JUNOSe Glossary

Numerics

3DES Triple DES encryption/decryption algorithm. An algorithm that encrypts data blocks with three different keys in succession. Data is encrypted with the first key, decrypted with the second key, and encrypted again with the third key. 3DES is one of the strongest encryption algorithms available for use in virtual private networks (VPNs). 3DES is slower than standard DES but provides greater security. 3DES is often implemented with cipher block chaining (CBC). Also called *triple DES*. *See also* DES.

10-gigabit small form-factor pluggable transceiver *See* XFP.

802.3ad link aggregation A process that enables grouping of Ethernet interfaces at the physical layer to form a single link layer interface, also known as a link aggregation group (LAG) or LAG bundle.

A

AAA authentication, authorization, and accounting. Each has an important but separate function.

- Authentication—Determines who the user is, then determines whether to grant that user access to the network. The primary purpose is to prevent intruders from networks. It uses a database of users and passwords.
- Authorization—Determines what the user can do by giving you the ability to limit network services to different users.
- Accounting—Tracks what the user did and when the user did it. You can use accounting for an audit trail or for billing for connection time or resources used.

See also redirected authentication.

AAA profile A set of characteristics that act as a pattern that you can assign to domain names. After you create an AAA profile, you can map it between a PPP client's domain name and certain AAA services on given interfaces. Using AAA profiles, you can:

- Allow or deny a domain name access to AAA authentication
- Map the original domain name to the mapped domain name for domain name lookup
- Use domain name aliases
- Force tunneling whenever a domain map contains tunnel attributes

- Manually set the NAS-Port-Type attribute (RADIUS attribute 61) for ATM and Ethernet interfaces
- Set the Profile-Service-Description attribute (RADIUS attribute 26-53)

An AAA profile contains a set of commands to control access for the incoming PPP subscriber. If no AAA profile is used, AAA continues as normal. The user's name and domain name are not changed as a result of an AAA profile mapping

AAL ATM Adaptation Layer. A collection of protocols that defines the conversion of user information into cells by segmenting upper-layer information into cells at the transmitter and reassembling them at the receiver. These protocols enable various types of traffic, including voice, data, image, and video, to run over an ATM network.

ABR
- area border router. A router that has interfaces in the OSPF boundary between two or more different areas. Both sides of any link always belong to the same OSPF area.
- available bit rate. A rate that is used in ATM for traffic sources that demand low loss ratios but can accept larger delays. ABR uses bandwidth that constant bit rate (CBR) and variable bit rate (VBR) does not use. ABR uses best effort to send the maximum number of cells but not guarantee cell delivery. *See also* CBR; VBR.

AC access concentrator. A device that receives and forwards data for a network point of presence (POP). It often acts as a server that supports multiple T1 or E1 lines over one port.

access lists A sequential collection of permit and deny conditions that you can use to filter inbound or outbound routes.

Files that provide filters that can be applied to route maps or distribution lists. They enable policies to be created, such as a policy to prevent forwarding of specified routes between the BGP-4 and IS-IS routing tables.

access messages Authorization and authentication (AA) messages that identify subscribers before the RADIUS server grants or denies them access to the network or network services. When an application requests user authentication, the request must have certain authenticating attributes, such as a user's name, password, and the particular type of service the user is requesting. This information is sent in the authentication request via the RADIUS protocol to the RADIUS server. In response, the RADIUS server grants or denies the request. *See also* accounting messages.

ACCM Async Control Character Map. An option negotiated by LCP that is used on asynchronous links, such as telephone lines, to identify control characters that must be escaped (replaced by a specific two-character sequence) to avoid them being interpreted by equipment used to establish the link.

accounting In RADIUS, the process of tracking what the user did and when the user did it. You can use accounting for an audit trail or for billing for connection time or resources used. *See also* broadcast accounting server; duplicate accounting server.

accounting messages AA messages that identify service provisions and use on a per-user basis. They keep track of when a particular service is initiated and terminated for a specific user. RADIUS attributes are used by each group of accounting messages. *See also* access messages.

ACFC Address and Control Field Compression. A compression method that enables routers to transmit packets without the two 1-byte address and control fields (0xff and 0x03) normal for PPP-encapsulated packets, thus transmitting less data and conserving bandwidth. ACFC is defined in RFC 1661, *The Point-to-Point Protocol (PPP)*. *See also* PFC.

active constituent A constituent that is monitored or controlled by the shared shaper mechanism. *See also* constituent; inactive constituent.

active state A state of an SRP module whereby data that was synchronized from the active SRP module to the standby SRP module during initialization remains synchronized through mirroring updates.

address pool In a NAT context, a group of IP addresses from which a NAT router obtains an address when dynamically creating a new translation.

Address Resolution Protocol *See* ARP.

address scope *See* scope.

adjacency The relationship between a pair of selected neighboring routers for exchanging routing information. Not every pair of neighboring routers is adjacent. A given router can have multiple adjacencies, but each adjacency consists of only two routers connected by one media segment. Packets that go between them do not have to pass through any other network devices. *See also* neighbor.

admission control An accounting mechanism that tracks resource information on a router-wide basis. Prevents requests from being accepted when sufficient resources are not available. Admission control determines whether a setup request can be honored for an MPLS LSP with traffic parameters.

administrative distance An integer (in the range 0–255) that is associated with each route known to a router. The distance represents how reliable the source of the route is considered to be. A lower value is preferred over a higher value. An administrative distance of 255 indicates no confidence in the source; routes with this distance are not installed in the routing table.

ADM Add/drop multiplexer. SONET functionality that enables lower-level signals to be dropped from a high-speed optical connection.

ADSL asymmetric digital subscriber line. A technology that enables more data to be sent over existing copper telephone lines, using the public switched telephone network (PSTN). ADSL supports data rates from 1.5 through 9 Mbps when receiving data (downstream rate) and from 16 through 640 Kbps when sending data (upstream rate).

AF assured forwarding. A DiffServ component that determines the degree of reliability given a packet within the DiffServ domain. AF values are set as part of Per-Hop Behavior (PHB) groups. *See also* PHB.

AFI
- authority and format identifier. A number that identifies the format and type of address being used.
- address family identifier. A number assigned by IANA used to identify the protocol associated with an address family. In an MP-BGP update message, AFI is used with SAFI to identify the network layer protocol associated with the network address of the next hop and the semantics of the NLRI that follows. *See also* SAFI.

agent *See* SNMP agent.

aggregation The process of accumulating data or logical interfaces into a single. larger, bundle (for example, higher-speed connections).

aggressive mode An IKE phase 1 negotiation mode that:
- Exposes identities of the peers to eavesdropping, making it less secure than main mode.
- Is faster than main mode because fewer messages are exchanged between peers. (Three messages are exchanged in aggressive mode.)
- Enables support for fully qualified domain names (FQDNs) when the router uses preshared keys

See also main mode.

AH authentication header. A part of an IP datagram that provides authentication of the sender and of data integrity.

AIS alarm indication signal. A signal transmitted instead of the normal signal to maintain transmission continuity and to indicate to the receiving equipment that a transmission interruption (fault) has occurred either at the equipment originating the AIS signal or upstream of that equipment.

AIS cell alarm indication signal cell. A type of ATM cell used to indicate a fault to the downstream endpoint.

alarm indication signal *See* AIS.

ALG application-level gateway. A security component used in a firewall or Network Address Translation (NAT). ALGs enable certain legitimate applications to pass through a firewall or between NAT realms without being stopped by security checks.

analyzer device A device that receives mirrored traffic from E-series routers during packet mirroring. Also called the mediation device.

analyzer port The IP interface in analyzer mode on E-series routers that is used to direct mirrored traffic to the analyzer device during packet mirroring.

ANCP Access Node Control Protocol. A protocol that is based on a subset of the General Switch Management Protocol (GSMP). With ANCP, IGMP is no longer terminated or proxied at the access node. Instead, IGMP passes through the access node transparently. Also known as layer 2 control (L2C).

ANSI American National Standards Institute. Private organization that coordinates the development and use of voluntary consensus standards in the United States. Representative for the United States to ISO. *See also*ISO.

anycast address A type of IPv6 address, used in IPv6; used to send a packet to one recipient out of a set of recipients.

Used for a set of interfaces on different nodes. An anycast transmission sends packets to only one of the interfaces associated with the address, not to all of the interfaces. This interface is typically the closest interface, as defined by the routing protocol.

API application programming interface. A set of routines, protocols, and tools for building software applications.

application layer

- The seventh and highest level in the seven-layer OSI reference model for network protocol design that manages communication between application processes. This layer is the main interface for users to interact with application programs such as electronic mail, database managers, and file-server software. *See also* OSI.
- The fifth and highest level in the five-layer TCP/IP protocol stack. This layer is used by most programs for network communication. Data is passed from the program in an application-specific format, then encapsulated into a transport layer protocol.

application-level gateway *See* ALG.

application programming interface *See* API.

application-specific integrated circuit *See* ASIC.

APS Automatic Protection Switching. Technology used by SONET ADMs to protect against circuit faults between the ADM and a router and to protect against failing routers. *See also* ADM.

area

- For IS-IS, a set of contiguous networks and hosts within an autonomous system that have been administratively grouped together.
- For OSPF, a collection of network segments interconnected by routers. A region in an OSPF routing domain. Also a unique number that identifies an area. Typically, formatted as an IP address.

area border router *See* ABR.

ARP Address Resolution Protocol. A protocol used to map a MAC address to an IP address. Dynamically binds the IP address (the logical address) to the correct MAC address.

AS autonomous system.

- A set of routers that use the same routing policy while running under a single technical administration. An AS runs interior gateway protocols (IGPs) such as RIP, OSPF, and IS-IS within its boundaries. ASs use exterior gateway protocols (EGPs) to exchange routing information with other ASs.
- A routing domain. Assigned a globally unique number called an Autonomous System Number (ASN).

AS boundary router autonomous system boundary router. An OSPF router that exchanges routing information with routers in other ASs. The AS boundary router redistributes routing information received from other ASs throughout its own AS.

ASCII American Standard Code for Information Interchange. A code for representing English characters as numbers, with each letter assigned a number in the range 0–127.

ASIC application-specific integrated circuit. Specialized processors that perform specific functions on the router.

AS number autonomous system number. A globally unique number assigned by the IANA that is used to identify an Autonomous System (AS). The AS number enables an AS to exchange exterior routing information with neighboring ASs.

assured forwarding *See* AF.

assured rate A rate when bandwidth is guaranteed until oversubscribed (JUNOSe QoS term)

asymmetric digital subscriber line *See* ADSL.

Asynchronous Transfer Mode *See* ATM.

Async Control Character Map *See* ACCM.

ATM Asynchronous Transfer Mode. A high-speed multiplexing and switching method utilizing fixed-length cells of 53 octets to support multiple types of traffic.

ATM Adaptation Layer *See* AAL.

ATM cell A package of information that is always 53 octets long, unlike a frame or packet, which has a variable length.

ATM subinterface A mechanism that enables a single physical ATM interface to support multiple logical interfaces.

attribute-value pair *See* AVP.

autodetection A process that determines the layers of each dynamic interface. The autodetection process occurs when the router conditionally constructs interface layers based on the encapsulation type of the incoming packet. Also called autosensing.

authentication

- In RADIUS, the process of determining who the user is, and then determining whether to grant that user access to the network. The primary purpose is to prevent intruders from networks. RADIUS authentication uses a database of users and passwords.
- A process that verifies that data is not altered during transmission and ensures that users are communicating with the individual or organization that they believe they are communicating with. *See also* IPSec.
- An SNMPv3 term related to the user-based security model (USM). Authentication provides the following benefits:
  - Only authorized parties can communicate with each other. Consequently, a management station can interact with a device only if the administrator configured the device to allow the interaction.
  - Messages are received promptly; users cannot save messages and replay them to alter content. This feature prevents users from sabotaging SNMP configurations and operations. For example, users can change configurations of network devices only if authorized to do so.

authentication, authorization, and accounting *See* AAA.

authentication header *See* AH.

authentication retry A feature of SSH that limits the number of times a user can try to correct incorrect information—such as a bad password—in a given connection attempt.

authority and format identifier *See* AFI.

authorization In RADIUS, the process of determining what the user can do by giving a network administrator the ability to limit network services to different users.

Automatic Commit mode A feature of JUNOSe software where the system automatically saves any change to the system configuration to NVS, without affecting the CLI prompt.

Automatic Protection Switching *See* APS.

autonomous system *See* AS.

AVP attribute value pair. A RADIUS attribute value carried in a RADIUS protocol message. The pair is a combination of a unique attribute—represented by an integer—and a value containing the actual value identified by the attribute.

B

backbone network A central network; a network that connects other networks together.

backdoor link A private link between two routers.

OSPF backdoor links typically serve as backup paths, providing a way for traffic to flow from one VPN site to the other only if the path over the backbone is broken.

However, when the OSPF backdoor link connects two sites that are in the same OSPF area, the undesired result is that the path over the OSPF backdoor link is always preferred over the path over the backbone.

backup DR backup designated router. An OSPF router on a broadcast segment that monitors the operation of the designated router and takes over its functions if the designated router fails. *See also* DR.

backup router The VRRP router available to take forwarding responsibility if the current master router fails. *See also* master router.

backward explicit congestion notification *See* BECN.

bandwidth management The part of policy management that rate-limits a classified packet flow at ingress to enforce ingress data rates below the physical line rate of a port. When the user configures a rate-limit profile, packets are tagged with a drop preference.

bandwidth oversubscription A feature of JUNOSe software that enables line modules to operate at a rate dependent on the resources available rather than having all line modules operate at full line rate performance. Oversubscription enables a much more extensive combination of line modules in the router. *See also* oversubscription.

Base64 A method used to encode digital certificate requests and certificates before they are sent to or from the certificate authority (CA)

baseline statistics *See* statistics baseline.

basic NAT The least secure type of traditional network address translation (NAT). Provides translation for IP addresses only and places the mapping into a NAT table. *See also* NAT.

B-channel Bearer channel. A 64 Kbps channel used for voice or data transfer on an ISDN interface. *See also* D-channel.

BECN backward explicit congestion notification. In a Frame Relay network, a header bit transmitted by the destination device requesting that the source device send data more slowly. BECN minimizes the possibility that packets will be discarded when more packets arrive than can be handled. *See also* FECN.

BER bit error rate. The percentage of received bits in error compared to the total number of bits received.

BERT bit error rate test. A test performed by sending a known pattern of bits and counting the number of bit errors received to measure the quality of data transmission. Can be run on E1, E3, T1, T3, and channelized (DS3, OC3, OC12, and STM1) interfaces to determine whether they are operating properly.

best effort The default traffic class for packets being forwarded across the device. Packets that are not assigned to a specific traffic class are assigned to the best-effort traffic class. Network forwards as many packets as possible in as reasonable a time as possible.

best-effort queue For a logical interface, the queue associated with the best-effort traffic class for that logical interface.

best-effort scheduler node The scheduler node associated with a logical interface and traffic class group pair, and where the traffic class group contains the best-effort traffic class. Also known as best-effort node.

best path The one route that BGP selects to a destination. When multiple routes to a given destination exist, BGP must determine which of these routes is the best. BGP puts the best path in its routing table and advertises that path to its BGP neighbors.

If only one route exists to a particular destination, BGP installs that route. If multiple routes exist for a destination, BGP uses tie-breaking rules to decide which one of the routes to install in the BGP routing table.

BFD Bidirectional Forwarding Detection. A protocol that uses control packets and shorter detection time limits to more rapidly detect failures in a network.

BGP Border Gateway Protocol. The exterior gateway protocol (EGP) used to exchange routing information among routers in different autonomous systems. Can act as a label distribution protocol for MPLS.

BGP messages Routing information that BGP speakers exchange with each other over a BGP session. BGP uses five message types:

- Open BGP messages—Messages used to establish and negotiate certain parameters for the BGP session after the underlying TCP session has been established.
- Update messages—Messages used to announce routes to prefixes that the speaker can reach and to withdraw routes to prefixes that it can no longer reach. The most important message in the BGP protocol.
- Keepalive messages—Periodic messages to determine whether the underlying TCP connection is still up.
- Notification messages—Messages sent to a BGP peer to terminate a BGP session (either because the speaker has been configured to do so or because it has detected some error condition).
- Route-refresh messages—Messages sent to BGP peers that advertise their route-refresh capability, which enables the BGP speaker to apply modified or new policies to the refreshed routes.

BGP peer A BGP neighbor that has been explicitly configured for a BGP speaker. BGP peers do not have to be directly connected to each other in order to share a BGP session.

BGP peer group Two or more BGP peers that share a common set of update policies. They are grouped together to reduce configuration overhead and to conserve system resources when updates are generated.

BGP route A prefix and a set of path attributes. Sometimes referred to as a path, although that term technically refers to one of the path attributes of that route.

BGP session A TCP connection over which routing information is exchanged according to the rules of the BGP protocol. When two BGP speakers are in the same autonomous system, the BGP session is called an *internal* BGP session, or IBGP session. When two BGP speakers are in different autonomous systems, the BGP session is called an *external* BGP session, or EBGP session. BGP uses the same types of message on IBGP and EBGP sessions, but the rules for when to send which message and how to interpret each message differ slightly. *See also* IBGP session; EBGP session

BGP speaker A router that has been configured to run the BGP routing protocol. Unlike some other routing protocols, BGP speakers do not automatically discover each other and begin exchanging information. Instead, each BGP speaker must be explicitly configured with a set of BGP peers with which it exchanges routing information.

Bidirectional Forwarding Detection *See* BFD.

bidirectional NAT A type of network address translation (NAT). Adds support for the DNS to basic NAT to allow public hosts to initiate sessions into the private network, usually to reach servers intended for public access.

bit error rate *See* BER.

bit error rate test *See* BERT.

BMA broadcast multiaccess. A network on which broadcast or multicast packets can be sent, enabling each device on a network segment to communicate directly with every other device on that segment. *See also* NBMA.

BOOTP Bootstrap protocol. A UDP/IP-based protocol that allows a booting host to configure itself dynamically and without user supervision. BOOTP provides a means to notify a host of its assigned IP address, the IP address of a boot server host, and the name of a file to be loaded into memory and executed. Other configuration information, such as the local subnet mask, the local time offset, the addresses of default routers, and the addresses of various Internet servers, can also be communicated to a host using BOOTP.

bootstrap protocol *See* BOOTP.

Border Gateway Protocol *See* BGP.

B-RAS Broadband Remote Access Server. An application that is responsible for aggregating the output from digital subscriber line access multiplexers (DSLAMs), providing user PPP sessions and PPP session termination, enforcing QoS policies, and routing traffic into an ISP's backbone network

bridged Ethernet interface A link layer protocol that allows multiple upper-layer interface types (IP, PPPoE, and CBF) to be simultaneously multiplexed over the same interface.

bridged IP A link layer protocol used to manage IP packets that are encapsulated inside and Ethernet frame running over a permanent virtual circuit (PVC)

bridge group A collection of bridge interfaces stacked on Ethernet layer 2 network interfaces (ports) to form a broadcast domain. Each bridge group has its own set of forwarding tables and filters and, as such, functions as a logical transparent bridging device.

bridge group interface An association of one or more network interfaces with a bridge group. Also called a bridge interface.

Broadband Remote Access Server *See* B-RAS.

broadcast accounting server In RADIUS, a server that sends the accounting information to a group of virtual routers. An accounting virtual router group can contain up to four virtual routers and the E-series router supports a maximum of 100 virtual router groups. The accounting information continues to be sent to the duplicate accounting virtual router, if one is configured. You might use broadcast accounting to send accounting information to a group of your private accounting servers. *See also* duplicate accounting server.

broadcast address An IPv4 type of address that enables a device to send a packet to all hosts on a subnetwork.

broadcast circuits Circuits that use designated routers and are represented as virtual nodes in the network topology. They require periodic database synchronization. By default, IS-IS treats the broadcast link as LAN media and tries to bring up the LAN adjacency even when the interface is configured as unnumbered or only a single neighbor exists on that link. *See also* point-to-point circuits.

broadcast multiaccess *See* BMA.

bulk configuration A process in which you configure a range of ATM permanent virtual circuits (PVCs) to support dynamic interfaces.

bundle *See* 802.3ad link aggregation.

bypass tunnel A single LSP used to back up a set of LSPs by bypassing specific links in the LSP. In the event of a failure in any link of the protected RSVP-TE LSP (the primary LSP), MPLS redirects traffic to the associated bypass tunnel in tens of milliseconds.

