundle or default bundle.

asymmetric digital subscriber line — A digital subscriber line (DSL) technology in which the transmission of data from server to client is much faster than the transmission from the client to the server.

Asynchronous Transfer Mode — International standard for cell relay in which multiple service types (such as voice, video, or data) are conveyed in fixed-length cells. Fixed-length cells allow cell processing to occur in hardware, thereby reducing transit delays. ATM is designed to take advantage of high-speed transmission media such as E3, SONET, and T3.

ATM adaptation layer 5 — This layer maps higher layer user data into ATM cells, making the data suitable for transport through the ATM network.

available bit rate — QoS class defined by the ATM Forum for ATM networks. ABR is used for connections that do not require timing relationships between source and destination. ABR provides no guarantees in terms of cell loss or delay, providing only best-effort service. Traffic sources adjust their transmission rates in response to information they receive describing the status of the network and its capability to successfully deliver data.

ATM — See Asynchronous Transfer Mode.

B

bandwidth — The range of frequencies a transmission line or channel can carry. The greater the bandwidth, the greater the information-carrying capacity of a channel. For a digital channel this is defined in bits. For an analog channel it is dependent on the type and method of modulation used to encode the data.

BGP — See Border Gateway Protocol.

Border Gateway Protocol — Interdomain routing protocol that exchanges reachability information with other BGP systems. It is defined in RFC 1163.

bps Bits per second. A standard measurement of digital transmission speeds.

bundle A logical group of permanent virtual circuits (PVCs) with one virtual interface connecting to a peer system.

C

CAR Committed access rate.

CBR See constant bit rate.

CBWFQ See Class-based Weighted Fair Queuing.

CEF See Cisco Express Forwarding.

child policy A policy map that defines one or more classes of traffic and the actions you want the router to take on the traffic, just as non-hierarchical policy maps do. However, in hierarchical policy maps, child policies are nested within a top-level parent policy, and then the parent policy is attached to the interface.

CIR See committed information rate.

Cisco express forwarding An advanced Layer 3 IP switching technology. Cisco express forwarding (CEF) optimizes network performance and scalability for networks with large and dynamic traffic patterns such as the Internet, on networks characterized by intensive Web-based applications, or interactive sessions.

class-based weighted fair queuing Extends the standard weighted fair queuing (WFQ) functionality to provide support for user-defined traffic classes. For class-based weighted fair queuing (CBWFQ), you define traffic classes based on match criteria including protocols, access control lists (ACLs), and input interfaces. Packets satisfying the match criteria for a class constitute the traffic for that class. A queue is reserved for each class and traffic belonging to a class is directed to the queue for that class. On the Cisco 10000 series router, the CBWFQ feature allows a virtual access interface (VAI) to inherit the service policy of the virtual circuit (VC) that the VAI uses.

class maps A modular QoS CLI element that you can use to define traffic classification rules or criteria. Class maps organize data packets into specific categories called classes that can, in turn, receive user-defined QoS policies. The traffic class defines the classification rules for packets received on an interface.

class of service The three most significant bits (the User Priority bits) of the 2-byte Tag Control Information field in the IEEE 802.1p portion of a Layer 2 IEEE 802.1Q frame header. QoS uses the User Priority bits for Layer 2 CoS information. IEEE 802.1p class of service-based packet matching and marking feature enables the Cisco 10000 series router to interoperate with switches to deliver end-to-end QoS. The IEEE 802.1p standard allows QoS to classify inbound Ethernet packets based on the value in the class of service (CoS) field and to explicitly set the value in the CoS field of outbound packets.

committed information rate The reserved bandwidth for the queue. The rate at which a Frame Relay network agrees to transfer information under normal conditions, averaged over a minimum increment of time. Committed information rate (CIR), measured in bits per second, is one of the key negotiated tariff metrics.

constant bit rate QoS class defined by the ATM Forum for ATM networks. Constant bit rate (CBR) is used for connections that depend on precise clocking to ensure undistorted delivery.

