Understanding Broadband Access Aggregation

First Published: May 2, 2005
Last Updated: May 2, 2005

Broadband access aggregation is the means by which connections are made among multiple technologies. These technologies include ISDN, DSL, cable, Ethernet, and wireless devices that are connected to corporate virtual private networks (VPNs), third-party applications, and the Internet. Subscriber demand for high-speed services, including multi-player gaming, video-on-demand, home security, digital audio, streaming video, and many other applications, require the delivery of IP services, regardless of the access medium.

Because so many different technologies are involved in broadband access aggregation, it is important that the service provider understand their network both in terms of the hardware that makes up the installation, which determines what type of sessions need to be established, but also in terms of what kinds of services their subscribers expect to receive. The demands placed on large service provider installations can often result in the need to contend with millions of sessions and provide flexible and reliable configurations for widely diverse consumer needs.

This module contains conceptual information about broadband access aggregation.

Contents

Information About Broadband Access Aggregation

To perform the necessary tasks in providing broadband access aggregation, you should understand the following concepts:

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

Encapsulation Protocols

Internet access has evolved from dialup modems to high-speed broadband. One of the most important considerations in setting up a broadband network is encapsulation. The key protocols include the tunneling protocol and the transport protocol. The tunneled protocol (the one to be encapsulated) gains one or more headers that can be used to identify different tunnels between a pair of devices and ultimately deliver the payload to a remote peer.

Tunneling protocols can be applied to protocols operating at the same layer of the Open Systems Interconnection (OSI) model or at different layers. A wide range of applications can be derived from various tunneling protocols, such as connecting isolated network segments, nondisruptive network renumber, Layer 2 transport, security, and controlling routing behavior.

Layer 2 Tunneling Protocol

Layer 2 Tunneling Protocol (L2TP) is one of the most used building blocks for broadband networks. It is an International Engineering Task Force (IETF) standard that combines aspects of two existing tunneling protocols: Cisco Layer 2 Forwarding (L2F) and Microsoft Point-to-Point Tunneling Protocol (PPTP).

The main component of L2TP is a reliable control channel that is responsible for session setup, negotiation, and teardown, and a forwarding plane that adds negotiated session IDs and forwards traffic. Layer 2 circuits terminate in a device called an L2TP access concentrator (LAC), and the PPP sessions terminate in an L2TP network server (LNS). The LNS authenticates the user and is the endpoint for PPP negotiation.

The LAC connects to the LNS using a LAN or a wide-area network (WAN) switch as a public or private ATM as shown in Figure 1. The LAC directs the subscriber session into L2TP tunnels based on the domain of each session. The LAC acts as one side of an L2TP tunnel endpoint and is a peer to the LNS on the other side of the tunnel. The LAC forwards packets to and from the LNS and a remote system.

Figure 1 L2TP Tunnel Between an LAC and an LNS

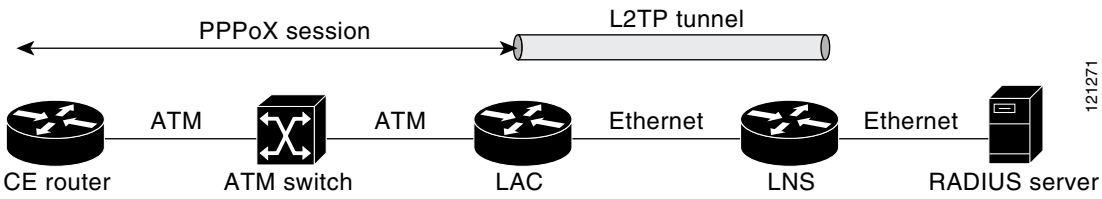

The LNS is a peer to the LAC and sits on one side of an L2TP tunnel. The LNS routes packets to and from the LAC and a destination network. The broadband aggregation server can be configured to terminate the PPP sessions and route client IP packets onto the ISP network toward their final destination. LNSs can also be configured to place sessions in VRFs before routing the packets.