C

CA certificate authority. A trusted third-party organization that creates, enrolls, validates, and revokes digital certificates. The CA guarantees a user's identity and issues public and private keys for message encryption and decryption (coding and decoding).

CAC

- call admission control (MPLS). A bandwidth and bandwidth-related resource monitoring and accounting facility that determines whether a setup request can be honored for an MPLS LSP with traffic parameters.
- connection admission control (ATM). A set of actions that the network takes during connection setup or renegotiation. ATM networks use CAC to determine whether to accept a connection request, based on whether allocating the connection's requested bandwidth would cause the network to violate the traffic contracts of existing connections.

call admission control *See* CAC.

CAM content-addressable memory. A category of hardware classifier. *See also* FPGA.

capability negotiation A method by which BGP peers determine whether they share the same capabilities, and whether the session will be maintained or terminated given the respective capabilities of the peers. BGP speakers advertise their capabilities in BGP open messages. *See also* cooperative route filtering.

carrier-of-carriers VPN A VPN that establishes a two-tiered relationship between a provider carrier and a customer carrier. The provider carrier provides a VPN backbone network for the customer carrier (Tier 1). The customer carrier, in turn, provides layer 3 VPN or Internet services to its end customers (Tier 2). For a carrier-of-carriers VPN, the customer's sites are configured within the same autonomous system (AS).

CBF connection-based forwarding. A method of forwarding frames in which forwarding decisions are made using only the identity of the ingress interface. No part of a packet's contents is used to determine how a packet should be forwarded.

CBR constant bit rate. An ATM service category that supports a constant and guaranteed rate to transport services such as video or voice as well as circuit emulation, which requires rigorous timing control and performance parameters. Data is serviced at a constant, repetitive rate. CBR is used for traffic that does not need to periodically burst to a higher rate, such as nonpacketized voice and audio.

CC cells continuity check cells. Cells that provide continual monitoring of a connection on a segment or from end to end.

CCITT International Telegraph and Telephone Consultative Committee. An organization that sets international data communications standards. Known since 1992 as ITU (International Telecommunication Union), a United Nations agency and the CCITT parent organization. CCITT no longer exists as a separate entity. *See also* ITU-T.

CDMA code division multiple access. A digital cellular technology that uses spread-spectrum techniques. Unlike competing systems that use TDMA (time division multiple access), such as GSM (Global System for Mobile Communications), CDMA does not assign a specific frequency to each user. Instead, every channel uses the full available spectrum. Individual conversations are encoded with a pseudo-random digital sequence. CDMA consistently provides better capacity for voice and data communications than other commercial mobile technologies, allowing more subscribers to connect at any given time.

CDV cell delay variation. The difference between a cell's expected and actual transfer delay. Determines the amount of jitter. (JUNOSe QoS term)

CDVT cell delay variation tolerance. The acceptable tolerance of CDV (jitter). (JUNOSe QoS term)

cell delay variation *See* CDV.

cell delay variation tolerance *See* CDVT.

cell loss priority *See* CLP.

certificate An electronic document that binds a person or entity to a public key using a digital signature

certificate authority *See* CA.

certificate revocation list *See* CRL.

Challenge Handshake Authentication Protocol *See* CHAP.

change of authorization *See* CoA.

channelized interface A wideband interface that is divided into many smaller channels to carry different streams of data.

CHAP Challenge Handshake Authentication Protocol. A protocol that authenticates remote users. CHAP is a server-driven, three-step authentication mechanism that depends on a shared secret password that resides on both the server and the client. *See also* PAP.

CIDR classless interdomain routing. An addressing method that replaces the traditional class structure of IP addresses. In CIDR, the boundary between the network and host portions of an IP address can be on any bit boundary. CIDR addresses have no class restrictions, enabling more efficient used of the IP address space. CIDR addresses are represented by a prefix and a notation that indicates the IP address and mask; for example, 10.12.8.3/16.

Cisco HDLC Cisco High-Level Data Link Control. A bit-oriented synchronous data-link layer protocol that governs information transfer. Developed by the International Organization for Standardization (ISO). It specifies a data encapsulation method on synchronous serial links using frame characters and checksums. *See also* SLARP.

Cisco High-Level Data Link Control *See* Cisco HDLC.

CISPR International Special Committee on Radio Interference. An IEC committee whose principal task is preparing standards that offer protection of radio reception from interference sources at the higher end of the frequency range (from 9 kHz upwards) such as electrical appliances of all types; the electricity supply system; industrial, scientific and electromedical RF; broadcasting receivers (sound and TV); and IT equipment (ITE).

CLACL classifier control list. A list that specifies the criteria by which the router defines a packet flow.

classification The process of taking a single data stream in and sorting it into multiple output substreams.

classifier A method of reading a sequence of bits in a packet header or label and determining how the packet is to be forwarded internally and scheduled (queued) for output.

classifier control list *See* CLACL.

classifier group The policy rules that make up a policy list.

classless interdomain routing *See* CIDR.

class of service *See* CoS.

clear to send *See* CTS.

CLEC competitive local exchange carrier. A company that competes with the already established local telecommunications business by providing its own network and switching.

CLI command-line interface. The interface that enables the configuration, monitoring, and management of hardware and software by entering commands on a terminal connected to the routing platform.

CLI access class The security level that grants access to specific CLI commands, such as for packet mirroring.

CLI-based packet mirroring A type of packet mirroring in which an authorized user uses the router CLI commands to configure and manage packet mirroring.

client *See* SNMP client.

CLNP Connectionless Network Protocol. An OSI network layer protocol used by CLNS to handle data at the transport layer; the OSI equivalent of IP.

CLNS Connectionless Network Service. An OSI network layer service that enables data transmission without establishing a circuit and that routes messages independently of any other messages.

CLP cell loss priority. An ATM cell bit that communicates the loss priority of the payload. A value of zero (0) specifies that the cell not be discarded if it encounters congestion as it moves through the network. A value of one (1) specifies that the network can drop the cell when congestion is encountered.

cluster A route reflector and its clients (BGP). Clients peer only with a route reflector and do not peer outside their cluster. Route reflectors peer with clients and other route reflectors within a cluster; outside a cluster they peer with other reflectors and other routers that are neither clients nor reflectors. *See also* route reflector; route reflector client.

CoA change of authorization. RADIUS messages that dynamically modify session authorization attributes, such as data filters.

code division multiple access *See* CDMA.

cold restart The result of a standby SRP module becoming active without high availability (HA) being configured (no switchover from active SRP). Similar to a cold start, except:

- The standby SRP becomes active much more quickly because the configuration is already loaded in the standby SRP memory and the device is running.
- Line module software is reloaded, so that it takes additional time for the newly active SRP to become fully operational.

See also graceful restart; warm restart.

color-aware rate limit A type of rate limit that can change the algorithm used, depending on the color of the incoming packet.

color-based thresholding In JUNOSe QoS, a process that assigns precedence to packets.Packets within the router are tagged with a drop precedence: committed—green; conformed—yellow; exceeded—red. When the queue fills above the exceeded threshold, the router drops red packets, but still queues yellow and green packets. When the queue fills above the conformed drop threshold, the router queues only green packets.

color-blind rate limit A type of rate limit that runs the same algorithm for all packets, regardless of their color. *See also* rate-limit hierarchy.

command-line interface *See* CLI.

command privileges A feature of the CLI in E-series routers. Command privileges fall within one of the following levels:

- 0—Allows you to execute the **help**, **enable**, **disable**, and **exit** commands
- 1—Allows you to execute commands in User Exec mode plus commands at level 0
- 5—Allows you to execute Privileged Exec **show** commands plus the commands at levels 1 and 0

- 10—Allows you to execute all commands except support commands, which may be provided by Juniper Networks Customer Service, or the **privilege** command to assign privileges to commands
- 15—Allows you to execute support commands and assign privileges to commands

committed action In a rate-limit profile, an action that drops, transmits, marks (IP and IPv6), or marks-exp (MPLS) when traffic flow does not exceed the rate. The mark value is not supported for hierarchical rate limits, and the transmit values conditional, unconditional, and final are supported only on hierarchical rate limits.

Common Open Policy Service *See* COPS.

Common Open Policy Service usage for policy provisioning *See* COPS-PR.

community In BGP, a logical group of prefixes that share some common attribute; used to simplify a routing policy. Community members can be on different networks and in different autonomous systems. BGP allows you to define the community to which a prefix belongs. A prefix can belong to more than one community. The community attribute lists the communities to which a prefix belongs.

community list A sequential collection of permit and deny conditions. Each condition describes the community number to be matched.

The router tests the community attribute of a route against the conditions in a community list one by one. The first match determines whether the router accepts (the route is permitted) or rejects (the route is denied) a route having the specified community. Because the router stops testing conditions after the first match, the order of the conditions is critical. If no conditions match, the router rejects the route.

competitive local exchange carrier *See* CLEC.

compound explicit shared shaper One of four types of shared shapers, in which the software selects constituents based on the shared priority and shared weight configured using a JUNOSe command. If no attributes are specified, the software supplies a shared priority consistent with the legacy scheduler configuration. *See also* compound implicit shared shaper; simple explicit shared shaper; simple implicit shared shaper.

compound implicit shared shaper One of four types of shared shapers, in which the software selects constituents automatically. If a node exists in a given traffic-class group, the node is active and the queues stacked above it are inactive constituents. *See also* compound explicit shared shaper simple explicit shared shaper; simple implicit shared shaper.

compound shared shaping A hardware-assisted mechanism that controls bandwidth for all scheduler objects associated with the subscriber logical interface. *See also* shared shaping; simple shared shaping.

concurrent routing and bridging *See* CRB.

confederation A set of sub-ASs established within an AS to reduce mesh overhead. BGP peers within each sub-AS are fully meshed, but the sub-ASs do not have to be fully meshed within the AS. *See also* route reflection.

configuration caching A mechanism that prevents the system from being partially configured with changes in the event of a reset. When a script or macro begins execution, the resulting configuration changes are automatically cached in system RAM rather than being committed to nonvolatile storage (NVS). When the script or macro completes execution, the cache is flushed as a background operation, saving the configuration changes to NVS.

conformed action In a rate-limit profile, an action that drops, transmits, marks (IP and IPv6), or marks-exp (MPLS) when traffic flow exceeds the rate but not the excess burst. The mark value is not supported for hierarchical rate limits, and the transmit values conditional, unconditional, and final are supported only on hierarchical rate limits.

connection admission control *See* CAC.

connection-based forwarding *See* CBF.

Connectionless Network Protocol *See* CLNP.

Connectionless Network Service *See* CLNS.

connectionless protocol A protocol, such as IP, that does not exchange control information to establish an end-to-end connection before transmitting data.

connection-oriented protocol A protocol that exchanges control information with the remote computer to verify that the remote computer is ready to receive data before the originating computer sends the data.

constant bit rate *See* CBR.

constituent A scheduler node or queue associated with a logical interface. A shared shaper is configured for a logical interface; all queues and scheduler nodes associated with that logical interface are constituents of the shared shaper. *See also* active constituent; inactive constituent.

Constraint-Based Routed Label Distribution Protocol *See* CR-LDP.

constraint-based routed label-switched path *See* CR-LSP.

constraint-based routing (MPLS) A mechanism to establish paths based on certain criteria (explicit route, QoS parameters). The standard routing protocols can be enhanced to carry additional information to be used when running the route calculation.

content addressable memory *See* CAM.

cooperative route filtering *See* ORF.

COPS Common Open Policy Service (protocol). A query-and-response protocol used to exchange policy information between a policy server and its clients.

COPS-PR COPS usage for policy provisioning. An IETF standard where the policy enforcement point (PEP) requests policy provisioning when the operational state of the interface and DHCP addresses change.

core dump file In E-series routers, the file that indicates which module has failed by referencing that module's hardware slot number. The hardware slot number is the slot number designation on the system backplane. This slot number is different from the chassis slot number that appears on the front of the chassis and in screen displays.

CoS class of service. A method of classifying traffic on a packet-by-packet basis using information in the type-of-service (ToS) byte to provide different service levels to different traffic. *See also* QoS.

CPE customer premises equipment. Data communications equipment (such as telephone, modem, and router) located at a customer site.

CRB concurrent routing and bridging. A mechanism where an E-series router can route a protocol among a group of interfaces in one bridge group and concurrently bridge the same protocol among a separate group of interfaces in a different bridge group on the router.

CRC cyclic redundancy check. An error-checking technique that uses a calculated numeric value to detect errors in transmitted data.

CRL certificate revocation list. A list of digital certificates that have been invalidated, including the reasons for revocation and the names of the entities that issued them. A CRL prevents usage of digital certificates and signatures that have been compromised.

CR-LDP Constraint-Based Routed Label Distribution Protocol. A traffic engineering signaling protocol for MPLS IP networks. CR-LDP provides mechanisms for establishing explicitly routed label switched paths (LSPs).

CR-LSP constraint-based routed label-switched path. An explicitly routed label switched path (LSP) established by means of CR-LDP.

CTS clear to send (signal). A signalling message transmitted in response to an RTS (Request to Send) message that enables the sender of the RTS message to begin data transfer.

customer premises equipment *See* CPE.

cyclic redundancy check *See* CRC.

D

data carrier detect *See* DCD.

data communication equipment *See* DCE.

Data Encryption Standard *See* DES.

datagram The packet format defined by IP.

data exchange interface *See* DXI.

data link layer The second level in the seven-layer OSI reference model for network protocol design and in the five-layer TCP/IP protocol stack. This layer provides the functional and procedural means to transfer data between network entities by splitting data into frames to send on the physical layer and receiving acknowledgement frames. It performs error checking, retransmitting frames not received correctly. In general, it controls the flow of information across the link, providing an error-free virtual channel to the network layer. *See also* OSI.

data service unit *See* DSU.

data set ready *See* DSP.

data stream inversion A collection of data bits in a data stream that are inverted for transmission

data terminal equipment *See* DTE.

data terminal ready *See* DTR.

DCD data carrier detect. A hardware signal defined by the RS-232C standard that indicates that the device, usually a modem, is online and ready for transmission.

DCE data communication equipment; data circuit-terminating equipment. A device, such as a modem, that provides the interface between a circuit and DTE (data terminal equipment).

D-channel Delta channel. A circuit-switched channel that carries signaling and control for B-channels. In Basic Rate Interface (BRI) applications, it can also support customer packet data traffic at speeds up to 9.6 kbps. *See also* B-channel.

DE discard eligibility (bit). A header bit in a Frame Relay packet that, when set, indicates that the frame can be discarded in preference to other frames without the DE bit set.

dead peer detection *See* DPD.

denial of service *See* DoS.

denial-of-service attack *See* DoS attack.

dense mode A multicast protocol mode where routers running dense-mode protocols forward multicast traffic except when explicitly requested not to do so. Dense mode forwarding assumes that most of the hosts on the network receive the multicast data. Routers flood packets and prune unwanted traffic every 3 minutes. *See also* sparse mode.

DES Data Encryption Standard. An encryption algorithm that uses a private 56-bit key that is applied to each 64-bit block of data. The sender and receiver of the data must each know the private key. *See also* 3DES.

designated intermediate system *See* DIS.

designated router *See* DR.

DF dont fragment (bit). One-bit flag in the IP datagram header that specifies whether or not to fragment the datagram. A value of zero (0) indicates to fragment the datagram. A value of one (1) indicates *not* to fragment the datagram.

DHCP Dynamic Host Configuration Protocol. A mechanism through which hosts using TCP/IP can obtain protocol configuration parameters automatically from a DHCP server on the network.

DHCP equal access mode A mode in which a DHCP local server works with the Juniper Networks SDX software to provide an advanced subscriber configuration and management service. In equal-access mode, the router enables access to non-PPP users. Non-PPP equal access requires the use of the E-series router DHCP local server and SDX software, which communicates with a RADIUS server.

DHCP external server A server that enables an E-series router that is not running DHCP relay or DHCP proxy server to monitor DHCP packets and to keep information for subscribers based on their IP address and MAC address. When the E-series router DHCP external server application is used, all DHCP traffic to and from the external server is monitored by the router. The services provided by integrating the E-series router DHCP external server application with SDX software are similar to those provided when the DHCP local server is integrated with SDX software. The router DHCP external application is used together with other features of the router to provide subscriber management.

DHCP proxy client The configuration that enables the router to obtain an IP address from a DHCP server for a remote PPP client. Each virtual router (acting as a DHCP proxy client) can query up to five DHCP servers. For PPP users, the router acts as a DHCP client to obtain an address for the user.

DHCP relay proxy An enhanced component of DHCP relay that manages host routes for DHCP clients, including selecting the single most appropriate offer from multiple DHCP servers.

DHCP standalone mode A mode in which the DHCP local server operates as a basic DHCP server. Clients are not authenticated by default; however, you can optionally configure the DHCP local server to use AAA authentication for the incoming clients.

dialed number identification service *See* DNIS.

dial-out route A route definition that contains the dial-out target, as well as a domain name and profile. The domain name is used in the initial Access-Request message. The profile is used to create the IP/Point-to-Point Protocol (PPP) stack for the dial-out session.

dial-out session A control entity for a triggered IP flow that is used to manage the establishment of an associated L2TP session for dial-out.

dial-out target A virtual router context and an IP address prefix, for which the arrival of an IP packet (a dial-out trigger) initiates a dial-out session.

dial-out trigger An IP packet that initiates a dial-out session

differentiated services *See* DiffServ.

Diffie-Hellman key exchange A feature of SSH that provides server authentication by protecting against hackers who interject mimics to obtain your password, so that you can be confident that you are connected to your own router

DiffServ Differentiated Services (based on RFC 2474). An architecture that provides assured forwarding and expedited forwarding by classifying packets into one of a small number of aggregated flows or traffic classes for which you can configure different QoS characteristics. The Juniper Networks QoS architecture extends DiffServ to support edge features such as high-density queuing.

digital signal *See* DS.

Digital Signal Standard *See* DSS.

digital subscriber line *See* DSL.

digital subscriber line access multiplexer *See* DSLAM.

direct server access The first authentication or accounting server that you configure in RADIUS. This server is treated as the primary authentication or accounting server, the next server configured is the secondary, and so on. *See also* round-robin server access.

DIS designated intermediate system. An IS-IS router that is elected by priority on an interface basis. In the case of a tie, the router with the highest MAC address becomes the DIS. DIS is analogous to the designated router in OSPF, although the election process and adjacencies within multiaccess media differ significantly. DIS assists broadcast routers to synchronize their IS-IS databases.

discard eligibility *See* DE.

Distance Vector Multicast Routing Protocol *See* DVMRP.

distance-vector routing A routing method that requires that each router to simply inform its neighbors of its routing table. For each network path, the receiving router picks the neighbor advertising the lowest metric, then adds this entry into its routing table for readvertisement. *See also* RIP.

distribution lists A list that controls routing information that is accepted or transmitted to peer routers. Distribution lists always use access lists to identify routes for distribution. For example, distribution lists can use access lists to specify routes to advertise. *See also* access lists.

DLCI data-link connection identifier. A 10-bit channel number that is attached to data frames to inform a Frame Relay network how to route the data in a Frame Relay virtual connection.

DNIS dialed number identification service. With DNIS, if users have a called number associated with them, the router searches the domain map for the called number. If it finds a match, the router uses the matching domain map entry information to authenticate the user. If the router does not find a match, it searches the domain map using normal processing.

DNS Domain Name System. An Internet client/server mechanism that maps domain names to IP addresses.

DNS-ALG Domain Name System—Application Level Gateway. An application-level gateway that facilitates name-to-address mapping over bidirectional NAT or twice NAT.

domain A collection of routers that are administered as a unit with common rules and procedures. Also, a collection of routers that use a common interior gateway protocol (IGP).

Domain Name System *See* DNS.

domain-specific part *See* DSP.

dont fragment (bit) *See* DF.

DoS denial of service. A system security breach in which network services become unavailable to users.