CoS See class of service.

D

data-link connection identifier A value that specifies a PVC in a Frame Relay network. In the basic Frame Relay specification, DLCIs are locally significant (connected devices might use different values to specify the same connection). In the LMI extended specification, DLCIs are globally significant (DLCIs specify individual end devices).

DLCI See data-link connection identifier.

downstream rate The line rate for return messages or data transfers from the network machine to the user's customer premises machine.

DSCP Differentiated service code point.

DSL Digital Subscriber Line.

E

encapsulation The technique used by layered protocols in which a layer adds header information to the protocol data unit (PDU) from the layer above.

endpoint discriminator A value a system uses when negotiating the use of Multilink PPP (MLP) with the peer system. The default value is the username that is used for authentication.

Ethernet One of the most common local area network (LAN) wiring schemes, Ethernet has a transmission rate of 10, 100, or 1000 Mbps.

F

fragmentation See link fragmentation and interleaving.

G

GE Gigabit Ethernet.

GRE Generic Route Encapsulation. A method of encapsulating any network protocol in another protocol.

H

hierarchical policy A QoS policy in which multiple policies are configured into a single QoS policy. The hierarchical policy combines one or more classes to apply specific actions on the aggregate traffic and to execute class-specific actions. The non-hierarchical policy, on the other hand, defines only class-specific actions. Hierarchical policies can have up to three levels configured.

hierarchical input policing A QoS model that defines policing policies at multiple levels of hierarchy for inbound packets. See also hierarchical policies.

HTML Hypertext Markup Language. The page-coding language for the World Wide Web.

http Hypertext Transfer Protocol. The protocol used to carry world-wide web (www) traffic between a www browser computer and the www server being accessed.

I

interleaving See link fragmentation and interleaving.

Internet A collection of networks interconnected by a set of routers that allow them to function as a single, large virtual network.

Internet Protocol The network layer protocol for the Internet protocol suite.

IP See Internet Protocol.

L

LFI See link fragmentation and interleaving.

link fragmentation and interleaving The ability to fragment network level datagrams (and possibly interleave them) at the link layer. Multilink inherently includes link fragmentation and interleaving capabilities, as do some other technologies such as ATM.

load balancing In routing, the capability of a router to distribute traffic over all its network ports that are the same distance from the destination address. Good load-balancing algorithms use both line speed and reliability information. Load balancing increases the use of network segments, thus increasing effective network bandwidth.

M

Management Information Base Database of network management information that is used and maintained by a network management protocol, such as SNMP or CMIP (Common Management Information Protocol). The value of a MIB object can be changed or retrieved using SNMP commands, usually through a Network Management System (NMS). MIB objects are organized in a tree structure that includes public (standard) and private (proprietary) branches.

marking A QoS tool used to differentiate packets based on designated markings. Using marking, you can partition your network into multiple priority levels or classes of service. Marking simplifies the network Qos design and QoS tools configuration, and reduces the overhead of packet classification by other QoS tools.

MIB See Management Information Base.

MLP See Multilink PPP.

MLP bundle A virtual interface that connects to a peer system. MLP bundles consist of multiple physical links.

MP See Multilink PPP.

Modular QoS command-line interface Also referred to as modular CLI. A platform-independent CLI for configuring QoS features on Cisco products.

MPLS See Multiprotocol Label Switching.

MPLS VPN MPLS-based virtual private network.

MQC See modular QoS command-line interface.

multicast Single packets copied by the network and sent to a specific subset of network addresses. These addresses are specified in the Destination Address Field.

Multilink PPP Multilink Point-to-Point Protocol, RFC 1990. Commonly abbreviated as MLP within Cisco; however, the correct abbreviation is MP. This protocol is a method of splitting, recombining, and sequencing datagrams across multiple logical data links.

multipoint subinterface Multipoint networks have three or more routers in the same subnet. For Dynamic Bandwidth Selection, if you put the PVC in a point-to-multipoint subinterface or in the main interface (which is multipoint by default), you need to either configure a static mapping or enable inverse Address Resolution Protocol (ARP) for dynamic mapping.

