Cisco – Cisco – PPPoE Baseline Architecture for the Cisco UAC

Table of Contents

Cisco − PPPoE Baseline Architecture for the Cisco UAC 6400

Introduction

This document describes an end−to−end Asymmetric Digital Subscriber Line (ADSL) architecture using Point−to−Point Protocol over Ethernet (PPPoE).

In the current environment of access technologies, it is highly desirable to connect multiple hosts at a remote site through the same customer premise access device. It is also essential to provide access control and billing functionality in a manner similar to dialup services using Point−to−Point Protocol (PPP). In many access technologies, the most cost−effective method for attaching multiple hosts to the customer premise access device is via Ethernet. In addition, it is desirable to keep the cost of this device as low as possible while requiring little or no configuration.

As customers deploy ADSL they must support PPP−style authentication and authorization over a large installed base of legacy bridging customer premises equipment (CPE). PPPoE provides the ability to connect a network of hosts over a simple bridging access device to a remote access concentrator or aggregation concentrator. With this model, each host uses its own PPP stack, thus presenting the user with a familiar user interface. Access control, billing, and type of service can be done on a per user, rather than a per site, basis.

Assumption

The baseline architecture assumes that the following are provided:

- High speed Internet access and corporate access to the end subscriber using PPPoE
- ATM as the core backbone technology, implemented by the Cisco 6400 Universal Access Concentrator (UAC)

This design implementation restriction may limit use of this architecture on other platforms, but PPPoE is constantly evolving. Read the latest release notes for related products to take advantage of new and updated features.

This paper is based on existing deployments as well as inhouse tests using the Cisco 6400 UAC. This paper is a continuation of the PPPoA Baseline Architecture paper and refers to it often. It is assumed that you have

read the PPPoA Baseline Architecture white paper and understand the fundamentals of PPP, and that you have read release notes for the latest software release.

Technology Brief

As specified in RFC 2516, PPPoE has two distinct stages: a discovery stage and a PPP session stage. When a host initiates a PPPoE session, it must first perform discovery to identify which server can meet the client's request, then identify the Ethernet MAC address of the peer and establish a PPPoE session id. While PPP defines a peer−to−peer relationship, discovery is inherently a client−server relationship.

In the discovery process, a host (the client) discovers one or more access concentrators (the servers) and selects one. When discovery completes successfully, both the host and the selected access concentrator have the information to build their point−to−point connection over Ethernet. After a PPP session is established, both the host and the access concentrator must allocate the resources for a PPP virtual interface (this may not be the case for all implementations). For more details on the PPPoE specification, please see RFC 2516.

Advantages and Disadvantages of PPPoE Architecture

PPPoE architecture inherits most of the advantages of PPP used in the dialup model and in PPPoA architecture. Some key advantages and disadvantages of PPPoE and how they differ from PPPoA are listed below.

Advantages

Some key advantages of PPPoE and how they differ from PPPoA include:

- Per session authentication based on Password Authentication Protocol (PAP) or Challenge Handshake Authentication Protocol (CHAP). This is the greatest advantage of PPPoE as authentication overcomes the security hole in a bridging architecture.
- Per session accounting is possible, which allows the service provider to charge the subscriber based on session time for various services offered. The service provider may also require a minimal access charge.
- PPPoE can be used on existing CPE installations that cannot be upgraded to PPP or that do not have the ability to run PPPoA, extending the PPP session over the bridged Ethernet LAN to the PC.
- PPPoE preserves the point−to−point session used by Internet Service Providers (ISPs) in the current dialup model. PPPoE is the only protocol capable of running point−to−point over Ethernet without requiring an intermediate IP stack.
- The Network Access Provider (NAP) or Network Service Provider (NSP) can provide secure access to a corporate gateway without managing end−to−end permanent virtual circuits (PVCs) and making use of Layer 3 routing and/or Layer 2 Tunneling Protocol (L2TP) tunnels. This makes the business model of selling wholesale services and virtual private networks (VPNs) scalable.
- PPPoE can provide a host (PC) access to multiple destinations at a given time. There can be multiple PPPoE sessions per PVC.
- The NSP can oversubscribe by deploying idle and session time−outs using an industry standard Remote Authentication Dial−In User Service (RADIUS) server for each subscriber.
- PPP can be used with the service selection gateway (SSG) feature.