DoS attack denial-of-service attack. Any attempt to deny valid users access to network or server resources by using up all the resources of the network element or server.

downstream-on-demand A method of label distribution whereby MPLS devices do not signal a FEC-to-label binding until requested to do so by an upstream device. Downstream-on-demand conserves labels in that they are not bound until they are needed and the LSR receives label mappings (also known as label bindings) from a neighbor that is the next hop to a destination; it is used when RSVP is the signaling protocol.

See also downstream-unsolicited; independent control ; ordered control.

downstream-unsolicited An MPLS label distribution method whereby MPLS devices do not wait for a request from an upstream device before signaling FEC-to-label bindings. As soon as the LSR learns a route, it sends a binding for that route to all peer LSRs, both upstream and downstream. Downstream-unsolicited does not conserve labels, because an LSR receives label mappings from neighbors that might not be the next hop for the destination; it is used by BGP or LDP when adjacent peers are configured to use the platform label space.

See also downstream-unsolicited; independent control; ordered control.

DPD dead peer detection. A keepalive mechanism that enables the E-series router to detect when communication to a remote IPSec peer has been disconnected. DPD enables the router to reclaim resources and to optionally redirect traffic to an alternate failover destination. If DPD is not enabled, the traffic continues to be sent to the unavailable destination. Also known as IKE keepalive.

DR designated router. A designated device (OSPF router) with which other routers form adjacencies, reducing the number of adjacencies required on a broadcast or NBMA network.

drop profile A template that controls the dropping behavior of a set of egress queues. The profile defines the range within the queue where random early detection (RED) operates, the maximum percentage of packets to drop, and sensitivity to bursts of packets. Weighted random early detection (WRED) is an extension to RED that enables an administrator to assign different RED drop profiles to each color of traffic.

DS

- digital signal. A discontinuous signal used in direct sequence spread spectrum modulation, also known as direct sequence code division multiple access (DS-CDMA) DS-CDMA is one of two approaches to spread spectrum modulation for digital signal transmission over the airwaves. In direct sequence spread spectrum, the stream of information to be transmitted is divided into small pieces, each of which is allocated across to a frequency channel across the spectrum.)

- Differentiated Services (field). The IPv4 header TOS octet or the IPv6 Traffic Class octet used to mark packets to enable differentiated services. *See also* DiffServ.

DS-BGP dual-stack Border Gateway Protocol router. A BGP router that runs both the IPv4 and the IPv6 protocol stack. DS-BGP routers are typically used to connect IPv6 islands across IPv4 clouds.

DSI dynamic subscriber interface. A subscriber interface that is associated with a primary IP interface and that is dynamically created in response to an external event, such as packet detection or a DHCP event.

DSL digital subscriber line. A technology that increases the digital capacity of standard telephone lines into the home or office and provides always-on Internet operation. *See also* ADSL; SDSL.

DSLAM digital subscriber line access multiplexer. A network device directly connected to subscriber premises that handles the copper termination and aggregates traffic into a higher-speed uplink. The output from DSLAMs is fed into the router through a DS3 or OC3 link.

DSP domain-specific part. The part of the Network Service Access Point (NSAP) address that uniquely identifies a system on the network.

DSR data set ready. One of the control signals on a standard RS-232C connector that indicates whether the DCE is connected and ready to start.

DSS Digital Signature Standard authentication algorithm. A cryptographic standard used for authenticating electronic documents, much as a written signature verifies the authenticity of a paper document.

DSU data service unit. A device used to connect a DTE to a digital phone line. DSU converts digital data from a router to voltages and encoding required by the phone line.

DTE data terminal equipment. A device, such as a computer, host, or terminal, that communicates with DCE. At the terminal end of a data transmission, DTE comprises the transmit and receive equipment. *See also* DCE.

DTR data terminal ready. The signal sent over a dedicated wire (RS-232 connection) from a computer (or terminal) to a transmission device to indicate that the computer is ready to receive data.

dual-stack Border Gateway Protocol *See* DS-BGP.

duplex mode The transmission and reception of signals in both directions. *See also* full-duplex mode; half-duplex mode.

duplicate accounting server In RADIUS, a server that sends the accounting information to a particular router. You might use duplicate accounting to send the accounting information to a customer's accounting server. *See also* broadcast accounting server.

DVMRP Distance Vector Multicast Routing Protocol. An interior gateway protocol (IGP) that supports operations within an autonomous system (AS), but not between autonomous systems. The multicast backbone of the Internet uses DVMRP to forward multicast datagrams. DVMRP is a dense-mode multicasting protocol and therefore uses a broadcast-and-prune mechanism. *See also* dense mode.

DVMRP tunnels Tunnels that allow the exchange of IP multicast traffic between routers separated by networks that do not support multicast routing.

DXI data exchange interface. A specification developed by the SMDS (switched megabit data services) interest group to define the interaction between internetworking devices and CSUs/DSUs that are transmitting over an SMDS access line.

dynamic encapsulation lockout A mechanism that temporarily prevents an ATM 1483 subinterface from autodetecting, accepting, and creating dynamic interface columns for a configurable time period.

Dynamic Host Configuration Protocol *See* DHCP.

dynamic interface A type of interface that is created through some external event, typically through the receipt of data over a lower-layer link, such as an ATM virtual circuit. The layers of a dynamic interface are created based on the packets received on the link and can be configured through RADIUS authentication, profiles, or a combination of RADIUS authentication and profiles. *See also* static interface.

dynamic oversubscription A mechanism that enables the router to vary queue thresholds based on the amount of egress buffer memory in use. *See also* bandwidth oversubscription; static oversubscription.

dynamic subscriber interface *See* DSI.

dynamic translation One of two NAT methods used to assign a translated IP address. This method uses access list rules and NAT address pools. Use it when you want the NAT router to initiate and manage address translation and session flows between address realms on demand.

dynamic tunnel-server ports *See* shared tunnel-server module.

E

EAP Extensible Authentication Protocol. An extension of the PPP protocol that enables peer authentication before network layer protocols can transmit over the link. EAP supports multiple authentication methods; the specific method used is negotiated between the EAP server and the peer.

EBGP external Border Gateway Protocol. A BGP configuration in which sessions are established between routers in different autonomous systems. *See also* EBGP session.

EBGP session A BGP session between two BGP speakers that are in different autonomous systems. EBGP sessions typically exist between peers that are physically connected. *See also* IBGP session.

ECC error checking and correction; error-checking code. The process of detecting errors during the transmission or storage of digital data and correcting them automatically. This usually involves sending or storing extra bits of data according to specified algorithms.

ECMP equal-cost multipath. A traffic load-balancing feature that enables traffic to the same destination to be distributed over multiple paths that have the same cost.

ECP Encryption Control Protocol. The protocol responsible for configuring and enabling data encryption algorithms on both ends of a PPP link.

EEPROM electrically erasable programmable read-only memory. A memory chip used to store small amounts of configuration data.

effective weight The result of a weight or an assured rate. Users configure the scheduler node by specifying either an assured rate or a weight within a scheduler profile. An assured rate, in bits per second, is translated into a weight, referred to as an effective weight.

EGP exterior gateway protocol. A protocol that distributes routing information to routers that connect separate autonomous systems.*See also* IGP.

electrically erasable programmable read-only memory *See* EEPROM.

electrostatic discharge *See* ESD.

E-LSP EXP-inferred-PSC LSP. One of two types of LSPs employed by MPLS to support differentiated services. The EXP field of the MPLS shim header is used to determine the per-hop behavior applied to the packet. *See also* L-LSP; shim header.

Encapsulating Security Payload *See* ESP.

encryption A software mechanism that makes data confidential by making it unreadable to everyone except the sender and the intended recipient. *See also* IPSec.

Encryption Control Protocol *See* ECP.

endpoint discriminator An LCP negotiation option that identifies the router transmitting the packet.

end system *See* ES.

Enterprise MIB An SNMP term for a MIB defined by a single vendor. In addition to providing consistency of management data representation across that vendor's product line, the enterprise MIB also accounts for proprietary functions and value-added features not addressed by standard MIBs.

equal access mode *See* DHCP equal access mode.

equal-cost multipath *See* ECMP.

error checking and correction *See* ECC.

error-checking code *See* ECC.

ES end system. Any nonrouting network node or host in OSI internetworking. *See also* intermediate system.

ESD electrostatic discharge. Stored static electricity that can damage electronic equipment and impair electrical circuitry when released.

ESP Encapsulating Security Payload. A protocol that provides data integrity, data confidentiality and, optionally, sender's authentication.

Ethernet link aggregation *See* 802.3ad link aggregation.

event categories Classification groups and severity levels for system events that can be used to track system changes. Severity levels (categories) include Emergency, Alert, Critical, Error, Warning, Notice, Info, and Debug.

Event MIB A MIB that defines a method for creating trigger conditions, testing those conditions, and determining which action to take when a trigger meets those conditions. The Event MIB enables you to define test conditions for object integers that are accessible in the SNMP agent, making it possible to monitor any aspect of a device without defining specific notifications. *See also* event table (mteEventTable); objects table (mteObjectsTable); SNMP Server Event Manager; trigger table (mteTriggerTable).

events *See* system events.

event table (mteEventTable) An SNMP term for a table that defines what action you want a device to take when a trigger occurs. This action can be in the form of a notification, setting a specified MIB object, or both. The results of these actions are controlled within two subordinate MIB tables—notification and set.

One of the three parts of the Event MIB. *See also* objects table (mteObjectsTable); trigger table (mteTriggerTable).

exceeded action In a rate-limit profile, an action that drops, transmits, marks (IP and IPv6), or marks-exp (MPLS) when traffic flow exceeds the rate. The mark value is not supported for hierarchical rate limits, and the transmit values conditional, unconditional, and final are supported only on hierarchical rate limits.

Exec modes *See* Privileged Exec mode; User Exec mode.

explicit routing A subset of constraint-based routing where the constraint is an explicit route. In explicit routing, the route the LSP takes is defined by the ingress node.

explicit shared shaper A type of shared shaper in which you select the active constituents in a scheduler profile. A subset of the interface traffic is shaped to the shared rate. *See also* implicit shared shaper; shared shaping.

export map A route map applied to a VRF to modify or filter routes exported from the VRF to the global BGP VPN RIB in the parent VR. *See also* import map.

Extensible Authentication Protocol *See* EAP.

external Border Gateway Protocol *See* EBGP.

External Data Representation Standard *See* XDR.

exterior gateway protocol *See* EGP.

F

facilities data link *See* FDL.

failover *See* switchover.

FDL facilities data link. A type of message that can be used to determine the status of a line and to display statistics for the remote end of a connection.

FEC forwarding equivalence class. A set of packets with similar or identical characteristics that are forwarded in the same manner, on the same path, and with the same forwarding treatment. Members of a FEC are bound to the same MPLS label.

FECN forward explicit congestion notification. In a Frame Relay network, a header bit transmitted by the source device requesting that the destination device slow down its requests for data. FECN and BECN minimize the possibility that packets will be discarded when more packets arrive than can be handled. *See also* BECN.

FIB forwarding information base. In the JUNOSe software, the IP routing table. Referred to in the context of BGP.

field-programmable gate array *See* FPGA.

field-replaceable unit *See* FRU.

FIFO first-in, first-out. A scheduling method in which the first data packet stored in the queue is the first data packet removed from the queue.

file system synchronization mode The default behavior mode for E-series routers that contain redundant SRP modules. Available only to SRP modules. In this mode:

- Files and data in nonvolatile storage (NVS) remain synchronized between the primary (active) SRP module and standby SRP module.
- SRP modules reload all line modules and restart from saved configuration files.
- If the active SRP module switches over to the standby SRP, the router cold-restarts as follows: all line modules are reloaded; user connections are lost; forwarding through the chassis stops until the router SRP module recovers; the standby SRP module boots from the last known good configuration from NVS.

See also high availability mode; switchover.

File Transfer Protocol *See* FTP.

firewall A means of controlling access to a network to protect it from costly misuse and malicious intent from other users (for example, denial-of-service attacks).

first-in, first-out *See* FIFO.

flooding The distribution and synchronization of the link-state database between OSPF routers.

forward explicit congestion notification *See* FECN.

forwarding equivalence class *See* FEC.

forwarding information base *See* FIB.

forwarding table A table of the best routes to all destinations reachable by the router. For each destination, the table has only the single best route to the destination selected from the IP routing table.

forwarding table entry *See* FTE.

FPGA field-programmable gate array. A semiconductor device that contains programmable logic components and interconnects. Also a category of hardware classifier. *See also* CAM.

FQDN fully qualified domain name. The hostname and domain name for a specific system.

fractional T1 channel A DS0 portion of a 24-DS0 T1 line. Fractional T1s enable you to separate out one DS0 line or combine several lines into *bundles* (usually in multilink PPP).

fragmentation The process of segmenting a large IP datagram into several smaller pieces. Required when IP must transmit a large packet through a network that transmits smaller packets, or when the MTU size of the other network is smaller.

fragmentation and assembly

- Frame Relay—A feature that reduces excessive delays of Frame Relay packets by breaking them up into smaller fragments and interleaving them with real-time frames.
- MLPPP— *Fragmentation* is the process by which a large packet is broken up into multiple smaller fragments for simultaneous transmission across multiple links of an MLPPP bundle. *Reassembly* is the process by which the destination router reassembles the fragments into the original packets.

Frame Relay A public, connection-oriented packet service based on the core aspects of the Integrated Services Digital Network. Frame Relay eliminates all processing at the network layer and greatly restricts data-link layer processing.

Frame Relay LMI Frame Relay local management interface. Provides the operator with configuration and status information relating to the Frame Relay VCs in operation. LMI specifies a polling mechanism to receive incremental and full-status updates from the network. The router can represent either side of the User-to-Network Interface (UNI) and supports unidirectional LMI. Bidirectional support for the Network-to-Network Interface (NNI) is also supported.

FRU field-replaceable unit. A router component that customers can replace onsite.

FTE forwarding table entry. Of all destinations reachable by the router, the single best route to a given destination selected from the IP routing table.

FTP File Transfer Protocol. An application protocol that is part of the TCP/IP protocol stack. Used for transferring files between network nodes. FTP is defined in RFC 959.

full-duplex mode A transmission mode that supports transmission and reception of signals in both directions simultaneously. *See also* duplex mode; half-duplex mode.

full-mesh VPN A VPN where each site in the VPN can communicate with every other site in that same VPN. *See also* hub-and-spoke VPN; overlapping VPN.

fully qualified domain name *See* FQDN.

G

general community *See* local-use community.

Generic Routing Encapsulation *See* GRE.

giaddr gateway IP address. The address that indicates a DHCP client's subnetwork. The giaddr is usually the IP address of a DHCP relay server.

Global Configuration mode A Privileged Exec mode from which you can set parameters or enable features. Within Global Configuration mode, you can apply features globally to a router, enable a feature or function, disable a feature or function, and configure a feature or function. *See also* Privileged Exec mode; User Exec mode.

global export map A route map applied to a VRF to modify and filter routes exported by the VRF to the global BGP non-VPN RIB in the parent VR. *See also* export map; global import map; import map.

global import map A route map applied to a VRF to modify and filter routes imported to the BGP RIB of the VRF from the global BGP non-VPN RIB in the parent VR. *See also* export map; global export map; import map.

global routing table A database maintained by IP on E-series router SRP modules. Contains at most one route per protocol to each prefix in the table. *See also* local routing table; forwarding table; routing table.

GRE Generic Routing Encapsulation. A method of encapsulating SMDS packets to enable data transmission through an IP tunnel. The resulting encapsulated packet contains a GRE header and a delivery header.

graceful restart A process that enables a router whose control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers (avoiding route flapping). Without graceful restart, a control plane restart disrupts services provided by the router. Implementation varies by protocol. Also called nonstop forwarding. *See also* cold restart; warm restart.

GRE generic routing encapsulation. A general tunneling protocol that can encapsulate many types of packets to enable data transmission through a tunnel. GRE is used with IP to create a virtual point-to-point link to routers at remote points in a network.

GRE tunnel An IP tunnel that uses GRE-encapsulated IP packets to enable data transmission. The resulting encapsulated packet contains a GRE header and a delivery header. Consequently, the packet requires more processing than an IP packet, and GRE can be slower than native routing protocols. GRE tunnels can be secured with IPSec.

group node A scheduler node associated with a {port interface, traffic-class group} pair. Because the logical interface is the port, only one such scheduler node can exist for each traffic-class group above the port. This node aggregates all traffic for traffic classes in the group.

group preshared keys A secure remote access method that uses L2TP/IPSec when connecting to networks that do not use a certificate authority (CA) to issue certificates.A group preshared key is associated with a local IP address in the E-series router and is used to authenticate L2TP/IPSec clients that target this IP address as their VPN server address.

Group preshared keys are not fully secure; they open to man-in-the-middle attacks. Digital certificates are preferred instead.

H

HA high availability. The idea of reducing or eliminating single points of failure. When applied to the E-series router, high availability provides both hardware-specific and software-specific methods to ensure minimal downtime and ultimately improve the performance of your network. *See also* high availability mode.

half-duplex mode A transmission mode that supports transmission and reception of signals in both directions, but not at the same time. *See also* duplex mode; full-duplex mode.

HAR hierarchical assured rate. A calculation process that dynamically adjusts bandwidth for scheduler nodes—a more powerful and efficient method of configuring assured rates than static assured rates.

Hashed Message Authentication Code *See* HMAC.

HDLC High-Speed Data Link Control. A protocol used by PPP for the PPP interface and for providing a packet-oriented interface for the network-layer protocols.

hello messages Messages used to detect adjacent peers and maintain adjacency.

hello protocol A protocol that establishes and maintains neighbor relationships and that communication between neighbors is bidirectional. The hello protocol also dynamically discovers neighboring routers on broadcast or point-to-point networks.

hierarchical assured rate *See* HAR.

hierarchical round-robin *See* HRR.

high availability *See* HA.

high availability mode A mode that ensures rapid SRP module recovery following a switchover. High availability mode uses an initial bulk file transfer and subsequent, transaction-based mirroring. This process is referred to as *stateful SRP switchover*. In addition to keeping the contents of NVS, high availability mode keeps state and dynamic configuration data from the SRP memory synchronized between the primary and standby SRP modules.

high-density Ethernet A process by which a module allows oversubscription of Ethernet packets. The module manages oversubscription by prioritizing and dropping certain packets.

high-density keepalive mode A mode whereby, when the keepalive timer expires, the interface first verifies whether any frames were received from the peer in the prior keepalive timeout interval. If so, the interface does not send an LCP echo request (keepalive). Keepalive packets are sent only if the peer is silent (that is, no traffic was received from the peer during the previous keepalive timeout interval). Also known as smart keepalive. *See also* low-density keepalive mode.

High-Level Data Link Control *See* HDLC.

high-speed serial interface *See* HSSI.

HMAC Hashed Message Authentication Code. A mechanism for message authentication that uses cryptographic hash functions. HMAC can be used with any iterative cryptographic hash function—for example, MD5 or SHA-1—in combination with a secret shared key. The cryptographic strength of HMAC depends on the properties of the underlying hash function. Defined in RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*.

HMAC MD5 authentication An authentication method for IS-IS that prevents unauthorized routers from injecting false routing information into your network or forming adjacencies with your router. The router creates secure digests of the packets, encrypted according to the HMAC MD5 message-digest algorithms. The digests are inserted into the packets from which they are created. Depending on the commands you issue, the digests can be inserted into hello packets, link-state PDUs, complete sequence number PDUs, and partial sequence number PDUs. Also called MD5 authentication.

hop count The number of routers that data packets must traverse between RIP networks. *See also* RIP metric.

hotfix One or more files that update an operational E-series router. Hotfixes can do any of the following: address one or more specific, critical software issues by replacing or adding functionality to one or more software components; enable the delivery of software updates without having to load an entire software release; or deploy debugging code to collect data that facilitates troubleshooting of software issues.

HRR hierarchical round-robin. A scheme for allocating bandwidth to queues in proportion to their weights.

HRR scheduler One part of the integrated scheduler used to extend ATM QoS functionality on all E-series router ASIC-enabled line modules. *See also* SAR scheduler.

HSSI high-speed serial interface. An interface that supports high-speed WAN switching services such as Frame Relay and Switched Multimegabit Data Service (SMDS) trunk encapsulation. You can configure an interface to act as data communications equipment (DCE) or data terminal equipment (DTE).

hub-and-spoke VPN A VPN where the spoke sites in the VPN can communicate only with the hub sites; they cannot communicate with other spoke sites, *See also* full-mesh VPN; overlapping VPN.