Multiprotocol Label Switching Switching method that forwards IP traffic using a label. This label instructs the routers and the switches in the network where to forward the packets based on preestablished IP routing information.

N

NAS See network access server.

nested hierarchical policy A QoS model that defines policies at two levels of hierarchy. See also hierarchical policies.

NetFlow A Cisco-proprietary IP statistics collection feature that collects information on IP flows passing through a router.

network access server Cisco platform (or collection of platforms) that interfaces between the packet world (for example, the Internet) and the circuit world (for example, PSTN).

NVRAM See Non-Volatile Random Access Memory.

Non-Volatile Random Access Memory The router uses this memory to store configuration information. The contents of this memory are not lost after a reboot or power cycle of the unit.

O

oversubscription A method of improving network utilization by assigning a committed rate that is greater than the speed of an interface or subinterface, or greater than the capacity of a VC port.

P

parent policy A QoS policy that defines the shape rate for the child traffic on an interface. The parent policy contains only the class-default class; it can contain no other classes.

PCR See peak cell rate.

peak cell rate Parameter defined by the ATM Forum for ATM traffic management.

permanent virtual circuit Virtual circuit that is permanently established. PVCs save bandwidth associated with circuit establishment and tear down in situations where certain virtual circuits must exist all the time. In ATM terminology, called a permanent virtual connection. See also virtual circuit (VC).

permanent virtual path A virtual path that consists of PVCs.

Point-to-Point Protocol The successor to SLIP, Point-to-Point Protocol (PPP) provides router-to-router and host-to-network connections over both synchronous and asynchronous circuits.

point-to-point subinterface With point-to-point subinterfaces, each pair of routers has its own subnet. If you put the PVC on a point-to-point subinterface, the router assumes that there is only one point-to-point PVC configured on the subinterface. Therefore, any IP packets with a destination IP address in the same subnet are forwarded on this VC. This is the simplest way to configure the mapping and is, therefore, the recommended method.

policing A traffic regulation mechanism that is used to limit the rate of traffic streams. Policing allows you to control the maximum rate of traffic sent or received on an interface. Policing propagates bursts of traffic and is applied to the inbound or outbound traffic on an interface. When the traffic rate exceeds the configured maximum rate, policing drops or remarks the excess traffic.

policy map A modular CLI (MQC) element that enables you to create QoS policies to tell the router the QoS actions and rules to apply to packets belonging to a particular traffic class.

PPP See Point-to-Point Protocol.

PPPoA PPP over ATM. Enables a high-capacity central site router with an Asynchronous Transfer Mode (ATM) interface to terminate multiple remote PPP connections.

PPPoE PPP over Ethernet. Allows a PPP session to be initiated on a simple bridging Ethernet connected client. Refers to a signaling protocol defined within PPPoE as well as the encapsulation method. See also RFC 2516.

PPPoEoA PPP over Ethernet over ATM. Allows tunneling and termination of PPP sessions over Ethernet links and allows for Ethernet PPP connections over ATM links.

PPPoEoE PPP over Ethernet over on Ethernet. Allows tunneling and termination of PPP sessions over Ethernet links and allows for Ethernet PPP connections over Ethernet links.

PPPoEo802.1Q VLAN PPP over Ethernet over IEEE 802.1Q VLANs. Allows tunneling and termination of Ethernet PPP sessions across VLAN links. IEEE 802.1Q encapsulation is used to interconnect a VLAN-capable router with another VLAN-capable networking device. The packets on the 802.1Q link contain a standard Ethernet frame and the VLAN information associated with that frame.

PPPoX PPP over PPPoA or PPPoE or both.