Disadvantages

Some key disadvantages of PPPoE and how they differ from PPPoA include:

- PPPoE client software must be installed on all hosts (PCs) connected to the Ethernet segment. This means that the access provider must maintain the CPE and the client software on the PC.
- Because PPPoE implementation uses RFC1483 bridging, it is susceptible to broadcast storms and possible denial−of−service attacks.

Implementation Considerations for PPPoE Architecture

Following are some key points to consider before implementing this type of architecture.

- The number of subscribers that will be supported. The number of PPPoE servers required depends on the number of sessions.
- Whether the PPP sessions are being terminated at the service provider's aggregation router or forwarded to other corporate gateways or ISPs.
- Whether the service provider or the final service destination is providing the IP address.
- In the case of more than one user, whether all users need to reach the same final destination or service, or they all have different service destinations. Do the end subscribers require simultaneous access to multiple destinations?
- The PPPoE client software that the access provider will use and whether the software has been tested, the operating system that the host will use, and whether that operating system can make an intelligent routing decision.
- How the service provider bills subscribers−based on a flat rate, per session usage, or services used.
- Deployment and provisioning of CPEs, DSLAMs and aggregation points of presence (POPs).
- The business model for the NAP. Does the model also include selling wholesale services like secure corporate access and value added services like voice and video? Are NAPs and NSPs the same entity?
- The business model of the company. Is it comparable to an independent local exchange carrier (ILEC), a competitive local exchange carrier (CLEC) or an ISP?
- The types of applications the NSP will offer to the end subscriber.
- The anticipated upstream and downstream volume of data flow. Consider NRP throughput, traffic engineering, and any QoS issues.

Keeping these points in mind, we will discuss how the PPPoE architecture will fit and scale to different business models for service providers and how the providers can benefit using this architecture.

Network Architecture

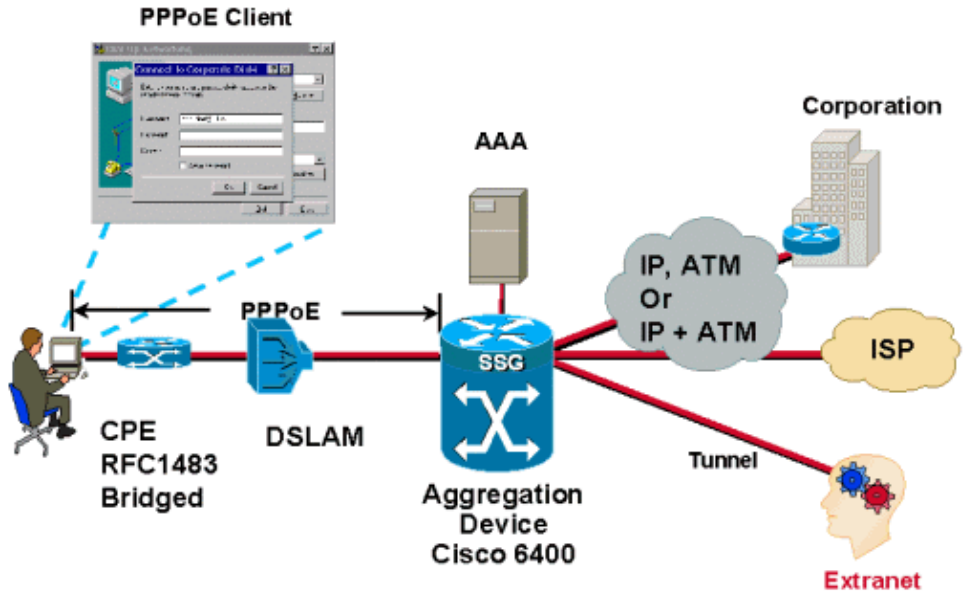

Design Considerations for PPPoE Architecture

This section covers issues that apply specifically to PPPoE Architecture.

Before deploying any architecture, it is essential to understand the business model of the service provider and what services the provider will offer. You should also know the client software being used on the PC. The most common software is from Routerware. Because the client software will be installed on a PC, the service provider technician should have a good understanding of that PC and its operating system.