I

I/O adapter *See* IOA.

I/O module A physical interface that pairs with line modules to provide connectivity to an ERX router. *See also* IOA.

IANA Internet Assigned Numbers Authority. A regulatory group that maintains all assigned and registered Internet numbers, such as IP and multicast addresses.

IAPP Inter Access Point Protocol. The IEEE 802.11F recommendation that describes optional extensions to IEEE 802.11, which defines wireless access-point communications among multivendor systems.

IBGP internal Border Gateway Protocol. A BGP configuration in which sessions are established between routers in the same autonomous system (AS). *See also* EBGP.

IBGP session A BGP session between two BGP speakers that are in the same autonomous system (AS). IBGP requires that BGP speakers within an autonomous system be fully meshed, meaning that there must be a BGP session between each pair of peers within the AS. IBGP does not require that all the peers be physically connected. *See also* EBGP session.

ICMP Internet Control Message Protocol. A protocol that provides a mechanism for a router or destination host to report an error in data traffic processing to the original source of the packet. Used in router discovery, ICMP allows router advertisements that enable a host to discover addresses of operating routers on the subnet.

ICMP Router Discovery Protocol *See* IRDP.

I-DAS integrated DHCP access server. A feature that enables you to use RADIUS start and stop attributes to track user events such as the lifetime of an IP address.

IDI initial domain identifier. The part of an ATM address format that contains the address fields describing the address allocation and issuing authority.

IDP initial domain part. The part of a CLNS address that consists of the AFI and IDI. *See also* AFI; IDI.

IEC International Electrotechnical Commission. An international standards organization that deals with electrical, electronic, and related technologies. *See also* ISO.

IEEE Institute of Electrical and Electronics Engineers. An international professional society for electrical engineers.

IETF Internet Engineering Task Force. An international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet.

IGMP Internet Group Management Protocol. A protocol that IP hosts use in IPv4 to report their multicast group memberships to neighboring routers. Similarly, multicast routers, such as E-series routers, use IGMP to discover which of their hosts belong to multicast groups and to determine whether group members are present.

IGMP proxy A method by which a router issues IGMP host messages on behalf of hosts that the router discovered through standard IGMP interfaces. The router acts as a proxy for its hosts.

IGP interior gateway protocol. A protocol that distributes routing information to routers within an autonomous system. *See also* EGP.

IKE Internet Key Exchange. A suite of protocols that provides ways to securely negotiate the shared private keys that the AH and ESP portions of IPSec need to function properly. IKE enables a pair of security gateways to:

- Dynamically establish a secure tunnel over which the security gateways can exchange tunnel and key information.
- Set up user-level tunnels or SAs, including tunnel attribute negotiations and key management. These tunnels can also be refreshed and terminated on top of the same secure channel.

IKE employs Diffie-Hellman methods and is optional in IPSec (the shared keys can be entered manually at the endpoints).

IKE endpoint The IP address of the entity that is one of two endpoints in an IKE/ISAKMP SA.

IKE policies Policies that define a combination of security parameters to be used during the IKE SA negotiation. IKE policies are configured on both security gateway peers, and there must be at least one policy on the local peer that matches a policy on the remote peer. Failing that, the two peers are not able to successfully negotiate the IKE SA, and no data flow is possible.

ILEC incumbent local exchange carrier. Any commercial telecom company that was in business after the breakup of AT&T in 1984 and before the Telecommunications Act of 1996.

ILMI Integrated Local Management Interface (a protocol). A specification developed by the ATM Forum that incorporates network management capabilities into the ATM user-to-network interface (UNI) and provides bidirectional exchange of management information between UNI management entities (UMEs).

implicit shared shaper A type of shared shaper in which the system automatically selects the active constituents. A shared-shaping rate is configured on the best-effort node or queue, and QoS locates the other constituents automatically. *See also* explicit shared shaper; shared shaping.

import map A route map applied to a VRF to modify and filter routes imported to the BGP RIB of the VRF from the global BGP VPN RIB in the parent VR. *See also* export map; global import map.

inactive constituent A constituent that is ignored by the shared shaper mechanism. *See also* active constituent; constituent.

InARP Inverse Address Resolution Protocol. A way of determining the IP address of the device at the far end of a circuit.

inbound traffic (IPSec) In the context of a secure interface, already secured traffic arriving on that interface (identified based on its SPI). This traffic is cleared and checked against the security parameters set for that interface.

incumbent local exchange carrier *See* ILEC.

independent control An MPLS label distribution method whereby the LSR sending the label acts independently of its downstream peer. It does not wait for a label from the downstream LSR before it sends a label to peers. *See also* ordered control.

initial domain identifier *See* IDI.

initial domain part *See*IDP.

input/output adapter *See* IOA.

input/output module *See* I/O module.

input policy A type of policy that evaluates a condition before the normal route lookup. *See also* output policy; policy; secondary input policy.

inside global address In a NAT context, the *translated* IP address of an inside host as seen by an outside host and network.

inside local address In a NAT context, the configured IP address that is assigned to a host on the inside network.

inside network In a NAT context, the local portion of a network that uses private, not publicly routable, IP addresses that you want to translate.

inside source translation The most commonly used NAT configuration. When an inside host sends a packet to the outside network, the NAT router translates the source information and, in the inbound direction, restores the original information. For outbound traffic, the NAT router translates the inside local address into the inside global address.

Institute of Electrical and Electronics Engineers *See* IEEE.

integrated DHCP access server *See* I-DAS.

integrated IS-IS An extended version of IS-IS that supports the routing of datagrams by means of IP or CLNS. Without the extensions, IS-IS routes datagrams only by means of CLNS.

Integrated Local Management Interface *See* ILMI.

integrated scheduler A type of QoS scheduler that provides extended ATM QoS functionality. The integrated scheduler consists of two schedulers in series—the hierarchical round robin (HRR) scheduler and the segmentation and reassembly (SAR) scheduler.

Integrated Services Digital Network *See* ISDN.

Inter Access Point Protocol *See* IAPP.

inter-AS routing Routing of packets among different autonomous systems (ASs). *See also* EBGP.

inter-AS services Services that support VPNs across AS boundaries.

interface label space A configurable pool of labels from which multiple smaller pools (ranges) of labels can be created. Interfaces are configured to use labels only from a particular pool.

interfaces Physical and logical channels on the router that define how data is transmitted to and received from lower layers in the protocol stack. *See also* subinterface.

interface specifier A label used in JUNOSe software to identify both the physical location (such as chassis slot and port number) of a particular interface type on the router and the logical interface, such as a channelized T3 interface. Used in conjunction with an interface type to uniquely identify the interface on the router. *See also* interface type.

interface type A label used in JUNOSe software to identify the type of interface you are configuring on the router. For example, gigabitEthernet indicates a Gigabit Ethernet interface. Used in conjunction with an interface specifier to uniquely identify the interface on the router. *See also* interface specifier.

interior gateway protocol *See* IGP.

intermediate system A router in OSI internetworking. *See also* ES.

internal Border Gateway Protocol *See* IBGP.

International Electrotechnical Commission *See* IEC.

International Organization for Standardization *See* ISO.

International Special Committee on Radio Interference *See* CISPR.

International Telecommunication Union — Telecommunication Standardization *See* ITU-T.

International Telegraph and Telephone Consultative Committee *See* CCITT.

Internet Assigned Numbers Authority *See* IANA.

Internet Control Message Protocol *See* ICMP.

Internet Engineering Task Force *See* IETF.

Internet Group Management Protocol *See* IGMP.

Internet Key Exchange *See* IKE.

Internet Protocol *See* IP.

Internet Protocol Control Protocol *See* IPCP.

Internet Protocol over Asynchronous Transfer Mode *See* IPoA.

Internet Protocol Security *See* IPSec.

Internet Protocol version 6 *See* IPv6.

Internet Security Association and Key Management Protocol *See* ISAKMP.

Internet service provider *See* ISP.

interprovider services *See* inter-AS services.

Inverse Address Resolution Protocol *See* InARP.

IOA input/output adapter. A physical interface that pairs with line modules to provide connectivity to E120 and E320 routers. *See also* I/O module.

IP Internet Protocol. A protocol that provides the functions necessary to deliver blocks of data (datagrams) from a source to a destination over an interconnected system of networks, where sources and destinations are identified by fixed length addresses. *See also* IP address; IPv6.

IP address A unique address that devices use to identify and communicate with each other across a network. IPv4 uses 32-bit (4 byte) addresses in a dotted-decimal notation (for example, 192.168.50.4). IPv6 uses 128-bit addresses in a hexadecimal notation of eight 16-bit components separated by colons (for example, 2001:DB8:0:0:8:822:210C:447F). *See also* IP; IPv6.

IP address classes Four classes that lend themselves to different network configurations, depending on the desired ratio of networks to hosts:

- Class A—The leading bit is set to 0, a 7-bit number, and a 24-bit local host address. Up to 125 class A networks can be defined, with up to 16,777,214 hosts per network.

- Class B—The two highest-order bits are set to 1 and 0, a 14-bit network number, and a 16-bit local host address. Up to 16,382 class B networks can be defined, with up to 65,534 hosts per network.
- Class C—The three leading bits are set to 1, 1, and 0, a 21-bit network number, and an 8-bit local host address. Up to 2,097,152 class C networks can be defined, with up to 254 hosts per network.
- Class D—The four highest-order bits are set to 1, 1, 1, and 0. Class D is used as a multicast address.

IPCP Internet Protocol Control Protocol. A network control protocol for establishing and configuring IP over a Point-to-Point Protocol (PPP) link. IPCP uses the same packet exchange mechanism as the Link Control Protocol (LCP).

IP multicast An Internet transmission method that enables a device to send packets to a group of hosts, rather than to a list of individual hosts. Routers use multicast routing algorithms to determine the best route and transmit datagrams throughout the network.

IPoA Internet Protocol over Asynchronous Transfer Mode. An interface stacking configuration supported on E-series routers. An IPoA interface is IP over ATM 1483 over ATM AAL5 over ATM.

IP reassembly A method of encapsulating and de-encapsulating packets as they enter and leave a tunnel.

IPSec Internet Protocol Security. A protocol that provides security to IP flows through the use of authentication and encryption.

- Authentication verifies that data is not altered during transmission and ensures that users are communicating with the individual or organization that they believe they are communicating with.
- Encryption makes data confidential by making it unreadable to everyone except the sender and intended recipient.

IPSec endpoint An IP address of the entity that is one of two endpoints in an IPSec SA

IPSec Service module *See* ISM.

IP tunnels A secure method of transporting datagrams between routers separated by networks that do not support all the protocols that those routers support. To configure an IP tunnel, you must first configure a TSM interface.

IPv6 Internet Protocol version 6. A version of IP designed to eventually supersede IP version 4 (IPv4). The intent is to enhance IP addressing and maintain other IPv4 functions that work well. *See also* IP address.

IRDP ICMP Router Discovery Protocol. A routing protocol used by DHCP clients.

ISAKMP Internet Security Association and Key Management Protocol. A protocol that allows the receiver of a message to obtain a public key and use digital certificates to authenticate the sender's identity. ISAKMP is key exchange independent; that is, it supports many different key exchanges. *See also* IKE.

ISAKMP SA Security associations used to secure control channels between security gateways. These are negotiated via IKE phase 1.

ISDN Integrated Services Digital Network. A set of digital communications standards that enable the transmission of information over existing twisted-pair telephone lines at higher speeds than standard analog telephone service. An ISDN interface provides multiple B-channels (bearer channels) for data and one D-channel for control and signaling information. *See also* B-channel; D-channel.

ISM IPSec Service module. A module that receives data from and transmits data to line modules that have ingress and egress ports. Does not pair with a corresponding I/O module that provides ingress and egress ports.

ISO International Organization for Standardization. A worldwide federation of standards bodies that promotes international standardization and publishes international agreements as International Standards.

ISP Internet service provider. A company that provides access to the Internet and related services.

ITU-T International Telecommunication Union—Telecommunication Standardization (formerly known as the CCITT). A group supported by the United Nations that makes recommendations and coordinates the development of telecommunications standards for the entire world.

J

JDBC Java Database Connectivity. An API that provides a standard means of database-independent connectivity between the Java platform and a wide range of databases

J-Flow A method by which you can collect IP traffic flow statistics on your routing devices. devices. J-Flow does not require any special protocol for connection setup. It also does not require any external changes to networked traffic, packets, or any other devices in the network.

K

keepalive message A signal from one endpoint to another that the first endpoint is still active. Keepalive messages are used to identify inactive or failed connections.

L

L2C layer 2 control. *See* ANCP.

L2TP access concentrator *See* LAC.

L2TP dial-out A way for corporate virtual private networks (VPNs) that use Broadband Remote Access Server (B-RAS) to dial out to remote offices that have only narrowband dial-up access.

L2TP network server *See* LNS.

L2TP tunnel switching The router configuration that enables you to switch packets between one session terminating at an L2TP LNS and another session originating at an L2TP LAC. A tunnel-switched LAC differs from a conventional LAC because it uses two interface columns: one for the incoming session (LNS) and one for the outgoing session (LAC). The router forwards traffic from the incoming session to the outgoing session and vice versa.

label edge router *See* LER.

label-switching router *See* LSR.

LAC L2TP access concentrator. A device that receives packets from a remote client and forwards them to an L2TP network server (LNS) on a remote network.

LACP Link Aggregation Control Protocol. A mechanism for exchanging port and system information to create and maintain LAG bundles.

LAG link aggregation group. A group of two or more network links bundled together to appear as a single link. Distributes MAC clients across the link layer interface and collects traffic from the links to present to the MAC clients of the LAG. Also known as a bundle.

LAN local area network.A computer network covering a local area, like a home, office or small group of buildings such as a college.

latency A delay in the transmission of a packet through a network from beginning to end.

layer 1 *See* physical layer.

layer 2 *See* data link layer.

layer 2 control *See* L2C.

layer 3 *See* network layer.

layer 4 *See* transport layer.

layer 5 *See* application layer; session layer.

layer 6 *See* presentation layer.

layer 7 *See* application layer.

LCP Link Control Protocol. A traffic controller used to establish, configure, and test data-link connections for the Point-to-Point Protocol (PPP).

LDP MD5 authentication A means of providing protection, using a shared secret (password), against spoofed TCP segments that can be introduced into the connection streams for LDP sessions. Authentication is configurable for both directly connected and targeted peers. Any given pair of peers must share the same password.

LER label edge router. A label-switching router serving as an ingress node or an egress node.

Level 1 routing Routing *within* an area:

- Level 1 routers (or intermediate systems) track all the individual links, routers, and end systems within a level 1 area.
- Level 1 routers do not know the identity of routers or destinations outside their area.
- A level 1 router forwards all traffic for destinations outside its area to the nearest level 2 router within its area.

Level 2 routing Routing *between* areas:

- Level 2 routers know the level 2 topology and know which addresses are reachable through each level 2 router.
- Level 2 routers track the location of each level 1 area.
- Level 2 routers are not concerned with the topology within any level 1 area (for example, the details internal to each level 1 area).
- Level 2 routers can identify when a level 2 router is also a level 1 router within the same area.
- Only a level 2 router can exchange packets with external routers located outside its routing domain.

line layer For a channelized OCx/STMx interface, the layer that manages the transport of SONET/SDH payloads, which are embedded in a sequence of STS/STM frames in the physical medium. This layer is responsible for multiplexing and synchronization. *See also* path layer; section layer.

line module *See* LM.

line module redundancy A configuration where an extra line module in a group of identical line modules provides redundancy if one of the modules fails. The process by which the router switches to the spare line module is called *switchover*. The requirements for line module redundancy depend on the type of router that you have.

Link Aggregation Control Protocol *See* LACP.

link aggregation group *See* 802.3ad link aggregation; LAG.

Link Control Protocol *See* LCP.

Link Integrity Protocol *See* LIP.

link layer *See* data link layer.

link-state advertisement *See* LSA.

link-state database *See* LSDB.

link-state PDU *See* LSP.

LIP Link Integrity Protocol. A protocol that runs on the member links of a Multilink Frame Relay (MLFR) bundle. Several types of LIP messages allow member links to join and leave the bundle.

LLC logical link control. The higher of two sublayers that make up the ISO/OSI data link layer. The LLC is responsible for managing communications links and handling frame traffic. *See also* data link layer; OSI.

L-LSP Label-only-inferred-PSC LSP. One of two types of LSPs employed by MPLS to support differentiated services. The per-hop behavior applied to the packet is determined from the packet label and the EXP field of the MPLS shim header. *See also* E-LSP; shim header.

LM line module. A module that acts as a frame forwarding engine for the physical interfaces (I/O modules and IOAs) and processes data from different types of network connections.

LMI Local Management Interface. Enhancements to the basic Frame Relay specifications, providing support for the following:

- A keepalive mechanism that verifies the flow of data
- A multicast mechanism that provides a network server with a local DLCI and multicast DLCI
- In Frame Relay networks, global addressing that gives DLCIs global instead of local significance
- A status mechanism that provides a switch with ongoing status reports on known DLCIs

LNS L2TP network server. A node that acts as one side of an L2TP tunnel endpoint and is a peer to the LAC. The logical termination point of a PPP connection that is being tunneled from the remote system by the LAC.

local address pool alias An alternate name for an existing local address pool. It consists of an alias name and a pool name.

local address server A server that allocates IP addresses from a pool of addresses stored locally on the router. A local address server is defined in the context of a virtual router. Local address servers exist as long as the virtual router exists or until you remove them by deleting all configured pools.

local area network *See* LAN.

local ATM passthrough The ability for the router to emulate packet-based ATM switching. Useful for customers who run IP in most of their network but still have to carry a small amount of native ATM traffic.

local authentication server An AAA local authentication server that enables the E-series router to provide local PAP and CHAP user authentication for subscribers. The router also provides limited authorization, using the IP address, IP address pool, and operational virtual router parameters. When a subscriber logs on to the E-series router that is using local authentication, the subscriber is authenticated against user entries in a local user database; the optional parameters are assigned to subscribers after the subscriber is authenticated.

local loopback The ability to loop the data back toward the router; on supported line modules. Also sends an alarm indication signal out toward the network

Local Management Interface *See* LMI.

local routing table A database local to the protocol that contains all the routes known by that protocol to the prefixes in the table. Also known as a routing information base, or RIB. *See also* global routing table; routing table.

local-use community A convenient way to categorize groups of routes to facilitate the use of routing policies. Also called private community or general community.

logical link control *See* LLC.

loopback *See* local loopback; network loopback; remote loopback.

loopback address An IP address type used by a node to send a packet to itself.

loose hop In the context of traffic engineering, a path that can use any router or any number of other intermediate (transit) points to reach the next address in the path. *See also* strict hop.

loose-source routing An MPLS routing method that specifies a set of hops that the packet must traverse, but not necessarily every hop in the path. That is, the specified hops do not have to be adjacent. *See also* strict-source routing.

low-density keepalive mode A mode in which, when the keepalive timer expires, the interface always sends an LCP echo request, regardless of whether the peer is silent. *See also* high-density keepalive mode.

LSA link-state advertisement. A packet used by link-state protocols, such as OSPF and IS-IS, that contain information about neighbors and path costs. LSAs are used by the receiving routers to maintain their routing tables.

LSDB link-state database. A computerized representation of the topology of an autonomous system. *See also* AS.

LSP

- link-state PDU (IS-IS). A PDU broadcast by link-state protocols that contains information about neighbors and path costs; used to maintain routing tables; also known as link-state advertisement
- label-switched path (MPLS). The path traversed by a packet that is routed by MPLS. Some LSPs act as tunnels. LSPs are unidirectional, carrying traffic only in the downstream direction from an ingress node to an egress node.
- link-state protocol. A routing protocol, such as OSPF and IS-IS, where each router shares information with other routers (by flooding information about itself to every reachable router in the routing area) to determine the best path. Link-state protocols use characteristics of the route such as speed and cost, as well as current congestion, to determine the best path. In link-state routing, every node receives a map of the connectivity of the network. Each node then independently calculates the best next hop from it for every possible destination in the network. The collection of best next hops forms the routing table for the node. Link state information is transmitted only when something has changed in the network. *See also* routing table.