PQ See priority queuing.

priority queuing A class queue that is given priority over other queues. The router dequeues and transmits priority queue data before it dequeues and sends packets in other queues. Using priority queuing, delay-sensitive data such as voice receives preferential treatment over other traffic.

PTA PPP termination aggregation. A method of aggregating IP traffic by terminating PPP sessions and aggregating the IP traffic into a single routing domain.

PTA-MD PTA Multi-Domain. A method of aggregating IP traffic by terminating PPP sessions and aggregating the IP traffic into a VPN or multiple IP routing domains.

PVC See permanent virtual circuit or connection.

PVP See permanent virtual path.

Q

QoS See quality of service.

QPPB QoS policy propagation through the Border Gateway Protocol.

quality of service Quality of service (QoS) refers to the ability of a network to provide improved service to selected network traffic over various underlying technologies including Frame Relay, ATM, Ethernet and 802.1 networks, SONET, and IP-routed networks. Cisco IOS QoS technology lets complex networks control and predictably service a variety of networked applications and traffic types.

queue depth A user-configurable limit on the number of packets that the router can place onto a queue. During periods of high traffic, a queue fills with packets waiting for transmission. When a queue reaches its queue depth and becomes full, by default the router drops packets until the congestion is eliminated and the queue is no longer full.

queuing delay Amount of time that data must wait before it can be transmitted onto a statistically multiplexed physical circuit.

R

RADIUS Remote Authentication Dial-In User Service (RADIUS). A client/server security protocol created by Livingston Enterprises. Security information is stored in a central location, known as the RADIUS server.

random early detection An alternative mechanism for avoiding congestion of Layer 3 queues. Random early detection (RED) implements a proactive queuing strategy that manages congestion before a queue reaches its queue depth or limit.

RIP Routing Information Protocol. An IGP used to exchange routing information within an autonomous system, RIP uses hop count as a routing metric.

route The path that network traffic takes from its source to its destination. The route a datagram follows can include many gateways and many physical networks. In the Internet, each datagram is routed separately.

router A system responsible for making decisions about which of several paths network (or Internet) traffic will follow. To do this, it uses a routing protocol to gain information about the network and algorithms to choose the best route based on several criteria known as "routing metrics."

routing table Information stored within a router that contains network path and status information. It is used to select the most appropriate route to forward information along.

S

SCR See sustainable cell rate.

sustainable cell rate Parameter defined by the ATM Forum for ATM traffic management. For VBR connnections, SCR determines the long-term average cell rate that can be transmitted.

SVC See switched virtual circuit.

switched virtual circuit A virtual circuit that is dynamically established on demand and is torn down when transmission is complete. SVCs are used in situations where data transmission is sporadic. Called a switched virtual connection in ATM terminology.

T

tail drop The default congestion avoidance mechanism for Layer 3 queues. Tail drop activates when a queue becomes full. After being activated, no packets make it to the queue. Tail drop treats all traffic equally and does not differentiate between classes of service.

three-level hierarchical policy A QoS model that defines policies at three levels of hierarchy. See also hierarchical policies.

token bucket An algorithm used o manage the maximum rate of traffic. This algorithm defines the maximum rate of traffic allowed on an interface at a given moment in time. The token bucket algorithm is especially useful in managing network bandwidth in cases where several large packets are sent in the same traffic stream. The algorithm puts tokens into the bucket at a certain rate. Each token is permission for the source to send a specific number of bits into the network.

ToS Type of service. First defined in RFC 791.

U

UBR See unspecified bit rate.

unspecified bit rate A QoS class defined by the ATM Forum for ATM networks. UBR allows any amount of data up to a specified maximum to be sent across the network, but there are not guarantees in terms of cell loss rate and delay.

upstream rate The line rate for message or data transfer from the source machine to a destination machine on the network.

V

VAI See virtual access interface.

VBR See variable bit rate.

VBR-nrt See variable bit rate-nonreal time.