As specified in RFC 2516, the maximum receive unit (MRU) option must not be negotiated to a size larger than 1492; Ethernet has a maximum payload size of 1500 octets. The PPPoE header is 6 octets and the PPP protocol ID is 2 octets, so the PPP maximum transmission unit (MTU) must not be greater than 1492. This can be achieved by configuring IP mtu 1492 for PPPoE virtual−template interfaces.

By default, no virtual access interface is precloned when a PPPoE VPDN group is configured. Users can change the maximum number of precloned virtual access interfaces by issuing the **virtual−template <number> pre−clone <number>** global command.

To protect the router against denial−of−service attacks, PPPoE by default allows only one session to be sourced from a MAC address over a VC. Users can change the defaults by issuing the pppoe session−limit per−mac and pppoe session−limit per−vc commands.

The accounting, authorization, and authentication process is the same as that of PPPoA. The only difference is that currently the VPI/VCI−based authentication, which is available for PPPoA and not available for PPPoE, can use the L2TP and SSG architectures for wholesale services.

Key Points of PPPoE Architecture

CPE

The CPE is configured for pure RFC1483 bridging. Each CPE will consume only one VPI/VCI pair and all PPPoE sessions initiated by hosts behind this CPE will be carried over in this single VC.

IP Management

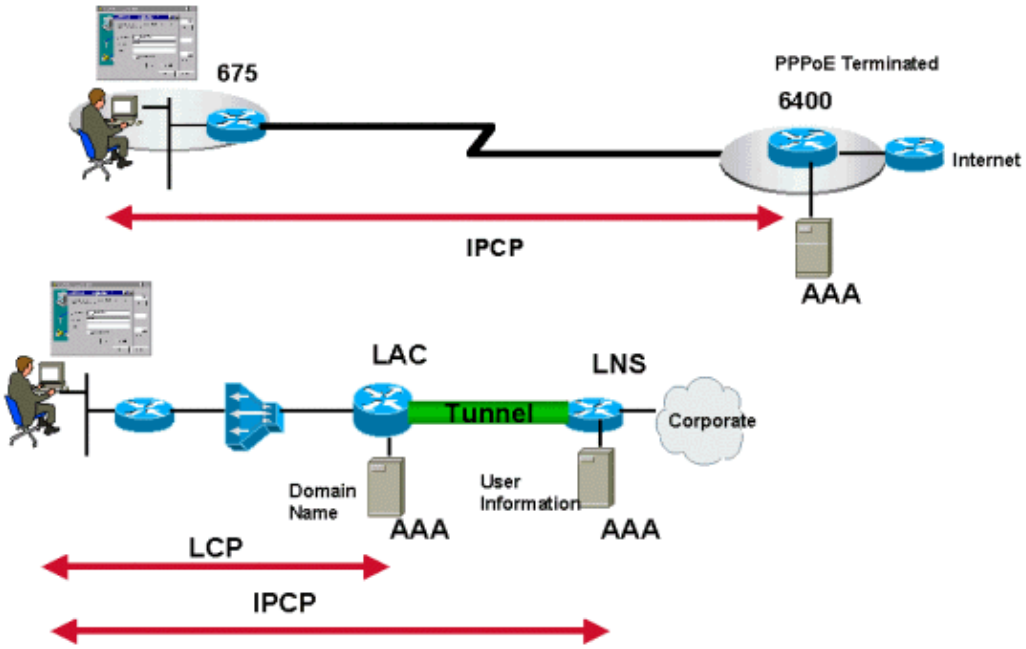

The IP address allocation for the individual host running the PPPoE client is based on the same principle of PPP in dial mode−IPCP negotiation. The IP address origin depends on the type of service the subscriber has purchased and where the PPP sessions are terminated. PPPoE makes use of the dialup networking feature of Microsoft Windows, and the IP address assigned is reflected in the PPP adapter.

The IP address assignment could come from the access concentrator terminating the PPPoE sessions or in the case of L2TP, from the home gateways. The IP address is assigned for each PPPoE session.

The CPE cannot do Network Address Translation/ Dynamic Host Configuration Protocol (NAT/DHCP) because it is bridged and there is no IP address allocated to it.

How the Service Destination is Reached

The service destination can be reached in the following ways:

- Terminating PPP sessions at the service provider
- L2TP tunneling
- Using SSG

Detailed explanations of these architectures will be covered in separate papers.