LSP priority level The relative importance of an LSP that determines which LSPs can preempt other LSPs.Priorities are in the range 0–7 in order of *decreasing* priority.

LSR label-switching router. A router on which MPLS is enabled and that can process label-switched packets. An MPLS node that can forward layer 3 packets based on their labels.

M

MAC message authentication code. A short piece of information used to authenticate a message.

MAC address validation A verification process performed on each incoming packet to prevent spoofing on IP Ethernet-based interfaces, including bridged Ethernet interfaces.

magic number A randomly generated number used to identify one end of a point-to-point connection. Each side negotiates its magic number, taking note of each other's magic number. If both sides discover that the magic numbers they are negotiating are the same, each side attempts to change its magic number. If they are not successful, and the magic numbers remain the same, the session terminates because of the loopback that is detected.

main mode A mode of IKE phase 1 negotiations that protects the identities of the peers during negotiations and enables greater proposal flexibility than aggressive mode. Main mode is more time consuming than aggressive mode because more messages are exchanged between peers. (Six messages are exchanged in main mode.) *See also* aggressive mode.

maintenance data link *See* MDL.

Manual Commit mode A feature of JUNOSe software where configuration changes affect only the current system configuration (the running configuration), without affecting the CLI prompt.

manual secure IP interfaces Interfaces that use a preconfigured set of SA parameters to secure traffic flowing through a secure IP interface. If SA parameters do not use a preconfigured, manual secure interface, the interface drops all traffic it receives. The router keeps statistics for dropped traffic. Both peer security gateways must contain a manually provisioned manual secure IP tunnel. *See also* signaled secure IP interface.

map tag A unique string used to identify a route map

master router The VRRP router that takes the responsibility of forwarding packets sent to the IP addresses associated with the virtual router, and that answers ARP requests for these IP addresses. If the IP address owner is available, it always becomes the master. *See also* backup router.

match clause The part of a route map that specifies the attribute values that determine whether a route matches the route map. A route that has the same attribute values passes the match condition. Routes that pass all the match conditions match the route map.

match policy list A list that is similar to a route map but that contains only match clauses and no set clauses. *See also* policy list.

MAU medium attachment unit. A small device that converts signals between an attachment unit interface (AUI) and coaxial cable.

maximum received reconstructed unit *See* MRRU.

MBGP multicast Border Gateway Protocol. An application of MP-BGP that enables BGP to carry IP multicast routes, permitting the configuration of a multicast routing topology different from one's unicast topology.

MBone multicast backbone. An interconnected set of subnetworks and routers that support the delivery of IP multicast traffic. The MBone is a virtual network that is layered on top of sections of the physical Internet.

MD5 Message Digest 5. A one-way hashing algorithm that produces a 128-bit hash used for generating message authentication signatures. MD5 is used in AH and ESP.

MD5 authentication *See* HMAC MD5 authentication.

MDL maintenance data link. A type of message that can be used to determine the status of a line and to display statistics for the remote end of a connection.

MED multiple exit discriminator. An optional BGP path attribute consisting of a metric value that is used to determine the exit point to a destination when all other factors determining the exit point are equal.

mediation device *See* analyzer device.

medium attachment unit *See* MAU.

Message Authentication Code *See* MAC.

Message Digest 5 *See* MD5.

midplane A hardware component that physically separates front and rear cavities inside the chassis, distributes power from the power supplies, and transfers packets and signals between router components, which plug into it. *See also* redundancy midplane.

mirrored interface The statically or dynamically configured interface on which traffic is being mirrored during packet mirroring on E-series routers.

mirrored user The user whose traffic is being mirrored during packet mirroring on E-series routers.

MLD Multicast Listener Discovery. An IPv6 protocol that hosts use to report their multicast group memberships to neighboring routers. Similarly, multicast routers, such as E-series routers, use MLD to discover which of their hosts belong to multicast groups.

MLD proxy A method by which the router issues MLD host messages on behalf of hosts that the router discovered through standard MLD interfaces. The router acts as a *proxy* for its hosts.

MP-BGP Border Gateway Protocol multiprotocol extensions (sometimes referred to as multiprotocol Border Gateway Protocol). Extensions to BGP that enable it to carry routing information for multiple network layer protocols instead of only for IP. Includes the ability to carry multicast routing information.

MPLS Multiprotocol Label Switching. A mechanism for engineering network traffic patterns that functions by assigning to network packets short labels that describe how to forward them through the network. Also called *label switching*. *See also* TE.

MPLS edge node An MPLS node that connects an MPLS domain with a node outside the domain that either does not run MPLS or is in a different domain

MPLS egress node An MPLS edge node in the role of handling traffic as it leaves an MPLS domain

MPLS FEC A set of packets that are all forwarded in the same manner by a given LSR

MPLS forwarding table A table that maps MPLS labels to next hops. MPLS looks up the outermost label in a received packet in the forwarding table to determine what labels to push on the packet's label stack and where to send the packet.

MPLS ingress node An MPLS edge node in the role of handling traffic as it enters an MPLS domain

MPLS node A router running MPLS. An MPLS node is aware of MPLS control protocols, operates one or more layer 3 routing protocols, and is capable of forwarding packets based on labels. Optionally, an MPLS node can be capable of forwarding native layer 3 packets.

MPLS traffic engineering The ability to establish LSPs according to particular criteria (constraints) in order to meet specific traffic requirements rather than relying on the path chosen by the conventional IGP. The constraint-based IGP examines the available network resources and calculates the shortest path for a particular tunnel that has the resources required by that tunnel. Traffic engineering enables you to make the best use of your network resources by reducing overuse and underuse of certain links.

mroute A multicast traffic flow entry used for forwarding multicast traffic

MRRU multilink maximum received reconstructed unit. Similar to the MTU, but is specific to link services interfaces such as MLPPP.

multicast address A type of IPv4 and IPv6 address used for sending packets to multiple destinations. Improves network efficiency by enabling a host to transmit a packet to a targeted group of receivers.

multicast Border Gateway Protocol *See* MBGP.

Multicast Listener Discovery *See* MLD.

multinetting A method for adding more than one IP address to an IP interface—that is, a primary address and one or more secondary addresses.

multiple exit discriminator *See* MED.

multipoint connection A single-source end system connected to multiple destination end systems. Multipoint indicates a nonbroadcast multiaccess (NBMA) interface.

Multiprotocol Border Gateway Protocol *See* MP-BGP.

Multiprotocol Label Switching *See* MPLS.

munged QoS profile The set of rules used for a given forwarding interface. This set results from a process in which rules from all the QoS profiles are combined.

N

NAS network access server. A device that provides connections to a single user, to a network or subnetwork, and to interconnected networks. In reference to TACACS + , the NAS is the E-series router.

NAT Network Address Translation. A method of concealing a set of host addresses on a private network behind a pool of public addresses. It allows conservation of registered IP addresses within private networks and simplifies IP address management tasks through a form of transparent routing. It increases network privacy by hiding internal IP addresses from external networks. It can be used as a security measure to protect the host addresses from direct targeting in network attacks. *See also* bidirectional NAT; traditional NAT; twice NAT.

NAT passthrough mode A mode where the router does not check UDP checksums. Used because a NAT device may change the IP address while the UDP header is encrypted. In this case, the UDP checksum cannot be recalculated. Not checking UDP checksums for a single remote user does not compromise security, because IPSec protects UDP with an authentication algorithm far stronger than UDP checksums. But NAT passthrough mode does not support secure access to the router by multiple remote users at locations such as hotels or airports where a NAT device resides between the router and the remote users. In addition, NAT passthrough mode does not provide secure access for groups of remote users at corporate locations where a NAT device resides between the company's intranet and the public IP network. *See also* NAT-T.

NAT-T NAT Traversal. IETF standards that allow secure router access for multiple remote hosts behind a NAT device.

NBMA nonbroadcast multiaccess. A network that connects two or more devices but does not permit broadcast or multicast addressing.*See also* BMA.

NEBS Network Equipment Building System. A set of guidelines originated by Bell Laboratories in the 1970s to assist equipment manufacturers in designing products that were compatible with the telecom environment.

neighbor An adjacent system reachable by traversing a single subnetwork. An immediately adjacent router. Also called a peer. *See also* adjacency.

Neighbor Discovery A method for determining the link layer addresses of neighbors that reside on attached links and overriding invalid cache entries. Neighbor Discovery is not a true protocol, but routers and hosts (nodes) use Neighbor Discovery messages to determine the link-layer addresses of neighbors that reside on attached links and to overwrite invalid cache entries. Hosts also use it to find neighboring routers that can forward packets on their behalf, and to actively track the ability to reach neighbors.

neighboring routers Routers that have interfaces to a common network.

nested profile assignment A profile that references another profile that configures attributes for a dynamic upper-interface encapsulation type.

NET network entity title. An ISO network address used by CLNS networks; an identifier of a network entity in an end system or intermediate system. A NET consists of an area address (routing domain), system identifier, and selector.

network access server *See* NAS.

Network Address Translation *See* NAT.

Network Address Translation Traversal *See* NAT-T.

network element In SNMP, a hardware device, such as a PC or a router. Also known as a managed device.

network entity title *See* NET.

Network Equipment Building System *See* NEBS.

network layer The third level in the seven-layer OSI reference model for network protocol design and in the five-layer TCP/IP protocol stack. This layer performs the basic task of routing data across the network (getting packets of data from source to destination).

network loopback The ability to loop data toward the network before the data reaches the frame.

network management station *See* NMS.

network management system *See* NMS.

network mask *See* subnet mask.

network service access point *See* NSAP.

Network Time Protocol *See* NTP.

network-to-network interface *See* NNI.

NMS network management system; network management station. A system that enables a user to configure and monitor network elements.

NNI network-to-network interface. An interface that makes connections possible between users connected to different Frame Relay networks. These separate Frame Relay networks can be considered as subnetworks within a complete network service.

nonbroadcast multiaccess *See* NBMA.

nonbroadcast network A network that has no broadcast capability but supports more than two routers

nonce A random value used to detect and protect against replay attacks (IPSec)

non-PPP equal access A method of allowing remote access in which the router provides IP addresses to subscribers' computers through Dynamic Host Configuration Protocol (DHCP). This method is particularly convenient for broadband (cable and DSL) environments or environments that use bridged Ethernet over ATM, because network operators can support one central system rather than an individual PPPoE client on each subscriber's computer.

nonstop forwarding *See* graceful restart.

nonvolatile storage *See* NVS.

notification In SNMP, a message that indicates a status change (equivalent to a trap).

not-so-stubby area *See* NSSA.

NSAP network service access point. A connection to a network that is identified with a hierarchical network address that specifies the point at which network services are made available to a transport layer entity in the OSI reference model. A valid NSAP address is unique and unambiguously identifies a single system. Also called ISO address.

NSF nonstop forwarding. *See* graceful restart.

NSSA not-so-stubby area. In OSPF, a type of stub area in which external routes can be flooded but that can also import selected external link-state advertisements (LSAs).

NTP Network Time Protocol. A protocol that provides a method of synchronizing the system clocks of hosts on the Internet to Universal Coordinated Time (UTC). You can configure your router to update its clock automatically by configuring it as a Network Time Protocol (NTP) client. Using NTP enables the system to record accurate times of events. You can view the log file of events to monitor the status of the network.

null interface A way to handle undesired traffic. The router creates it automatically. It is always up, cannot be deleted, and acts as a data sink. That is, it cannot forward or receive traffic.

NVS nonvolatile storage. Memory that retains stored information even when power is lost to the device.

NVS card A memory card on an SRP module that stores system software, configuration files, and core dumps.

O

OAM operations, administration, and management. An ATM Forum specification for monitoring ATM virtual connections. OAM performs standard loopback, fault detection and notification, and remote defect identification for each connection, verifying that the connection is up and the router is operational.

objects table (mteObjectsTable) An SNMP term for a table that defines objects that you want to add to event messages. That is, you can create a list of user-specified objects and bind them to a trigger event. This can provide a snapshot of other values on a router when the trigger occurs. You can bind objects to a specific trigger, a type of test (for example, existence or Boolean tests), or a type of event (for example, rising or falling events). One of the three parts of the Event MIB. *See also* event table (mteEventTable); trigger table (mteTriggerTable).

ODBC Open Database Connectivity. A standard or open application programming interface (API) for accessing a database.

one-rate rate-limit profile A type of profile in which. when the committed rate is exceeded, the rate limiter drops a single packet and then resumes transmission up to a configurable burst window. *See also* rate-limit profile; two-rate rate-limit profile.

opaque LSAs LSAs that provide a generalized way of extending OSPF. The router generates opaque LSAs to carry traffic engineering information, accepts them from other routers, and floods them accordingly. OSPF uses the traffic engineering information to build a database from which paths can be computed for MPLS label-switched paths.

Open Database Connectivity *See* ODBC.

Open Shortest Path First *See* OSPF.

Open Systems Interconnection *See* OSI.

operational virtual router For a secure IP tunnel, the VR in which a secure IP tunnel exists. *See also* transport virtual router.

Operation, Administration, and Maintenance *See* OAM.

ordered control An MPLS label distribution method whereby an LSR does not advertise a label for a FEC unless it is the egress LSR for the FEC, or until it has received a label for the FEC from its downstream peer. In this manner, the entire LSP is established before MPLS begins to map data onto the LSP, preventing inappropriate (early) data mapping from occurring on the first LSR in the path. JUNOSe software does not support ordered control when LDP or BGP is the signaling protocol. *See also* downstream-on-demand; independent control.

ORF outbound route filter; outbound route filtering. A BGP capability that enables a BGP speaker to send its inbound route filter to a peer, which then installs that filter to apply after its own outbound route filter. The BGP peer then sends to the BGP speaker only routes desired by that speaker, thus minimizing the number of unwanted routing updates sent.

OSI Open Systems Interconnection. The standard reference model for how messages are transmitted between two points on a network.

OSPF Open Shortest Path First. A interior gateway protocol (IGP) that advertises the states of local network links within an autonomous system (AS) and makes routing decisions based on the shortest-path-first (SPF) algorithm. OSPF is a link-state routing protocol, similar to the Intermediate System–to–Intermediate System (IS-IS) routing protocol. OSPF was designed expressly for the TCP/IP Internet environment, including explicit support for classless interdomain routing (CIDR) and the tagging of externally derived routing information. *See also* AS.

outbound route filter (filtering) *See* ORF.

outbound traffic (IPSec) In the context of a secure interface, the clear traffic forwarded to the interface (either by policy or by routing) that is typically secured according to security parameters set for that interface.

output policy A type of policy that is applied to packets before they leave an interface. *See also* input policy; policy; secondary input policy.

outside global address In a NAT context, the configured, publicly routable IP address assigned to a host on the outside network.

outside local address In a NAT context, the *translated* IP address of an outside host as it appears to the inside network.

outside network In a NAT context, the public portion of a network that uses legitimate, publicly routable IP addresses to which you want private hosts to connect.

outside source information Information used in NAT configuration only when addresses of external hosts might create a conflict on a private network. When an outside host sends a packet inbound to the inside network, the NAT router translates the source information and, in the outbound direction, restores the original information. For inbound traffic, the NAT router translates the outside global address into the outside local address.

overlapping VPN A VPN where a site is a member of more than one VPN. An overlapping VPN is often used to provide centralized services. The central site might contain DNS servers or WWW servers or management stations that need to be reachable from multiple VPNs. Overlapping IPv4 and IPv6 VPNs are supported by the same route-target mechanism. *See also* full-mesh VPN; hub-and-spoke VPN.

oversubscription A method that allows provisioning of more bandwidth than the line rate of the physical interface. *See also* bandwidth oversubscription.

P

packet detection For GRE tunnel interfaces, an event when the router receives a packet with a source IP address that is not in the demultiplexer table, which triggers dynamic creation of subscriber interfaces. In this case, the primary IP interface must be in autoconfiguration mode. Packet detection is the only method of dynamically creating subscriber interfaces on GRE tunnel interfaces; you cannot use DHCP local server or DHCP external server.

packet mirroring A JUNOSe software feature that enables you to send a copy of a packet to an external host for analysis. Packet mirroring has many uses, including traffic debugging and troubleshooting user networking problems. With it you can mirror traffic traversing a specific interface or traffic that is to or from a particular user. Packet mirroring is always transparent to users and does not affect the delivery of the original traffic. In some cases, the means and authority for conducting packet mirroring can depend on the regulations of specific countries. *See also* CLI-based packet mirroring; RADIUS-based packet mirroring.

packet over SONET/SDH The serial transmission of data over SONET frames through the use of a protocol such as PPP.

packet-switching network A network that uses the addressing information in packets to switch packets from one physical network to another, moving each packet toward its final destination.

PADI PPPoE Active Discovery Initiation packet. A Point-to-Point Protocol over Ethernet (PPPoE) initiation packet that is broadcast by the client to start the discovery process.

PADM PPPoE Active Discovery Message. A control message that servers send to clients.

PADN PPPoE Active Discovery Network. A message that a PPPoE server sends to a client. The information sent associates the PPPoE sessions with a set of routes. The client can use this set of routes to determine which session to use based on the destination IP address.

PADO PPPoE Active Discovery Offer packet. A Point-to-Point Protocol over Ethernet (PPPoE) offer packet that is sent to the client by one or more access concentrators in reply to a PPPoE Active Discovery Initiation (PADI) packet.

PADR PPPoE Active Discovery Request packet. A Point-to-Point Protocol over Ethernet (PPPoE) packet sent by the client to one selected access concentrator to request a session.

PADS PPPoE Active Discovery Session packet. A Point-to-Point Protocol over Ethernet (PPPoE) packet sent by the selected access concentrator to confirm the session.

PADT PPPoE Active Discovery Termination packet. A Point-to-Point Protocol over Ethernet (PPPoE) packet sent by either the client or the access concentrator to terminate a session.

PAP Password Authentication Protocol. A security protocol that uses password protection to authenticate a user to a network or host. *See also* CHAP.

partial sequence number PDU (protocol data unit) *See* PSNP.

passive interface An interface that only advertises its IP address in its LSPs; it does not send or receive IS-IS packets.

passive peers BGP peers that a BGP speaker accepts inbound BGP connections from but that never initiates an outbound BGP connection to the peers. This passive status conserves CPU and TCP connection resources when the neighbor does not exist.

Password Authentication Protocol *See* PAP.

path layer For a channelized OCx/STMx interface, the layer that maps the user payload into a SONET/SDH format suitable for the line layer. This layer transports the actual network services (such as T3s) between SONET/SDH multiplexing devices and provides end-to-end transmission. *See also* line layer; section layer.

PBX private branch exchange. A private telephone system that enables telephone extensions within the system to connect with each other as well as with the public telephone system.

PCMCIA Personal Computer Memory Card International Association. An industry group that promotes standards for credit card-size memory and I/O devices.

PCR peak cell rate. The maximum allowable rate, measured in cells per second, at which cells can be transported along a connection in an ATM network.

PDP Policy decision point. The COPS server, which makes policy decisions for itself and for clients that request decisions. The SRC (formerly called SDX) application is the PDP.

PDU protocol data unit. The OSI term equivalent to packet, containing protocol control information and, possibly, user data.

PE *See* PE router.

peak cell rate *See* PCR.

peer *See* BGP peer.

pending state A state of an SRP module that the system transitions to when an unsupported application is configured. When a transition to the pending state occurs, the system generates SNMP traps and log messages. How the router behaves depends on which high availability state the application is in when it shifts to a pending state.

penultimate hop popping *See* PHP.

PEP Policy enforcement point. The COPS client, which enforces policy decisions. The JUNOSe software COPS interface is a PEP.

perfect forward secrecy *See* PFS.

per-hop behavior *See* PHB.

permanent virtual channel; permanent virtual circuit; permanent virtual connection *See* PVC.

PE router provider edge router. A router in the service provider's network that is connected to a customer edge (CE) device and participates in a virtual private network (VPN). *See also* P router.

persistent tunnel A tunnel that is configured to remain available. Persistent tunnels have only local significance; that is, they apply only to the end of the tunnel where they are set. If the other end of the tunnel chooses to terminate the tunnel, the tunnel is removed.

Personal Computer Memory Card International Association *See* PCMCIA.

