

Cisco – PPPoE Baseline Architecture for the Cisco UAC 6400

Document ID: 12915

Contents

Introduction

Assumption

Technology Brief

Advantages and Disadvantages of PPPoE Architecture

Advantages

Disadvantages

Implementation Considerations for PPPoE Architecture

Key Points of PPPoE Architecture

Conclusion

References

Related Information

Introduction

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

In the current environment of Access technologies, it is 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 that use Point-to-Point Protocol (PPP). In many Access technologies, the most cost-effective method to attach 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 and the configuration requirement less or none.

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. Therefore, it presents the user with a familiar user interface. You can access control, billing, and type of service on a per user, rather than a per site, basis.

Assumption

The baseline architecture assumes that these items are provided:

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

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

This paper is based on current deployments as well as inhouse tests that use 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 in order to identify which server can meet the request of the client. Secondly, it needs to 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 in order 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 is probably not the case for all implementations). For more details on the PPPoE specification, refer to 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. These sections list some key advantages and disadvantages of PPPoE and how they differ from PPPoA.

Advantages

These are some key advantages of PPPoE and how they differ from PPPoA:

- 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 can also require a minimal access charge.
- You can use PPPoE on current CPE installations that cannot be upgraded to PPP or that do not have the ability to run PPPoA, that extends 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 to run point-to-point over Ethernet without the requirement of an intermediate IP stack.
- The Network Access Provider (NAP) or Network Service Provider (NSP) can provide secure access to a corporate gateway without the management of end-to-end permanent virtual circuits (PVCs) and without the use of Layer 3 routing and/or Layer 2 Tunneling Protocol (L2TP) tunnels. This makes the business model of the sale of wholesale services and virtual private networks (VPNs) scalable.
- PPPoE can provide a host (PC) access to multiple destinations at a given time. You can have multiple PPPoE sessions per PVC.
- The NSP can oversubscribe by the deployment of idle and session time-outs with the help of an industry standard Remote Authentication Dial-In User Service (RADIUS) server for each subscriber.
- You can use PPP with the service selection gateway (SSG) feature.

Disadvantages

These are some key disadvantages of PPPoE and how they differ from PPPoA:

- You must install PPPoE client software on all hosts (PCs) that connect to the Ethernet segment. This means that the access provider must maintain the CPE and the client software on the PC.
- Since PPPoE implementation uses RFC 1483 bridging, it is susceptible to broadcast storms and possible denial-of-service attacks.


Implementation Considerations for PPPoE Architecture

These are some key points to consider before you implement this type of architecture.

- The number of subscribers that is supported. The number of PPPoE servers required depends on the number of sessions.
- Whether the PPP sessions are terminated at the aggregation router of the service provider or forwarded to other corporate gateways or ISPs.
- Whether the service provider or the final service destination provides the IP address.
- In the case of more than one user, whether all users need to reach the same final destination or service, or do they all have different service destinations. Do the end subscribers require simultaneous access to multiple destinations?
- The PPPoE client software that the access provider uses and whether the software has been tested, the operating system that the host uses, 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 the provision of CPEs, DSLAMs and aggregation points of presence (POPs).
- The business model for the NAP. Does the model also include the sale of 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 offers to the end subscriber.
- The anticipated upstream and downstream volume of data flow. Consider NRP throughput, traffic engineering, and any QoS issues.

This document discusses how the PPPoE architecture fits and scales to different business models for service providers and how the providers can benefit with the help of this architecture.

Network Architecture



Design Considerations for PPPoE Architecture

This section covers issues that apply specifically to PPPoE Architecture.

Before the deployment of any architecture, it is essential to understand the business model of the service provider and what services the provider offers. You need to know the client software that is used on the PC. The most common software is from Routerware. Since the client software is installed on a PC, the service provider technician needs to have a good knowledge of that PC and its operating system.

As specified in RFC 2516, the maximum receive unit (MRU) option must not negotiate 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 is achieved with the configuration of 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.

In order 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 issue the **pppoe session-limit per-mac** and **pppoe session-limit per-vc** commands in order to change the defaults.


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 RFC 1483 bridging. Each CPE consumes only one VPI/VCI pair and all PPPoE sessions initiated by hosts behind this CPE is carried over in this single VC.

IP Management



The IP address allocation for the individual host that runs 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 purchases and where the PPP sessions terminate. 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 can come from the access concentrator that terminates 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

These are the ways to reach the service destination:

- The termination of PPP sessions at the service provider
- L2TP tunneling
- With the use of SSG

Detailed explanations of these architectures are 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 Ethernet address of the peer, which together uniquely define the PPPoE session. These are the steps:

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 that indicates what service type it requests.

2. One or more access concentrators send offer packets.

When the access concentrator or the router receives a PADI that it can serve, it sends 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. Since the PADI was broadcast, the host can 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 is based on the services offered by each access concentrator. The host then sends one PADR packet to the access concentrator it chooses. The destination_addr field is set to the unicast Ethernet address of the access concentrator or the router that sends 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 sends 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 can be sent by either the host or the access concentrator any time after a session is established in order to indicate that a PPPoE session has been terminated.

For a more detailed explanation, refer to RFC 2516.

Conclusion

For ADSL, PPPoE gains popularity, and is second only to PPPoA.

References

- RFC 2516 A method to transmit 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 Technical Support**
- **Technical Support – Cisco Systems**

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2014 – 2015 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Jun 01, 2005

Document ID: 12915