Operational Description of PPPoE

This release of PPPoE client software supports the discovery and session stages described in RFC 2516. There are four steps to the discovery stage. When it completes, both peers know the PPPoE session id and the peer's

Ethernet address, which together uniquely define the PPPoE session. The steps are:

1. The host broadcasts an initiation packet.

   The host sends the PPPoE active discovery initiation (PADI) packet with the destination_addr set to the broadcast address; the PADI consists of one tag indicating what service type it is requesting.
2. One or more access concentrators send offer packets.

   When the access concentrator or the router receives a PADI that it can serve, it replies by sending a PPPoE active discovery offer (PADO) packet; the destination_addr is the unicast address of the host that sent the PADI. If the access concentrator cannot serve the PADI, it must not respond with a PADO. Because the PADI was broadcast, the host may receive more than one PADO.

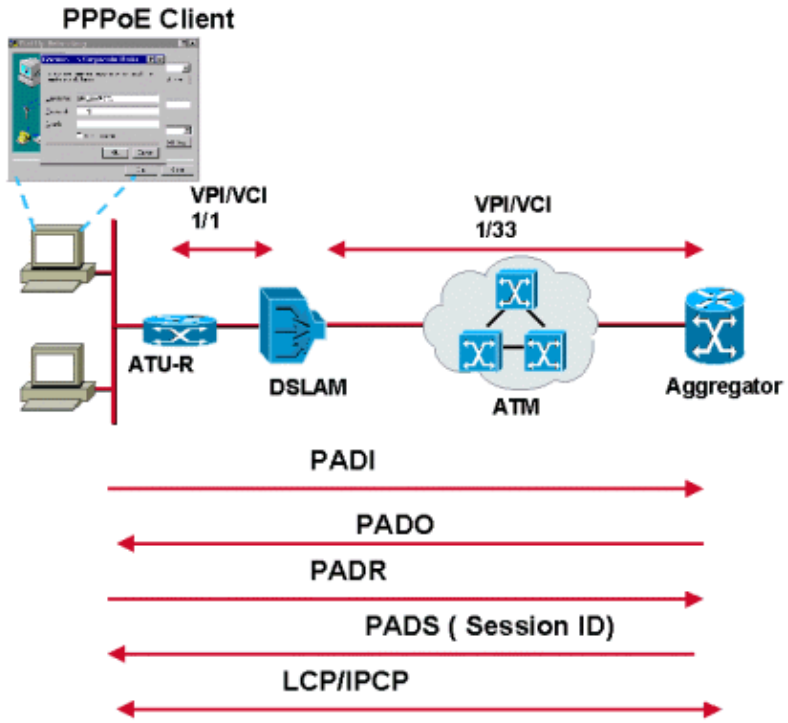

3. The host sends a unicast session request packet.

   The host looks through the PADO packets it receives and chooses one. The choice can be based on the services offered by each access concentrator. The host then sends one PADR packet to the access concentrator it has chosen. The destination_addr field is set to the unicast Ethernet address of the access concentrator or the router that sent the PADO.
4. The selected access concentrator sends a confirmation packet.

   When the access concentrator receives a PADR packet, it prepares to begin a PPP session. It generates a unique session id for the PPPoE session and replies to the host with a PPPoE active discovery session−confirmation (PADS) packet. The destination_addr field is the unicast Ethernet address of the host that sent the PADR.

Once the PPPoE session begins, PPP data is sent as in any other PPP encapsulation. All Ethernet packets are unicast.

A PPPoE active discovery terminate (PADT) packet may be sent by either the host or the access concentrator

any time after a session is established to indicate that a PPPoE session has been terminated.

For a more detailed explanation, please refer to RFC 2516.

Conclusion

For ADSL, PPPoE is gaining in popularity, second only to PPPoA.

References

- RFC 2516 A Method for Transmitting PPP over Ethernet (PPPoE)
- RFC 1483 Multiprotocol Encapsulation over ATM Adaptation Layer 5
- RFC 2364 Point−to−Point over AAL5

Related Information

- **PPPoA Baseline Architecture**
- **DSL (Digital Subscriber Line) Technical Support**
- **Technical Support − Cisco Systems**