PFC Protocol Field Compression. For all protocols with identifiers from 0x0000 through 0x00ff, a compression method that enables routers to compress the protocol field to one byte, as defined in RFC 1661, *The Point-to-Point Protocol (PPP)*. PFC allows you to conserve bandwidth by transmitting less data. Normally, PPP-encapsulated packets are transmitted with a two-byte protocol field. *See also* ACFC.

PFS perfect forward secrecy. An optional feature that causes every newly refreshed key to be completely unrelated to the previous key. PFS provides added security, but requires extra processing for a new Diffie-Hellmann key exchange on every key refresh.

PHB per-hop behavior. Traffic conditioning applied to traffic at each node in a differentiated services domain. The PHB provides the scheduling behavior and drop probability required by the traffic.

PHP penultimate hop popping. A mechanism used in an MPLS network that allows the transit router before the egress router to perform a label pop operation and forward the remaining data (often an IPv4 packet) to the egress router.*See also* UHP.

physical layer The first and lowest level in the seven-layer OSI reference model for network protocol design and in the five-layer TCP/IP protocol stack. This layer defines all the electrical and physical specifications for devices and provides the transmission of bits over the network medium. It includes the physical media: cables, microwaves, and networking equipment such as hubs and repeaters.

physical layer convergence procedure *See* PLCP.

PIB Policy Information Base. A collection of sets of attributes that represent configuration information for a device.

PIM Protocol Independent Multicast. A protocol that enables multicast routers to identify other multicast routers to receive packets.

PIM dense mode Protocol Independent Multicast dense mode. A dense-mode multicast protocol that uses a reverse-path multicast, flood-and-prune mechanism. *See also* dense mode.

PIM sparse-dense mode Protocol Independent Multicast sparse-dense mode. A protocol the router uses to send data when a rendezvous point (RP) is not known for a group. However, if the router discovers an RP or you configure an RP statically, PIM sparse mode takes over.

PIM sparse mode Protocol Independent Multicast sparse mode. A sparse-mode multicast protocol, which uses *shared trees*. In a shared tree, sources forward multicast datagrams to a directly connected router, the designated router. The designated router encapsulates the datagram and unicasts it to an assigned rendezvous point router, which then forwards the datagram to members of multicast groups. *See also* sparse mode.

PIM sparse mode remote neighbors Neighbors that are used to run multicast services over BGP/MPLS virtual private networks.

PIM SSM Protocol Independent Multicast source-specific multicast. An extension of the PIM protocol where a client can receive multicast traffic directly from the source. PIM SSM uses PIM sparse mode functionality to create a shortest-path tree (SPT) between the client and the source, but builds the SPT without using a rendezvous point.

pipe (and short-pipe) model A tunneling model whereby any traffic conditioning (in a pure JUNOSe environment, a change in traffic class/color combination) that is applied when traffic goes through the tunnel has no effect on the EXP bits coding in the inner header. That is, when traffic exits an LSP (when a label is popped) or when traffic enters an LSP, the inner header's EXP bits coding is not changed.

The pipe and short-pipe models differ in the header that the tunnel egress uses when it determines the PHB of an incoming packet. With the short-pipe model, the tunnel egress uses an inner header that is used for forwarding. With the pipe model, the outermost label is always used. Because of this, you cannot use PHP with the pipe model.

See also uniform model.

PKCS Public-Key Cryptography Standards. A series of standards established by RSA Laboratories

PKCS10 PKCS #10. A syntax used for digital certificate certification requests.

platform label space A large, single, unconfigurable pool of labels that can be shared by the platform—all MPLS interfaces on a given virtual router. *See also* interface label space.

PLCP physical layer convergence procedure. A protocol defined by IEEE 802.6 that is used for DS3 transmission of ATM. ATM cells are encapsulated in a frame defined by the PLCP, which is defined by the DS3 M-frame.

point of presence *See* POP.

point-to-point circuits In IS-IS, circuits that have less overhead than broadcast circuits, because they do not use designated routers, the link-state database has no representation of the pseudonode or network LSA, and they do not require periodic database synchronization. However, if more than two routers are connected on the LAN media, routing information in the network is reduced. *See also* broadcast circuits.

point-to-point connection A standard connection; for example, a connection between two ATM end stations.

Point-to-Point Protocol *See* PPP.

policy A condition and an action that are attached to an interface. The condition and action cause the router to handle the packets passing through the interface in a certain way. *See also* input policy; output policy; secondary input policy.

policy decision point *See* PDP.

policy enforcement point *See* PEP.

Policy Information Base *See* PIB.

policy list In policy management, a set of rules, each of which specifies a policy action.

policy management A feature that allows network service providers to implement packet forwarding and routing specifically tailored to their customer's requirements. Using policy management, customers can implement policies that selectively cause packets to take different paths.

policy routing A routing method that redefines a classified packet flow to a destination port or IP address.

policy rule A policy action optionally combined with a classification. A set of policy rules defines what specialized treatment to apply to classified traffic flows.

POP point of presence. The demarcation point between two networks (for example, between a LAN and a WAN)

port shaping A method for shaping the aggregate traffic through a port or channel to a rate that is less than the line or port rate.

POS packet over SONET. A communications protocol for transmitting packets over SDH or SONET, which are both circuit switched protocols.

PPP Point-to-Point Protocol. A link-layer protocol that provides multiprotocol encapsulation. PPP is used for link-layer and network-layer configuration. Provides a standard method for transporting multiprotocol datagrams over point-to-point links. Defined in RFC 1661.

PPPoE Active Discovery Initiation *See* PADI.

PPPoE Active Discovery Message *See* PADM.

PPPoE Active Discovery Network *See* PADN.

PPPoE Active Discovery Offer *See* PADO.

PPPoE Active Discovery Request *See* PADR.

PPPoE Active Discovery Session *See* PADS.

PPPoE Active Discovery Termination *See* PADT.

PPPoE service name table A collection of service name tags, as defined in RFC 2516, for an access concentrator (AC) such as an E-series router. PPPoE clients use service name tags to request that an AC support certain services. Configuring PPPoE service name tables enables the AC to support multiple service name tags in addition to the empty service name tag.*See also* service name tag.

precedence level The order in which the effectiveness of CLI privilege levels of E-series routers is implemented. The CLI uses the following order of precedence:

1. Privilege level set for all commands within a mode, including modes that are accessed from another mode; for example, Global Configuration mode
2. Privilege level set for all commands that begin with the same keyword; for example, **snmp** commands
3. Privilege level set for individual commands; for example, **snmp-server community**.

prefix The first part of a BGP route, which describes a set of IP addresses that can be reached using the route. Prefixes are made possible by classless interdomain routing (CIDR).

prefix list A sequential collection of permit and deny conditions that apply to IP or IPv6 addresses. Like an access list, the router tests addresses one by one against the conditions in a prefix list. Unlike an access list, the prefix list specifies a base IP or IPv6 address and a length. The tested address is matched against the prefix.

prefix tree A nonsequential collection of permit and deny conditions that apply to IP addresses. Like a prefix list, the prefix tree specifies a base IP address and a length, the number of bits applied to the base to determine the network prefix. The tested address is matched against the prefix. The prefix tree also enables route summarization. However, the prefix tree does not match addresses one by one in sequence against the listed conditions. The router performs a binary search against the tree structure of the entries. The prefix tree provides a faster search methodology and matches the test address more closely than either the access list or the prefix list.

prepended header A header created by the policy-mirroring action during packet mirroring, and used for demultiplexing at the analyzer to sort through the multiple mirrored streams that arrive from different sources. During a packet mirroring session, the router prepends a special UDP/IP header to each mirrored packet that is sent to the analyzer port.

presentation layer The sixth level in the seven-layer OSI reference model for network protocol design. This layer transforms data to provide a standard interface for the application layer.

primary IP address An IP address configured as primary from the set of real interface addresses. VRRP advertisements are always sent (by the master router) using the primary IP address as the source of the IP packet.

primary IP interface A normal IP interface on a supported layer 2 interface, such as Ethernet. You create a primary interface by assigning an IP address to the Ethernet interface.

private line aggregation The consolidation of multiple high-speed access lines into one access point.

private community *See* local-use community.

Privileged Exec mode A User Exec mode that provides privileged-level access. Privileged Exec commands allow you to perform such functions as displaying system information, setting operating parameters, and gaining access to Global Configuration mode. *See also* User Exec mode.

privileged level A level of access in the CLI of E-series routers that enables you to view router configuration, change a configuration, and run debugging commands. You need a password to access this level. This level gives you full CLI privileges. The CLI has the ability to map any command to one of 16 levels of command privilege (in the range 0–15). When you access Privileged Exec mode, you have access to those commands that map to your access level or below.

profile A set of characteristics that act as a pattern. Defined through CLI commands to configure dynamic interfaces.

programmable read-only memory *See* PROM.

progress indicator An animated representation of how much progress has been made on a CLI operation that does not finish within the expected completion time. This type of status indicator is supported for the file system synchronization application and the file copy application. The progress indicator displays a series of dots that represents the time required to complete the operation. The dots are followed by the actual percentage of the total that has been completed and by an oscillating asterisk that indicates ongoing activity. As the application progresses, the dots are replaced with asterisks, starting at the left, to represent how much of the operation is finished.

PROM programmable read-only memory. A form of digital memory in which each bit is locked by using a fuse or antifuse action to store information permanently.

promiscuous peer group A BGP peer group that accepts incoming BGP connections from any remote address that matches an access list. Promiscuous peers are useful when the remote address of the peer is not known ahead of time. An example is in B-RAS applications, in which interfaces for subscribers are created dynamically and the remote address of the subscriber is assigned dynamically from a local pool or by using RADIUS or some other method.

protect interface A type of interface that provides the redundant connection on modules that have APS/MSP or that otherwise enable port redundancy.

Protocol Independent Multicast *See* PIM.

Protocol Independent Multicast source-specific multicast *See* PIM SSM.

protocol data unit *See* PDU.

Protocol Field Compression *See* PFC.

P router A router within a service provider core that connects directly to PE routers or other P routers and does not connect directly to a customer edge (CE) device. *See also* PE router.

provider core router *See* P router.

provider edge router *See* PE router.

proxy ARP proxy Address Resolution Protocol. A protocol that enables an E-series router to respond to ARP requests on behalf of an Ethernet end node.

PSNP partial sequence number PDU (protocol data unit). A PDU sent by designated router to acknowledge and request link-state information.

Public-Key Cryptography Standards *See* PKCS.

PVC permanent virtual circuit; permanent virtual connection (when referring to ATM). A virtual circuit that is permanently established. PVCs save bandwidth associated with circuit establishment and teardown in situations where certain virtual circuits must exist all the time.

Q

QoS quality of service. A suite of features that configure queuing and scheduling on the forwarding path of an E-series router. QoS provides a level of predictability and control beyond the best-effort delivery that the router provides by default. (Best-effort service provides packet transmission with no assurance of reliability, delay, jitter, or throughput.) *See also* CoS.

QoS administrator The individual responsible for implementing a QoS queuing architecture by defining the scheduler profiles and referencing them from QoS profiles. A QoS administrator also configures parameter definitions that control the parameters, interfaces, and ranges of values that QoS clients, using QoS parameters, can assign.

QoS client The individual responsible for configuring services for individual subscribers by creating parameter instances. The parameter instances that a QoS client creates depend on the settings that the QoS administrator defined in parameter definitions. QoS clients can use the CLI, Service Deployment System (SDX), IP multicast bandwidth adjustment, RADIUS, or Service Manager to manage these services.

QoS parameters Special parameters that enable you to configure a queuing architecture without specifying numeric subscriber rates and weights in scheduler profiles. You then use the same QoS and scheduler profiles across all subscribers who use the same services but at different bandwidths, reducing the total number of QoS profiles and scheduler profiles required.

QoS port-type profile A QoS profile that is automatically attached to ports of the corresponding type if you do not explicitly attach a QoS profile.

QoS profile A collection of QoS commands that specify queue profiles, drop profiles, scheduler profiles, and statistics profiles in combination with interface types.

QoS profile attachment A reference that applies the rules in the QoS profile to a specific interface.

quadruple play The addition of mobile phone service to triple play. *See also* triple play.

quality of service *See* QoS.

queue The first-in first-out (FIFO) set of buffers that control packets on the data path.

queue profile The template that specifies the buffering and tail-dropping behavior of an egress queue.

R

RADIUS Remote Authentication Dial-In User Service. A distributed client/server that protects networks against unauthorized access. RADIUS clients running on an E-series router send authentication requests to a central RADIUS server. The central RADIUS server stores all the required user authentication and network access information. RADIUS informs the router of the privilege levels for which RADIUS-authenticated users have enable access. The router permits or denies enable access accordingly.

RADIUS-based packet mirroring A type of packet mirroring in which a RADIUS administrator uses RADIUS attributes to configure packet mirroring of a particular user's traffic without regard to how often the user logs on or off or which E-series router or interface the user uses. Is particularly appropriate for large networks and for debugging network problems related to mobile users, who do not always log on to a particular router.

RADIUS Services

- Authentication—Determines whether or not a user is allowed to access a specific service or resource.
- Authorization—Associates connection attributes or characteristics with a specific user.
- Accounting—Tracks service use by subscribers

rate-limit hierarchy A type of rate limiting that enables lower-priority traffic to access unused bandwidth allocated for real-time traffic during times when no real-time traffic is flowing. *See also* color-aware rate limit; color-blind rate limit.

rate limiting A method of applying rate limits on bandwidth and burst size for traffic on a particular interface. *See also* one-rate rate-limit profile; two-rate rate-limit profile.

rate-limit profile A set of bandwidth attributes and associated actions. Provides a variety of services, including tiered bandwidth service where traffic conforming to configured bandwidth levels is treated differently than traffic that exceeds the configured values, and a hard-limit service where a fixed bandwidth limit is applied to a traffic flow. Also provides a TCP-friendly rate-limiting service that works in conjunction with TCP's native flow-control functionality. *See also* one-rate rate-limit profile; two-rate rate-limit profile.

rate shaping A mechanism that throttles the rate at which an interface can transmit packets.

RDI cell remote defect indication cell. A cell received from the remote endpoint of the virtual path (VP) or virtual connection (VC) that indicates an interruption in the cell transfer capability of the VP/VC.

Real-Time Streaming Protocol *See* RTSP.

receive window size *See* RWS.

redirected authentication A service that helps offload AAA activity on the router, by providing the domain-mapping-like feature remotely on the RADIUS server.

redundancy *See* line module redundancy. *See also* HA; switchover.

redundancy midplane A hardware component that provides additional connectivity so the spare line module can take control of the I/O module associated with any failed line module in the redundancy group. *See also* midplane.

relative strict-priority scheduling A process that provides strict-priority scheduling within a shaped aggregate rate.Relative strict priority differs from true strict priority in that it can implement the aggregate shaping rate for both strict and nonstrict traffic. With true strict priority, you can shape the nonstrict or the strict traffic separately, but you cannot shape the aggregate to a single rate.

relay proxy *See* DHCP relay proxy.

Remote Authentication Dial-In User Service *See* RADIUS.

remote loopback The ability to request that remote devices enter into loopback; the ability to be placed in loopback by remote devices.

remote neighbors RIP neighbors that enable the router to establish neighbor adjacencies through unidirectional interfaces, such as MPLS tunnels, rather than the standard practice of using the same interface for receipt and transmission of RIP packets. The remote neighbor can be more than one hop away through intermediate routes that are not running RIP. RIP uses the interface associated with the best route to the remote neighbor to reach the neighbor. A best route to the neighbor must exist in the IP routing table.

rendezvous point *See* RP.

requesting authority The group that is authorized to request or conduct packet mirroring (E-series routers)

Resource Reservation Protocol *See* RSVP.

Resource Reservation Protocol–Traffic Engineering *See* RSVP-TE.

resource threshold monitor *See* RTM.

Response Time Reporter *See* RTR.

reverse path forwarding *See* RPF.

RIB routing information base. A logical data structure used by BGP to store routing information. This information includes routes BGP learned from peers, local routes resulting from the application of BGP policies to the learned routes, and the routes that BGP advertises to its peers. *See also* routing table.

RIP Routing Information Protocol. An interior gateway protocol (IGP) typically used in small, homogeneous IPv4 networks. RIP uses distance-vector routing to route information based on hop count. *See also* distance-vector routing.

RIP messages Messages sent from the RIP port that contain routing information. RIP exchanges routing information by means of User Datagram Protocol (UDP) data packets. Each RIP router sends and receives datagrams on UDP port number 520, the RIP version 1/RIP version 2 port. All communications intended for another router's RIP process area are sent from the RIP port.

RIP metric The metric that RIP uses (also known as cost) to compare the value of different routes, based on hop count. The hop count is the number of routers that data packets must traverse between RIP networks. Metrics range from 0 for a directly connected network to 16 for an unreachable network. This small range prevents RIP from being useful for large networks.

root certificate The self-signed public key certificate for a root CA; root certificates are used to verify other certificates.

round-robin server access A method of access for RADIUS servers. The first configured server is treated as a primary for the first request, the second server configured as primary for the second request, and so on. When the router reaches the end of the list of servers, it starts again at the top of the list until it comes full cycle through the list. *See also* direct server access.

route flap dampening A mechanism for minimizing instability caused by route flapping. The router stores a penalty value with each route. Each time the route flaps, the router increases the penalty by 1000. If the penalty for a route reaches a configured *suppress* value, the router suppresses the route. That is, the router does not include the route as a forwarding entry and does not advertise the route to BGP peers. *See also* route flapping.

route flapping A condition of network instability where a route is announced and withdrawn repeatedly, often as the result of an intermittently failing link.

route leakage The process of allowing routes from one protocol or area to be learned by another protocol or area. Routes can be leaked into OSPF or from OSPF as follows

- Route leakage into OSPF—When another routing protocol adds a new route to the routing table, or when a static route is added to the routing table, OSPF can be informed through the **redistribute** commands. When OSPF learns the new route, it floods the information into the routing domain by using external LSAs.
- Route leakage from OSPF—OSPF adds routing information to the routing table, which is used in forwarding IP packets.

route maps Maps that modify the characteristics of a route (generally to set its metric or to specify additional attributes) as it is transmitted or accepted by a router. Route maps can use access lists to identify the set of routes to modify. Route maps control and modify routing information and define conditions for redistributing routes between routing domains.

In BGP, route maps consist of match clauses and set clauses. Match clauses specify the attribute values that determine whether a route matches a route map. Set clauses modify the specified attributes of routes that pass all match clauses in the route map.

route reflection An alternative to confederations as a strategy to reduce IBGP meshing. BGP specifies that a BGP speaker cannot advertise routes to an IBGP neighbor if the speaker learned the route from a different IBGP neighbor. In route reflection, a BGP speaker (the route reflector) advertises routes learned from each of its IBGP neighbors to its other IBGP neighbors. Routes are reflected among IBGP routers that are not meshed. *See also* cluster; confederation; route reflector; route reflector client.

route reflector A BGP speaker that advertises routes learned from each of its IBGP neighbors to its other IBGP neighbors; routes are reflected among IBGP routers that are not meshed. The route reflector's neighbors are called *route reflector clients*. The clients are neighbors only to the route reflector, not to each other. Each route reflector client depends on the route reflector to advertise its routes within the AS; each client also depends on the route reflector to pass routes to the client.

A route reflector and its clients are collectively referred to as a *cluster*. Clients peer only with a route reflector and do not peer outside their cluster. Route reflectors peer with clients and other route reflectors within a cluster; outside a cluster they peer with other reflectors and other routers that are neither clients nor reflectors. Route reflectors and nonclient routers must be fully meshed. *See also* route reflector client.

route reflector client A route reflector's neighbor. The clients are neighbors only to the route reflector, not to each other. Each route reflector client depends on the route reflector to advertise its routes within the AS;. Each client also depends on the route reflector to pass routes to the client. *See also* cluster; route reflector.

route-refresh capability A lower-cost alternative to soft reconfiguration as a means to change policies without major disruptions. The router advertises the route-refresh capability when it establishes a BGP session with a peer to indicate that it is capable of exchanging BGP route-refresh messages. *See also* cooperative route filtering; soft reconfiguration.

router ID A 32-bit number that uniquely identifies a router within an autonomous system; for example, 10.10.1.5.

route tag

- A field in a RIP message that allows boundary routers in an autonomous system (AS) to exchange information about external routes. Route tags provide a method of separating internal RIP routes (routes within the RIP routing domain) from external RIP routes, which may have been imported from an EGP (exterior gateway protocol) or another IGP (interior gateway protocol).