The following user encapsulations can go into an L2TP tunnel:

- PPP sessions encapsulated in L2TP tunnels (LNS-side support only)
- PPPoE termination over ATM
- PPPoA termination
- PPPoEoE or PPPoEo802.1Q

Cisco's broadband aggregation routers function as the service provider's network access server when configured as the LAC. Subscribers can use a local or PPP connection to initiate a PPPoA or PPPoE session to the LAC. The LAC terminates the physical connection and forwards the PPP session to the provider's LNS.

ATM Services

ATM networks provide the following ATM services, which provide delivery of the subscriber sessions to the service providers access concentrators:

- Permanent virtual circuits (PVC)
- Switched virtual circuits (SVC)

A PVC allows direct connectivity between sites. In this way a PVC is similar to a leased line. PVCs generally guarantee availability of a connection, and no call setup procedures are required between ATM switches. However, PVCs provide a static connectivity and require manual administration to set up.

An SVC is created and released dynamically and remains in use only as long as data is being transferred. In this way it is similar to a telephone call. Dynamic call control requires a signaling protocol between the ATM endpoint and the ATM switch. SVCs provide connection flexibility and call setup that can be automatically handled by a networking device. Setting up the connection requires extra time and overhead.

ATM supports two types of connections:

- Point-to-point
- Point-to-multipoint

A point-to-point ATM connection connects two ATM end systems and can be unidirectional (one-way communication) or bidirectional (two-way communication).

A point-to-multipoint ATM connection connects a single source end-system (known as the Root node) to multiple destination end-systems (known as leaves). Such connections are unidirectional only. Root nodes can transmit to leaves, but leaves cannot transmit to the root or to each other on the same connection.

PPPoE

PPP over Ethernet (PPPoE) provides the ability to connect hosts on a network over a simple bridging device to a remote aggregation concentrator. PPPoE is the predominant access protocol in broadband networks worldwide. PPPoE typically is deployed with a software stack housed on the end-customer's (subscriber's) PC. This software allows the network service provider to "own" the customer as the PPP session runs from the customer PC to the service provider network.

PPPoEoE/PPPoEo802.1q

PPPoEoE is a variant of PPPoE where the Layer 2 transport protocol is now Ethernet or 802.1q VLAN instead of ATM. This encapsulation method is generally found in Metro Ethernet or Ethernet digital subscriber line access multiplexer (DSLAM) environments. The common deployment model is that this encapsulation method is typically found in multi-tenant buildings or hotels. By delivering Ethernet to the subscriber, the available bandwidth is much more abundant and the ease of further service delivery is increased.

PPPoA

With PPP over ATM (PPPoA), a customer premises equipment (CPE) device encapsulates the PPP session based on RFC 1483 for transport across the DSLAM. PPPoA is commonly used in SOHO and branch office type environments although it is not limited to them. It has greater flexibility for the home than the average PPPoE deployment because the customer LAN behind the CPE is under the complete control of the customer and the CPE acts as a router as opposed to a bridge for PPPoE.

When you configure PPP over ATM, a logical interface known as a *virtual access interface* associates each PPP connection with an ATM VC. You can create this logical interface by configuring an ATM PVC or SVC. This configuration encapsulates each PPP connection in a separate PVC or SVC, allowing each PPP connection to terminate at the router ATM interface as if received from a typical PPP serial interface.

The virtual access interface for each VC obtains its configuration from a virtual interface template (virtual template) when the VC is created. Before you create the ATM VC, it is recommended that you create and configure a virtual template as described in the "Preparing for Broadband Access Aggregation" module.

Once you have configured the router for PPP over ATM, the PPP subsystem starts and the router attempts to send a PPP configure request to the remote peer. If the peer does not respond, the router periodically goes into a "listen" state and waits for a configuration request from the peer. After a timeout (typically 45 seconds), the router again attempts to reach the remote router by sending configuration requests.