VBR-rt See variable bit rate-real time.

VC See virtual circuit.

VCI See virtual channel identifier.

VCL See virtual channel link.

variable bit rate A QoS class defined by the ATM Forum for ATM networks. Variable bit rate (VBR) is subdivided into a real time (rt) class and nonreal time (nrt) class. See also variable bit rate-nonreal time and variable bit rate-real time.

variable bit rate-nonreal time A QoS class defined by the ATM Forum for ATM networks. Variable bit rate-nonreal time (VBR-nrt) is used for connections in which there is no fixed timing relationship between samples, but that still need a guaranteed QoS.

variable bit rate-real time A QoS class defined by the ATM Forum for ATM networks. Variable bit rate-real time (VBR-rt) is used for connections in which there is a fixed timing relationship between samples.

virtual access interface An instance of a unique virtual interface that is created dynamically and exists temporarily. Virtual access interfaces can be created and configured differently by different applications, such as virtual profiles and virtual private dialup networks (VPDNs).Virtual access interfaces are cloned from virtual template interfaces.

virtual channel identifier A 16-bit field in the header of an ATM cell. The virtual channel identifier (VCI), together with the VPI, is used to identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination. ATM switches use the VPI/VCI fields to identify the next network VCL that a cell needs to transmit on its way to its final destination. The function of the VCI is similar to that of the DLCI in Frame Relay.

virtual channel link Connection between two ATM devices. A virtual channel connection is made up of one or more VCLs.

virtual circuit Logical circuit created to ensure reliable communication between two network devices. A virtual circuit is defined by a VPI/VCI pair, and can be either permanent (PVC) or switched (SVC). Virtual circuits are used in Frame Relay and X.25. In ATM, a virtual circuit is called a virtual channel. Sometimes abbreviated VC.

virtual LAN A group of devices on one or more local area networks (LANs) that are configured (using management software) so that they can communicate as if they were attached to the same wire, when in fact they are located on a number of different LAN segments. Because virtual LANs (VLANs) are based on logical instead of physical connections, they are extremely flexible.

virtual path One of two types of ATM circuits identified by a VPI. A virtual path is a bundle of virtual channels, all of which are switched transparently across an ATM network based on a common VPI.

virtual path identifier An 8-bit field in the header of an ATM cell. The VPI, together with the VCI, is used to identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination. ATM switches use the VPI/VCI fields to identify the next VCL that a cell needs to transmit on its way to its final destination. The function of the VPI is similar to that of the DLCI in Frame Relay.

virtual template interface A logical interface configured with generic configuration information for a specific purpose or configuration common to specific users, plus router-dependent information. The template takes the form of a list of Cisco IOS interface commands that are applied to virtual access interfaces, as needed.

VLAN See virtual LAN.

VPI See virtual path identifier.

W

WAN See wide area network.

weighted fair queuing Congestion management algorithm that identifies conversations (in the form of traffic streams), separates packets that belong to each conversation, and ensures that capacity is shared fairly between these individual conversations. WFQ is an automatic way of stabilizing network behavior during congestion and results in increased performance and reduced retransmission.

weighted random early detection A mechanism for avoiding congestion of Layer 3 queues. Weighted random early detection (WRED) combines the capabilities of the random early detection (RED) mechanism with IP precedence, differential services code point (DSCP), and discard-class to provide preferential handling of higher priority packets. WRED attempts to anticipate and avoid congestion. WRED implements a proactive queuing strategy that manages congestion before a queue reaches its queue depth. By selectively dropping packets, WRED prevents packets from enqueuing to the Layer 3 queue.

WFQ See weighted fair queuing.

wide area network A data communications network that spans any distance and is usually provided by a public carrier (such as a telephone company or service provider).

WRED See Weighted Random Early Detection.

X

xDSL Various types of digital subscriber lines. Examples include ADSL, HDLS, and VDSL.