- In IS-IS, a numeric value assigned to the IP addresses on an IS-IS route before the route is propagated to other routers in an IS-IS domain. You can use this tag to control IS-IS route redistribution, route leaking, or route summarization by referencing it in a route map.

route target A type of BGP extended community that you use to define VPN membership. The route target appears in a field in the update messages associated with VPN-IPv4.

You create route-target import lists and route-target export lists for each VRF. The route targets that you place in a route target export list are attached to every route advertised to other PE routers. When a PE router receives a route from another PE router, it compares the route targets attached to each route against the route-target import list defined for each of its VRFs. If any route target attached to a route matches the import list for a VRF, then the route is imported to that VRF. If no route target matches the import list, then the route is rejected for that VRF.

routing domain A collection of contiguous networks that provide full connectivity to all end systems located within them. A routing domain is partitioned into areas.

routing information base *See* RIB.

Routing Information Protocol *See* RIP.

routing policy A way to control flow of routes into and out of the router. Determines how the system handles the routes it receives from and sends to neighboring routers.

In many cases, routing policy consists of filtering routes, accepting certain routes, accepting and modifying other routes, and rejecting some routes.

routing table A table maintained by IP that contains only the single best route for each protocol to a given destination. Because each protocol typically has multiple routes to a destination, the IP routing table has the one best route per protocol. *See also* global routing table; local routing table.

RP rendezvous point. For PIM sparse mode, a core router acting as the root of the distribution tree in a shared tree.

RPF reverse-path forwarding. An algorithm that checks the unicast routing table to determine whether there is a shortest path back to the source address of the incoming multicast packet. Unicast RPF helps determine the source of denial-of-service attacks and rejects packets from unexpected source addresses.

RSA Rivest-Shamir-Adleman (encryption algorithm). An algorithm for public key encryption.

RSVP Resource Reservation Protocol. A signaling protocol that establishes a session between two routers to transport a specific traffic flow. *See also* RSVP-TE.

RSVP MD5 authentication A method of authentication that provides hop-by-hop security against message spoofing and replay attacks. When authentication is configured, RSVP embeds an integrity object within secure cleartext RSVP messages sent between peers. The integrity object includes a key ID unique to the sender, a message sequence number, and keyed message digest. These attributes enable verification of both packet content and sender.

RSVP-TE RSVP-traffic engineering. RSVP with traffic engineering extensions (as defined by RFC 3209) that allow RSVP to establish label-switched paths (LSPs) in MPLS networks. *See also* MPLS; RSVP.

RTM resource threshold monitor. A CLI mode that enables you to set the rising and falling thresholds and trap hold-down times for certain interfaces. You can also view the resource threshold information.

RTR Response Time Reporter. A feature that enables you to monitor network performance and resources by measuring response times and the availability of your network devices. The primary objective of RTR is to collect statistics and information about network performance.

RTSP Real-Time Streaming Protocol. An application-level protocol for control over the delivery of data with real-time properties. RTSP provides an extensible framework to enable controlled, on-demand delivery of real-time data, such as audio and video. Sources of data can include both live data feeds and stored clips. This protocol is intended to control multiple data delivery sessions, provide a means for choosing delivery channels such as UDP, multicast UDP and TCP, and provide a means for choosing delivery mechanisms based upon RTP.

RWS receive window size. The number of packets that an L2TP peer can transmit without receiving an acknowledgment from the router. L2TP uses the RWS to implement a sliding window mechanism for the transmission of control messages. If the RWS is not configured for the L2TP tunnel, the router determines the RWS and uses this value for all new tunnels on both the L2TP access concentrator (LAC) and the L2TP network server (LNS).

S

SA security association. The set of security parameters that dictates how IPSec processes a packet. The SA defines what rules to use for authentication and encryption algorithms, key exchange mechanisms, and secure communications between two parties. A single secure tunnel uses multiple SAs. *See also* SA parameters.

SAFI subsequent address family identifier. A number that further identifies an address family identified by an AFI. In an MP-BGP update message, SAFI is used with AFI to identify the network layer protocol associated with the network address of the next hop and the semantics of the NLRI that follows. *See also* AFI.

salt encryption A random string of data used to modify a password hash.

SA parameters The actual session parameters used to secure a specific data flow associated with a specific secure IP interface. How SA parameters are set depends on how the IP interfaces are secured:

- For manual secure IP interfaces, the system administrator sets SA parameters. Manually setting SA parameters allows provisioning of IP security to destinations that do not support SA negotiation via IKE.

- For signaled secure IP interfaces, the two security gateway peers negotiate SA parameters; the system administrator cannot set any of the parameters. For some of these parameters, such as session keys, the system administrator does not have even read access.

SAR scheduler One part of the integrated scheduler used to extend ATM QoS functionality. The commercial SAR scheduler enables you to configure traditional ATM cell-based QoS. *See also* HRR scheduler.

SC system controller. A subsystem located on the SRP modules on the E320 router that controls the overall operations on the router.

scheduler hierarchy A hierarchical, tree-like arrangement of scheduler nodes and queues. The router supports up to three levels of scheduler nodes stacked above a port (level 0), with a final level of queues stacked above the nodes. A traffic-class group uses a scheduler level at level 1.

scheduler node An element within the hierarchical scheduler that implements bandwidth controls for a group of queues. Queues are stacked above scheduler nodes in a hierarchy. The root node is associated with a channel or physical port.

scheduler profile A collection of commands that configures the bandwidth at which queues drain as a function of relative weight, assured rate, and shaping rate.

scope A value used in some unicast and multicast IPv6 addresses that identifies the application suitable for the address.

SCR sustained cell rate. An upper bound on the conforming average rate of an ATM connection over a sustained time interval that is longer than the time interval for which the PCR is defined.

SCSI small computer system interface. A standard interface and command set for transferring data between devices over a computer bus.

SDH Synchronous Digital Hierarchy. An international standard defined by the International Telecommunication Union for transmitting bits over fiber-optic cable.

SDRAM synchronous dynamic random access memory. A type of RAM that is stored on dual in-line memory modules (DIMMs) and synchronized with the system clock.

SDSL symmetric digital subscriber line. A version of digital subscriber line (DSL) where the upload speeds and download speeds are the same, typically in the range 144 Kbps–1.5 Mbps. SDSL uses one cable pair and does not share lines with analog phones.

SDX software Service Deployment System (formerly SSC) software. A customizable Juniper Networks product with which service providers can rapidly deploy IP services—such as video on demand (VoD), IP television, stateful firewalls, Layer 3 VPNs, and bandwidth on demand (BoD)—to hundreds of thousands of subscribers over a variety of broadband access technologies.

secondary input policy A type of policy that evaluates conditions after a route lookup. *See also* input policy; output policy; policy.

section layer For channelized OCx/STMx interfaces, the layer that manages the transport of STS/STM frames across the physical path. This layer is responsible for frame alignment, scrambling, error detection, error monitoring, signal reception, and signal regeneration. *See also* line layer; path layer.

secure IP interfaces Virtual IP interfaces that you can configure to provide confidentiality and authentication services for the data flowing through such interfaces. The software provides these services using mechanisms created by the suite of IPSec standards established by the IETF.

secure policy A policy that is created with a mirror action and that contains information about where to forward mirrored traffic during packet mirroring. *See also* packet mirroring.

secure tunnel A virtual connection between two security gateways used to exchange data packets in a secure way. A secure tunnel is made up of a local SA and a remote SA, where both are negotiated in the context of an ISAKMP SA.

security association *See* SA.

security policy database *See* SPD.

Serial Line Address Resolution Protocol *See* SLARP.

Service Deployment System *See* SDX software.

Service line module *See* SM.

service name tag An entry in a PPPoE service name table that specifies a particular service that an access concentrator (AC), such as an E-series router, can provide to a PPPoE client. An empty service name tag of zero length indicates that any service is acceptable. *See also* PPPoE service name table.

session layer The fifth level in the seven-layer OSI reference model for network protocol design. This layer controls the dialogues and connections (sessions) between computers. It establishes, manages and terminates the connections between the local and remote application. The OSI model made this layer responsible for "graceful close" of sessions, which is a property of TCP, and also for session checkpointing and recovery, which is not usually used in the Internet protocols suite. Session layers are commonly used in application environments that make use of remote procedure calls (RPCs).

set clause Part of a route map that defines how the attributes are modified for matching routes. The set conditions apply only to routes that pass all the match conditions (or a route map with no match conditions). When a route passes all the match conditions, all set conditions are applied.

SFM switch fabric module. A module that works with the SFP module to create a shared memory fabric for the E320 router.

SFP small form-factor pluggable transceiver. A transceiver that provides support for optical or copper cables. SFPs are hot-insertable and hot-removable. *See also* XFP.

(S,G) Source (S) of the multicast packet and the destination multicast group address (G)

shared IP interface One of a group of IP interfaces that are created over the same layer 2 logical interface, which enables multiple IP interfaces to share the same logical resources.

shared local address pool A group of available addresses that enables a local address server to distribute addresses that are allocated from DHCP local server address pools within the same virtual router. The addresses are configured and managed within DHCP. Therefore, thresholds are not configured on the shared pool, but are instead managed by the referenced DHCP local server pool.

A shared local address pool references one DHCP address pool, and can then obtain addresses from the referenced DHCP address pool and from any DHCP address pools that are linked to the referenced DHCP address pool.

shared shaper constituent *See* constituent.

shared shaping A mechanism that enables dynamic sharing of logical interface bandwidth for traffic that is queued through separate scheduler hierarchies. Also called shared rate shaping. *See also* compound shared shaping; simple shared shaping.

shared tunnel-server module A type of module that supports dynamic tunnel-server ports. It provides both tunnel services and regular access services.

shim header An MPLS header, which is located (shimmed) between the layer 2 header and the data packet. The header includes the MPLS label, class-of-service information (EXP bits), a bit that indicates whether the label is at the bottom of the stack, and time-to-live bits.

shim interface AN MPLS interface stacked on a layer 2 interfaces to provide layer 2 services over MPLS, or to create local cross-connects by cross-connecting the layer 2 interface to another layer 2 interface.

shortest path first *See* SPF.

shortest-path tree *See* SPT.

short-pipe model *See* pipe (and short-pipe) model.

signaled secure IP interface An interface that negotiates an SA on demand with the remote security gateway. The remote security gateway must also support SA negotiation; otherwise the gateway drops traffic. The router keeps statistics for dropped traffic. *See also* manual secure IP interfaces.

simple authentication An authentication method in IS-IS that uses a text password (authentication key) that can be entered in encrypted or unencrypted form. The receiving router uses this authentication key to verify the packet.

simple explicit shared shaper One of four types of shared shapers, in which the weight and priority attributes of the **shared-shaping-constituent** command are ignored, because the simple shared shaper does not allocate bandwidth among constituents; instead it controls just the best-effort queue or node. *See also* compound explicit shared shaper; compound implicit shared shaper; simple implicit shared shaper.

simple implicit shared shaper One of four types of shared shapers, in which constituents are best-effort node or queues, and all nodes and queues in named traffic-class groups. *See also* compound explicit shared shaper; compound implicit shared shaper; simple explicit shared shaper.

simple shared shaping A software-assisted mechanism that measures the rate of active constituents, and can shape the best-effort node or queue associated with a logical interface to a shared rate. *See also* compound shared shaping; shared shaping.

SLA service level agreement. A formal agreement between a service provider and its customers (as part of a networking service contract) to provide a certain level of service (usually a level of performance).

SLARP Serial Line Address Resolution Protocol. A simple control protocol provided by the Cisco High-Level Data Link Control implementation that maintains serial link keepalives. *See also* Cisco HDLC.

sleep A feature of SSH that prevents a user who has exceeded the authentication retry limit from connecting from the same host within the specified period.

slot group A group of adjacent chassis (module) slots. The number of slots and number of slots per group depend on the system.

SM Service line module. A tunnel-service line module that does not pair with a corresponding I/O module that provides ingress and egress ports. Receives data from and transmits data to line modules that have ingress and egress ports.

small computer system interface *See* SCSI.

small outline dual inline memory module *See* SODIMM.

smart keepalive *See* low-density keepalive mode; high-density keepalive mode.

SMDS Switched Multimegabit Data Service. A connectionless, wide-area networking service designed for LAN interconnection. An SMDS network is composed of a series of SMDS switches inside a service provider's network, a series of channel service units/data service units (CSUs/DSUs) that connect subscribers to the network, and routers and gateways to connect to each CSU/DSU.

SNMP agent A managed device, such as a router, that collects and stores management information. The SNMP agent (SNMPv3) recognizes up to 32 usernames that can have one of the following security levels: no authentication and no privacy, authentication only, authentication and privacy.

SNMP client A device that executes management applications that monitor and control network elements. Sometimes called a network management station (NMS) or simply a manager. The SNMP client runs on a network host and communicates with one or more SNMP servers on other network devices, such as routers, to configure and monitor the operation of those network devices.

SNMP community A logical group of SNMP-managed devices and clients in the same administrative domain.

SNMP community name A name that acts as a password and is used to authenticate messages sent between an SNMP client and a router containing an SNMP server. The community name is sent in every packet between the client and the server.

SNMP event A condition or state change that might cause the generation of a trap message

SNMP group A set of users with the same access privileges to the router. Three predefined groups are available: admin, public, and private. Applies to SNMPv3.

SNMP managed object A characteristic of something that can be managed, such as a list of currently active TCP circuits in a device.

SNMP MIB A tree-structured schema that specifies the format of managed data for a device function. The goal of a MIB is to provide a common and consistent management representation for that function across networking devices. E-series routers support both standard and enterprise SNMP MIBs. *See also* Enterprise MIB; standard MIB.

SNMP notification A message that indicates a status change (equivalent to a trap in SNMPv1). Applies to SNMPv3.

SNMP privilege level A MIB access level that allows increasing levels of privilege:

- Read-only—Read-only access to the entire MIB except for SNMP configuration objects
- Read-write—Read-write access to the entire MIB except for SNMP configuration objects
- Admin—Read-write access to the entire MIB

SNMP secure packet mirroring trap A type of SNMP trap that enables the administrator to capture and report packet mirroring information to an external device. The secure information can then be viewed on the remote device. *See also* packet mirroring.

SNMP server A managed device, such as a router, that collects and stores management information. The SNMP server operates on a network device, such as a router, a switch, or a workstation. It responds to SNMP requests received from SNMP clients and generates *trap messages* to alert the clients about notable state changes in the network device. *See also* SNMP client.

SNMP Server Event Manager An application that works in conjunction with the Event MIB (RFC 2981) to allow many management functions such as fault detection, configuration management, accounting management, and performance management. These functions are traditionally performed by the network management station (NMS). However, by using the SNMP Server Event Manager, you can distribute some of these functions to E-series routers and automate them. *See also* Event MIB.

SNMP trap A message sent by an SNMP server to a client to indicate the occurrence of a significant event, such as a specifically defined condition or a threshold that was reached. Managed devices use traps to asynchronously report certain events to clients. *See also* SNMP server.

SNMP trap severity level A level of severity that an SNMP trap message can have. From most severe to least severe, the trap severity levels are Emergency, Alert, Critical, Warning, and Notice.

SNMP user An individual who accesses the router. The router may provide authentication and privacy for the user through SNMPv3. Each user is associated with a group. Applies to SNMPv3.

SNMP view The management information available to the user: read, write, or notification. Three predefined views are available for each group:

- everything—Includes all MIBs associated with the router
- user—Includes all MIBs associated with the router, except standard and enterprise MIBs used to configure SNMP operation
- nothing—Excludes all MIBs

Applies to SNMPv3.

SNTP Simple Network Time Protocol. An adaptation of the Network Time Protocol (NTP) used to synchronize computer clocks in the Internet. SNTP can be used when the ultimate performance of the full NTP implementation described in RFC-1305 is not needed or justified. When operating with current and previous NTP and SNTP versions, SNTP Version 4 involves a clarification of certain design features of NTP that allow operation in a simple, stateless remote-procedure call (RPC) mode with accuracy and reliability expectations similar to the UDP/TIME protocol described in RFC-868.

SODIMM small outline dual inline memory module. A memory module that is approximately half the size of a standard DIMM.

soft reconfiguration A way to reapply inbound policies to stored BGP routes without clearing the BGP sessions and therefore disrupting the network.

sparse mode A multicast protocol mode where routers running sparse-mode protocols forward multicast traffic only when explicitly requested to do so. *See also* dense mode.

SPD security policy database. An ordered list of policy entries that specifies what services are to be offered to IP datagrams and in what fashion. The SPD must discriminate among traffic that is afforded IPSec protection and traffic that is allowed to bypass IPSec. This applies to the IPSec protection to be applied by a sender and that must be present at the receiver. The SPD requires distinct entries for inbound and outbound traffic. For any outbound or inbound datagram, three processing choices are possible: discard, bypass IPSec, or apply IPSec.

SPF shortest path first. An algorithm used by IS-IS and OSPF to make routing decisions based on the state of network links.

split horizon A mechanism to aid in preventing routing loops when distance-vector routing protocols such as RIP are employed in broadcast networks. When split horizon is enabled, the router cannot advertise information about routes on an interface from which the information originates.

spoof checking MPLS forwarding table behavior, whereby MPLS determines that an MPLS packet received from an upstream neighbor does not contain an MPLS label that was advertised to that neighbor. The packet is dropped.

MPLS supports the following types of spoof checking:

- Router spoof checking—MPLS packets are accepted only if they arrive on an MPLS major interface that is in the same virtual router as the MPLS forwarding table.
- Interface spoof checking—MPLS packets are accepted only if they arrive on the particular MPLS major interface identified in the spoof check field.

SPT Shortest-path tree. An algorithm that builds a network topology that attempts to minimize the path from one router (the root) to other routers in a routing area.

SRP switch route processor. An ERX router module that performs system management, routing table calculations and maintenance, forwarding table computations, statistics processing, configuration storage, and other control plane functions,

SSH timeout The maximum time allowed for a user to be authenticated, starting from the receipt of the first SSH protocol packet.

stacked virtual local network *See* S-VLAN.

standalone mode *See* DHCP standalone mode.

standard MIB A MIB defined by a body such as the IETF that fosters consistency of management data representation across many vendors' networking products.

stateful access control A way to address firewall issues. After a firewall for a protocol is configured, all packets that belong to those applications, which, in turn, use that protocol, are subject to stateful monitoring. Stateful access control guards a network by allowing traffic only in the trusted direction.

stateful SRP switchover *See* high availability mode.

stateless access control A way to address firewall issues. You can use the E-series policy manager to provide solutions to access problems, such as address spoofing. E-series routers automatically provide some stateless checks as part of their normal forwarding feature set.

static interface A type of interface that is created through an existing configuration mechanism such as the command-line interface (CLI) or Simple Network Management Protocol (SNMP). *See also* dynamic interface.

static oversubscription A process that enables the router to vary queue thresholds based on the number of queues currently configured, which is relatively static. *See also* bandwidth oversubscription; dynamic oversubscription.

static translation One of two NAT methods used to assign a translated IP address. Establishes a one-to-one mapping between a local and global address. Entered as a direct configuration setting that remains in the translation table until you remove them. Used when you must initiate connections from both the inside and outside interfaces or when the translation is not subject to change. *See also* dynamic translation.

static tunnel-server port A virtual port that is always present on dedicated tunnel-server modules. No explicit configuration is required for this type of port.

statistics baseline The starting point for statistics collection after resetting protocol or application statistics and counters to zero.

statistics profile A template that specifies rate statistics and event-gathering characteristics. A statistics profile enables you to gather statistics for the rate at which packets are forwarded out of a queue and for the rate at which committed, conformed, or exceeded packets are dropped. Statistics profiles also enable you to use events to monitor the rate statistics.

strict hop In MPLS explicit routing, a next hop defined by the ingress node that is connected to the previous node in the path. *See also* loose hop.

strict-priority scheduling A process that designates the traffic class (queue) that receives top priority for transmission of its packets through a port.It is implemented with a special strict-priority scheduler node that is stacked directly above the port.

strict-source routing An MPLS routing mechanism that specifies every hop that the packet must traverse. The specified path consists of adjacent hops.

stub area An area that does not get flooded with external link-state advertisements (LSAs) but does carry intra-area and interarea routes and a default route. *See also* NSSA.

subchannel A group of T1 timeslots. Subchannel numbers are in the range 1–24 but do not necessarily correspond to DS0 timeslots. The subchannel number identifies a fractional T1 channel.

subinterface A mechanism that allows a single physical interface to support multiple logical interfaces or networks. Several logical interfaces or networks can be associated with a single physical interface. Configuring multiple virtual interfaces, or subinterfaces, on a single physical interface provides greater flexibility and connectivity on the network.

subnet addressing A type of addressing used in IP addresses. A subset of a class A, B, or C network. Subnets cannot be used with class D (multicast) addresses. *See also* IP address classes.

subnet mask A 32-bit number that is used to separate the network information from the host information in an IP address.In binary notation, a series of 1s followed by a series of contiguous 0s. The 1s represent the network number; the 0s represent the host number. Use of masks can divide networks into subnetworks by extending the network portion of the address into the host portion. Subnetting increases the number of subnetworks and reduces the number of hosts.

subscriber (client) bridge interface A type of bridge interface, where the traffic flow direction is downstream—from the server (trunk) to the client (subscriber). *See also* trunk (server) bridge interface.

subscriber interfaces An extension of a shared IP interface. Subscriber interfaces are bidirectional—they can both receive and transmit traffic, in contrast to shared IP interfaces, which are unidirectional—they can transmit but not receive traffic.

subscriber policy A set of forwarding and filtering rules that defines how to handle various packet or attribute types.

subsequent address family identifier *See* SAFI.

sustained cell rate *See* SCR.