The virtual access interface remains associated with a VC as long as the VC is configured. If you remove the configuration of the VC, the virtual access interface is marked as deleted. If you shut down the associated ATM interface, you will also cause the virtual access interface to be marked as down (within 10 seconds), and you will bring the PPP connection down. If you set a keepalive timer of the virtual template on the interface, the virtual access interface uses the PPP echo mechanism to verify the existence of the remote peer.

The following three types of PPP over ATM connections are supported:

- IETF-compliant MUX encapsulated PPP over ATM
- IETF-compliant LLC encapsulated PPP over ATM
- Cisco-proprietary PPP over ATM

Routed Bridge Encapsulation

ATM routed bridge encapsulation (RBE) is used to route IP over bridged RFC 1483 Ethernet traffic from a stub-bridged LAN.

The ATM subinterface on a head-end router is configured to function in ATM routed-bridge encapsulation mode. This configuration is useful when a remote bridged Ethernet network device needs connectivity to a routed network by way of a device bridging from an Ethernet LAN to an ATM RFC 1483 bridged encapsulation.

The bridged ATM interface supports ATM PVCs and ATM SVCs.

Because PVCs are statically configured along the entire path between the end systems, it would not be suitable to route bridged encapsulated traffic over them when the user wants to configure the VCs dynamically and tear down the VCs when there is no traffic.

The Subscriber Service Switch was developed in response to a need by Internet service providers for increased scalability and extensibility for remote access service selection and Layer 2 subscriber policy management. This Layer 2 subscriber policy is needed to manage tunneling of PPP in a policy-based bridging fashion.

Cisco Subscriber Service Switch

The Cisco Subscriber Service Switch provides flexibility on where and how many subscribers are connected to available services and how those services are defined. In the past, remote access service selection was largely determined by the telephone number dialed or the PPP username and password entered during a PPP authentication cycle. However, emerging broadband, cable, virtual private network (VPN), and wireless access methods have created an environment where PPP sessions may be tunneled over a variety of protocols and media. The multitude of protocols, management domains, network infrastructure, and variety of services has created a complex environment for directing a subscriber to a given service or application. The problem is further complicated by the much greater density of total PPP sessions that can be transported over shared media versus traditional point-to-point links. The Subscriber Service Switch can provide a flexible and extensible decision point linking an incoming subscriber (typically a PPP session over some physical or virtual link) to another tunneled link or local termination for Layer 3 processing.

The Subscriber Service Switch is also scalable in situations where a subscriber's Layer 2 service is switched across virtual links. Examples include switching between PPPoA, PPPoE, L2TP, Layer 2 Forwarding Protocol (L2F), Point-to-Point Tunneling Protocol (PPTP), generic routing encapsulation (GRE), and General Packet Radio Service (GPRS) Tunneling Protocol (GTP wireless data standard).

As networks grow beyond the campus, network security increases in importance and administration complexity. Customers need to protect networks and network resources from unauthorized access by remote users. Cisco Systems uses a strategy known as authentication, authorization, and accounting (AAA) for verifying the identity of, granting access to, and tracking the actions of remote users. In today's networks, the TACACS+ and RADIUS protocols are commonly used to provide AAA solutions. Support for RADIUS along with TACACS+ enables Cisco to deliver tremendous flexibility and choice to organizations in AAA functionality.

RADIUS Support in Cisco IOS

Cisco Systems introduced support for RADIUS in Cisco IOS Release 11.1 in its network access server (NAS) devices.

The RADIUS protocol is an access server authentication and accounting protocol. RADIUS has gained support among a wide customer base, including Internet service providers (ISPs).

The RADIUS protocol is based on a client/server model. An NAS operates as a client of RADIUS. The client is responsible for passing user information to a designated RADIUS server and then acting on the response that is returned.

A RADIUS server (or daemon) can provide authentication and accounting services to one or more client NAS devices. RADIUS servers are responsible for receiving user connection requests, authenticating users, and then returning all configuration information necessary for the client to deliver service to the users. A RADIUS access server is generally a dedicated workstation connected to the network.