SVC switched virtual circuit (or connection, if referring to ATM). 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.

S-VLAN stacked virtual local area network. A type of VLAN that provides a two-level VLAN tag structure, which extends the VLAN ID space to more than 16 million VLANs. Creating an S-VLAN requires the use of a second encapsulation tag. The router performs decapsulation twice, once to get the S-VLAN tag and once to get the VLAN tag. This double tagging approach enables more than 16 million address possibilities, which more than satisfies the scaling requirement for Ethernet B-RAS applications.

S-VLAN oversubscription The ability to configure up to the maximum number of S-VLANs supported on an I/O module or IOA, knowing that no more than the maximum number of supported PPP sessions can be connected to the router at any one time.

S-VLAN tunnel A special type of stacked VLAN that uses a single interface to tunnel traffic from multiple VLANs across an MPLS network. The S-VLAN tunnel enables multiple VLANs, each configured with a unique VLAN ID tag, to share a common S-VLAN ID tag when they traverse an MPLS network.

Switched Multimegabit Data Service *See* SMDS.

switched virtual circuit *See* SVC.

switch fabric module *See* SFM.

switchover In a redundant configuration, the process by which the router switches to the spare line module. During switchover, the line, circuit, and IP interfaces on the I/O module or IOAs appear to go down temporarily. The duration of the downtime depends on the number of interfaces and the size of the routing table, because the router must reload the interface configuration and the routing table from the SRP module. *See also* high availability mode.

switch route processor *See* SRP.

symmetric digital subscriber line *See* SDSL.

synchronization

- A process that prevents a redundant NVS card from overwriting saved files on the primary NVS card if the primary SRP module fails and the redundant SRP module takes control. *See also* file system synchronization mode.
- A mechanism for ensuring that a BGP speaker does not advertise routes to its EBGP peers before all the BGP routes have been redistributed into all routers within its AS that are running an IGP and are not running BGP. When BGP is not synchronized with the IGPs, the IGP routers cannot forward all traffic received from another AS. The BGP speaker cannot propagate a BGP route that it learned from a peer until an IGP route to the prefix has been installed in the BGP speaker's IP routing table.
- The method that NTP uses to ensure accurate time. There are three stages to synchronization:
  - Preliminary synchronization---A stage during which the system evaluates the initial time situation and decides how to proceed with longer-term synchronization
  - Frequency calibration---A stage that takes place the first time you use NTP or when you reboot the system. During this stage, the system evaluates the frequency error of its clock by measuring change in the offset error. A frequency calibration takes 15 minutes.
  - Progressive synchronization---A stage during which the system continues to synchronize to a server after it had established initial NTP parameters.

Synchronous Digital Hierarchy *See* SDH.

synchronous dynamic random access memory *See* SDRAM.

system controller *See* SC.

system events System changes that can be classified into log event categories and that can be used for tracking purposes.

T

table map A mechanism for applying a route map to an IS-IS route as a way to filter and manipulate route attributes before the route is added to the routing table. Issuing the JUNOSe **table-map** command (in Router Configuration mode) applies a specified route map as a policy filter on the route before the route is installed in the routing table.

TACACS Terminal Access Controller Access Control System. A security protocol that provides centralized validation of users who are attempting to gain access to a router or NAS.

TACACS+ Terminal Access Controller Access Control System Plus. An authentication method for validating users who attempt to access the router using telnet.

TACACS+ accounting service A service that enables you to create an audit trail of User Exec sessions and command-line interface (CLI) commands that have been executed within these sessions. For example, you can track user CLI connects and disconnects, when configuration modes have been entered and exited, and which configuration and operational commands have been executed.

TACACS+ host The security server on which the TACACS+ process is running. Also referred to as a TACACS+ server.

TACACS+ process A program or software running on a security server that provides AAA services using the TACACS+ protocol. The program processes authentication, authorization, and accounting requests from an NAS. When processing authentication requests, the process might respond to the NAS with a request for additional information, such as a password.

TCP Transmission Control Protocol. A data communication protocol that creates connections between hosts for the exchange of data. Guarantees packets are transmitted in their original sequence from sender to receiver.

TE traffic engineering. The ability to control the path taken through a network or portion of a network based on a set of traffic parameters (bandwidth, QoS parameters, and so on). Traffic engineering enables performance optimization of operational networks and their resources. *See also* MPLS traffic engineering; RSVP-TE.

Terminal Access Controller Access Control System (Plus) *See* TACACS; TACACS+.

TFTP Trivial File Transfer Protocol. A network application that is simpler than the File Transfer Protocol (FTP) but less capable. TFTP uses the User Datagram Protocol (UDP) to transfer small files between hosts on a network. TFTP does not support any security features.

TLV type-length-value. An element inside a data communications protocol used to encode optional information. These fields are used as follows:

- type—A 1-4 byte numeric code that indicates the kind of field that this part of the message represents
- length—A 1-4 byte field that denotes the size of the value field, typically in bytes
- value—A variable-sized set of bytes that contains the data for this part of the message

ToS type of service. The method of handling traffic using information extracted from the fields in the ToS byte to differentiate packet flows.

totally stubby area An OSPF area type that prevents Type 3, 4, and 5 link-state advertisements (LSAs) from entering the non-backbone area. Although it blocks type 3 summary LSAs from flowing into the area, type 3 LSAs carrying default route information alone are injected into the area. *See also* NSSA; stub area.

traditional NAT The most common method of using network address translation (NAT). Primary use is translating private addresses to legal addresses for use in an external network. There are two types of traditional NAT: basic NAT and NAPT. *See also* basic NAT; NAT.

traffic class A chassis-wide collection of buffers, queues, and bandwidth that you can allocate to provide a defined level of service to packets in the traffic class for JUNOSe QoS.

traffic-class group A separate hierarchy of scheduler nodes and queues over a port. Traffic classes belong to the default group unless they are specifically assigned to a named group. Organizing traffic into multiple traffic-class groups enables you to manage and shape traffic—by service class, for example—when the traffic classes are distributed across different virtual circuits. The router supports up to four traffic-class groups. A traffic class cannot belong to more than one group.

traffic engineering *See* TE.

transform sets Sets composed of security parameters that provide a required security level to a particular data flow. Transform sets are used during user SA negotiation to find common agreement between the local and the remote security gateway on how to protect that specific data flow.

A transform set includes encapsulation protocols and transforms; for example, encryption/decryption/authentication algorithms.

transient black hole A condition where a transit router running both IS-IS and BGP drops traffic because not all of the information required to reach some external destinations is yet available.

Transmission Control Protocol *See* TCP.

transparent bridge A data-link layer (layer 2) relay device that connects two or more networks or network systems. Transparent bridging is configured when you create one or more bridge groups on an E-series router. *See also* bridge group; bridge group interface.

transport layer The fourth level in the seven-layer OSI reference model for network protocol design and in the five-layer TCP/IP protocol stack. This layer provides communication between applications residing in different hosts and reliable transparent data transfer between end users. It is the first layer to address reliability.

transport virtual router For a secure IP tunnel—the VR in which both of the secure tunnel endpoints, the source and destination, are routable addresses. Normally, the transport VR is the default ISP routing infrastructure on top of which VPNs are provisioned.

trap *See* SNMP trap.

trigger A RADIUS attribute that identifies a user whose traffic is to be mirrored. Packet mirroring starts when a trigger is detected. *See also* packet mirroring.

trigger table (mteTriggerTable) An SNMP term for a table that lists any currently defined trigger conditions. Triggers fall into three categories—existence, Boolean, and threshold. One of three parts of the Event MIB. *See also* event table (mteEventTable); objects table (mteObjectsTable).

triple play The provisioning of three services (data, voice, and video) over a single broadband connection. *See also* quadruple play.

Trivial File Transfer Protocol *See* TFTP.

trunk (server) bridge interface A type of bridge interface; the traffic flow direction is upstream—from the client (subscriber) to the server (trunk). *See also* subscriber (client) bridge interface. "

trusted network An internal network (for instance, an intranet) or your personal computer. *See also* untrusted network.

TSM Tunnel Service Module. A line module, but one that does not pair with a corresponding I/O module that provides ingress and egress ports. A TSM receives data from and transmits data to line modules that have ingress and egress ports.

tunnel Generally, a private, secure path through an otherwise public network. More specifically, it is an LSP that is used by an IGP to reach a destination, or an LSP that uses traffic engineering.

Tunnel Service line module *See* TSM.

twice NAT A type of network address translation (NAT). Both the source and destination addresses are subject to translation as packets traverse the NAT router in either direction. *See also* NAT.

two-rate rate-limit profile A type of rate-limit profile that enables the user to build tiered rate-limit services and to specify different treatments for packets at different rates. *See also* one-rate rate-limit profile; rate-limit profile.

type, length, and value *See* TLV.

type of service *See* ToS.

U

U Unit. A standard unit of measurement for rack-mounted equipment (a U equals 1.75 in., or 4.44 cm).

UBR unspecified bit rate. An ATM service category that does not specify traffic-related service guarantees. Specifically, UBR does not define a per-connection negotiated bandwidth.

UDP User Datagram Protocol. In TCP/IP, a connectionless transport layer protocol that exchanges datagrams without acknowledgments or guaranteed delivery, requiring that error processing and retransmission be handled by other protocols.

UHP ultimate hop popping. When the egress router advertises the explicit null label or a non-null label to its upstream neighbor. This advertisement, performed by the signaling protocol, either LDP or RSVP-TE, ensures that all MPLS packets traversing the LSP to the egress router include a label. *See also* PHP.

ultimate hop popping *See* UHP.

unchannelized interface An interface that is not fragmented into channels.

UNI user-to-network interface. An ATM Forum specification that defines an interoperability standard for the interface between a router or an ATM switch located in a private network and the ATM switches located within the public carrier networks. Also used to describe similar connections in Frame Relay networks.

unicast address An IPv4 and IPv6 user-to-user addressing protocol used to send a datagram to a single recipient.

uniform model A tunnellng method that renders MPLS transparent to the differentiated services operation. From the diff-serv perspective, it is as if MPLS is not used. In the uniform model, if traffic conditioning is applied somewhere along the LSP, the EXP bits of the inner header must be changed at the egress when the inner header becomes the outer header (because of the pop of the outer label). *See also* pipe (and short-pipe) model.

unspecified bit rate *See* UBR.

untrusted network An external network, such as the Internet. *See also* trusted network.

User Datagram Protocol *See* UDP.

User Exec mode The CLI mode you are in after you log in to the system. By default, the commands you can execute from User Exec mode provide only user-level access. The User Exec commands allow you to perform such functions as changing terminal settings on a temporary basis, performing **ping** and **trace** commands, displaying system information, and accessing Global Configuration mode. *See also* Global Configuration mode; Privileged Exec mode; privileged level.

user level A level of access in the CLI of E-series routers that enables you to view router status. This level restricts you to User Exec mode.

user-network interface *See* UNI.

USM user-based security model. A method for providing SNMP message level security using authentication protocols and privacy protocols.

V

V.35 interface A type of interface that provides synchronous operation between data communication equipment (DCE) and data terminal equipment (DTE) for data communication over the telephone network.

variable bit rate *See* VBR.

VBR variable bit rate. An ATM service category that supports variable bit rate data traffic with average and peak traffic parameters. The VBR service category has two subcategories: VBR-NRT and VBR-RT.

VBR-NRT variable bit rate, non–real time. A subcategory of the VBR service category that is used for bursty or other non-time-sensitive transmissions. VBR-NRT guarantees minimum delay and cell loss.

VBR-RT variable bit rate, real time. A subcategory of the VBR service category that is used for time-sensitive connections such as video or voice. VBR-RT guarantees minimum delay and cell loss.

VCC virtual channel connection. A connection that uses all the addressing bits of a cell header to move traffic from one link to another. The VCC is formed by joining a series of virtual channels, which are logical circuits uniquely identified for each link of the network.

VCC cell relay encapsulation A method for the router to emulate ATM switch behavior by forwarding individual ATM cells over an MPLS pseudowire (also referred to as an MPLS tunnel) created between two ATM VCCs, or as part of a local ATM passthrough connection between two ATM 1483 subinterfaces on the same router.

VCD virtual circuit descriptor. A unique number that identifies a virtual circuit.

VCI virtual channel identifier. A 16-bit field in an ATM cell header; the VCI value is unique on a single link. *See also* VPI.

VDSL very-high-bit-rate digital subscriber line. A DSL technology providing faster data transmission over short distances, usually between 1000 and 4500 feet (300 and 1500 meters), of twisted pair copper wire. The shorter the distance, the faster the connection rate.

VE router VPLS edge device. A router that is analogous to a provider edge (PE) router in a BGP/MPLS VPN configuration, and performs similar functions.

very-high-bit-rate digital subscriber line *See* VDSL.

virtual channel connection *See* VCC.

virtual channel identifier *See* VCI.

virtual circuit descriptor *See* VCD.

virtual local area network *See* VLAN.

virtual path *See* VP.

virtual path connection *See* VPC.

virtual path identifier *See* VPI.

virtual private network *See* VPN.

virtual router *See* VR.

virtual router identifier *See* VRID.

virtual routing and forwarding instance *See* VRF.

VLAN virtual local area network. A logical group of user end stations, servers, and other network devices that appear to be on the same LAN, regardless of their physical location. VLANs allow network administrators to resegment their networks without physically rearranging the devices or network connections.

voice over Internet Protocol *See* VoIP.

VoIP voice over Internet Protocol. A protocol that enables people to use the Internet as the transmission medium for telephone calls by sending voice data in packets using the Internet Protocol instead of over traditional telephony circuits.

VP virtual path. A unidirectional logical association or bundle of VCs

VPC virtual path connection. A concatenation of VPIs between Virtual Path Terminators (VPTs). VPCs are unidirectional.

VPI virtual path identifier. An 8-bit field in the header of an ATM cell that indicates the virtual path the cell takes. *See also* VCI.

VPLS domain A set of VPLS edge routers running VPLS instances that participate in that domain. Typically is associated with customers who want to use Ethernet-based layer 2 VPNs to connect geographically dispersed sites in their organization across an MPLS-based service provider core, also known as an MPLS backbone. To provide signaling for VPLS, BGP builds a full mesh of label-switched paths (LSPs) among all of the VPLS instances on each of the VPLS edge routers participating in a particular VPLS domain.

VPLS instance A new or existing bridge group that has additional VPLS attributes configured. A single VPLS instance is analogous to a distributed learning bridge (also known as a bridge group) used for transparent bridging, and performs similar functions.

A bridge group is a collection of bridge interfaces stacked on Ethernet layer 2 interfaces to form a broadcast domain. Similarly, a VPLS instance is a collection of network interfaces stacked on Ethernet layer 2 interfaces that transmits packets between the router, or VE device, and the CE device located at the edge of the customer's network. In addition, the VPLS virtual core interface enables a VPLS instance to forward traffic not only between bridge interfaces, like a bridge group, but also between a bridge (network) interface and the service provider core.

VPN virtual private network. A private data network that uses a public TCP/IP network, typically the Internet, while maintaining privacy with a tunneling protocol, encryption, and security procedures.

VP tunneling Tunneling that enables traffic shaping to be applied to the aggregation of all VCs within a single virtual path. Thus, VP tunnels can be used to ensure that the total traffic transmitted on a VP does not exceed the specified peak cell rate.

VR virtual router. Multiple distinct logical routers within a single router, which enables service providers to configure multiple, separate, secure routers within a single chassis. Each virtual router has its own separate set of IP interfaces, forwarding table, and instances of routing protocols. Applications for this function include the creation of individual routers dedicated to wholesale customers, corporate virtual private network (VPN) users, or a specific traffic type.

VRF VPN routing and forwarding instance; sometimes known as a virtual router and forwarding instance. A VRF exists within the context of a VR. VRFs are used to create VPNs. In this case, the VRF forwarding table includes only routes to sites that have at least one VPN in common with the site that is associated with the VRF.The router looks up a packet's destination in the VRF associated with the interface on which the packet is received. In general, any application that can be enabled in a VR can be enabled in a VRF.

VRID virtual router identifier. A number in the range 1—255 that identifies a VRRP instance.

VRRP router A router that is running VRRP. It might participate in one or more virtual router IDs (VRIDs).

W

warm restart The result of a redundant, standby SRP module becoming active when high availability (HA) is configured.

- The line modules remain enabled and forwarding remains active.
- The newly active SRP module recovers dynamic state information from mirrored storage.

BGP and other routing protocols typically use graceful restart to avoid route flapping during an SRP warm restart. *See also* cold restart; graceful restart.

weight
- In BGP, a preference for a particular route over other routes to a destination. The higher the assigned weight, the more preferred the route. By default, the route weight on E-series routers is 32768 for paths originated by the router, and 0 for other paths.
- In QoS, a data unit that specifies the relative weight for queues in the traffic class.

weighted random early detection *See* WRED.

weighted round-robin *See* WRR.

WEP Wired Equivalent Privacy protocol. A security protocol that encrypts data exchanged on wireless networks. Defined in the original IEEE 802.11 standard.

Wired Equivalent Privacy *See* WEP.

wireless local area network *See* WLAN.

WLAN wireless local area network. A type of LAN in which mobile users can connect to the network through a wireless (radio) connection. The IEEE 802.11 standard specifies the technologies for wireless LANs, including the Wired Equivalent Privacy (WEP) encryption algorithm.

working interface An interface that provides the primary connection on modules that have APS/MSP or that otherwise enable redundancy

WRED weighted random early detection. A congestion avoidance technique that signals end-to-end protocols such as TCP that the router is becoming congested along a particular egress path. The intent is to trigger TCP congestion avoidance in a random set of TCP flows before congestion becomes severe and causes tail dropping on a large number of flows.

WRR weighted round-robin. A scheme used to decide the queue from which the next packet is to be transmitted.

X

X.21 interface A type of interface that provides synchronous operation between data communication equipment and data terminal equipment on public data networks.

xDSL A combined term used to refer to ADSL, HDSL, SDSL, and VDSL.

XDR External Data Representation Standard. A standard for the description and encoding of data. XDR can be used to transfer data between computers.

XFP 10-gigabit small form-factor pluggable transceiver. A type of hot-swappable optical transceiver. *See also* SFP.

Part Number: 162-01855-00,
Revision A00

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA

Phone 408 745 2000
or 888 JUNIPER
Fax 408 745 2100

JUNOSe™ Software for E-series™ Routing Platforms Glossary

Writing: Brian Wesley Simmons, Bruce Gillham, Diane Florio, Fran Singer, John Borelli, Mark Barnard, Sarah Lesway-Ball
Editing: Ben Mann, Fran Mues

Revision History
29 February 2008—Revision 1