Additional References

Related Documents

Related Topic	Document Title
Configuring a PVC range	"Configuring PVCs" in the *Cisco IOS Wide-Area Networking Configuration Guide*
Creating a virtual template	"Configuring Virtual Template Interfaces" chapter in the *Cisco IOS Dial Technologies Configuration Guide*

Standards

Standards	Title
None	—

MIBs

MIBs	MIBs Link
None	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFCs	Title
1483	Multiprotocol Encapsulation over ATM Adaptation Layer 5

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

Glossary

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

ACR—allowed cell rate. A parameter defined by the ATM Forum for ATM traffic management. ACR varies between the MCR and the PCR, and is controlled dynamically using congestion control mechanisms.

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

MCR—minimum cell rate. Parameter defined by the ATM Forum for ATM traffic management. MCR is defined only for ABR transmissions, and specifies the minimum value for the ACR.

NAS—network access server. Cisco platform (or collection of platforms, such as an AccessPath system) that interfaces between the packet world (for example, the Internet) and the circuit world (for example, the PSTN).

PCR—peak cell rate. Parameter defined by the ATM Forum for ATM traffic management. In Constant Bit Rate (CBR) transmissions, PCR determines how often data samples are sent. In ABR transmissions, PCR determines the maximum value of the ACR.

PPP—Point-to-Point Protocol. PPP is the successor to Serial Line Internet Protocol (SLIP) that provides router-to-router and host-to-network connections over synchronous and asynchronous circuits. Whereas SLIP was designed to work with IP, PPP was designed to work with several network layer protocols, such as IP, IPX, and ARA. PPP also has built-in security mechanisms, such as CHAP and PAP. PPP relies on two protocols: Link Control Protocol (LCP) and Network Control Protocol (NCP).

PPPoA—Point-to-Point Protocol over ATM. The PPPoA feature enables a high-capacity central site router with an Asynchronous Transfer Mode (ATM) interface to terminate multiple remote Point-to-Point Protocol (PPP) connections.

PPPoE—Point-to-Point Protocol over Ethernet. PPPoE allows a PPP session to be initiated on a simple bridging Ethernet connected client.

PPPoX—Point-to-Point Protocol over Protocol. PPPoX indicates that the point-to-point protocol terminates on another protocol which could be ATM or Ethernet.

PVC—permanent virtual circuit. A virtual circuit that is permanently established. PVCs save bandwidth associated with circuit establishment and tear down in situations where certain virtual circuits must exist all the time. In ATM terminology, called a permanent virtual connection.

QoS—quality of service. Cisco IOS QoS technology lets complex networks control and predictably service a variety of networked applications and traffic types.

RADIUS—Remote Authentication Dial-in User Service

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

UBR—unspecified bit rate. QoS class defined by the ATM Forum for ATM networks. UBR allows any amount of data up to a specified maximum to be sent across the network but there are no guarantees in terms of cell loss rate and delay.

VBR—variable bit rate. QoS class defined by the ATM Forum for ATM networks. VBR is subdivided into a real time (rt) class and non-real time (nrt) class. VBR (rt) is used for connections in which there is a fixed timing relationship between samples. VBR (nrt) is used for connections in which there is no fixed timing relationship between samples but that still need a guaranteed QoS.

VPDN—virtual private dialup network. A VPDN is a network that extends remote access to a private network using a shared infrastructure. VPDNs use Layer 2 tunnel technologies (L2F, L2TP, and PPTP) to extend the Layer 2 and higher parts of the network connection from a remote user across an ISP network to a private network. VPDNs are a cost effective method of establishing a long distance, point-to-point connection between remote dial users and a private network.

VSA—vendor-specific attribute. An attribute that has been implemented by a particular vendor. It uses the attribute Vendor-Specific to encapsulate the resulting AV pair: essentially, Vendor-Specific = protocol:attribute = value.

Note See *Internetworking Terms and Acronyms* for terms not included in this glossary.