

Routed Bridged Encapsulation Baseline Architecture

Document ID: 12917

- Introduction**
- Assumption**
- Technology Brief**
- Operational Description**
- RBE Advantages**
- Implementation Considerations**
- Network Architecture**
- Design Considerations for RBE Architecture**
- Key Points of RBE**
 - CPE
 - IP Management
- How a Service Destination is Reached**
 - Providing Internet Access
 - Wholesale Services
 - Corporate Access
 - Service Selection Capabilities
- Conclusion**
- Related Information**

Introduction

This document describes an end-to-end asymmetric digital subscriber line (ADSL) architecture that uses the Routed Bridged Encapsulation (RBE) feature for the Cisco 6400 Universal Access Concentrator (UAC). RBE was developed to address the known RFC1483 bridging issues, including broadcast storms and security. Except for the fact that it operates exclusively over ATM, the RBE feature functions identically to half-bridging. Additional scalability, performance, and security can be achieved by using the unique characteristics of xDSL subscribers.

Assumption

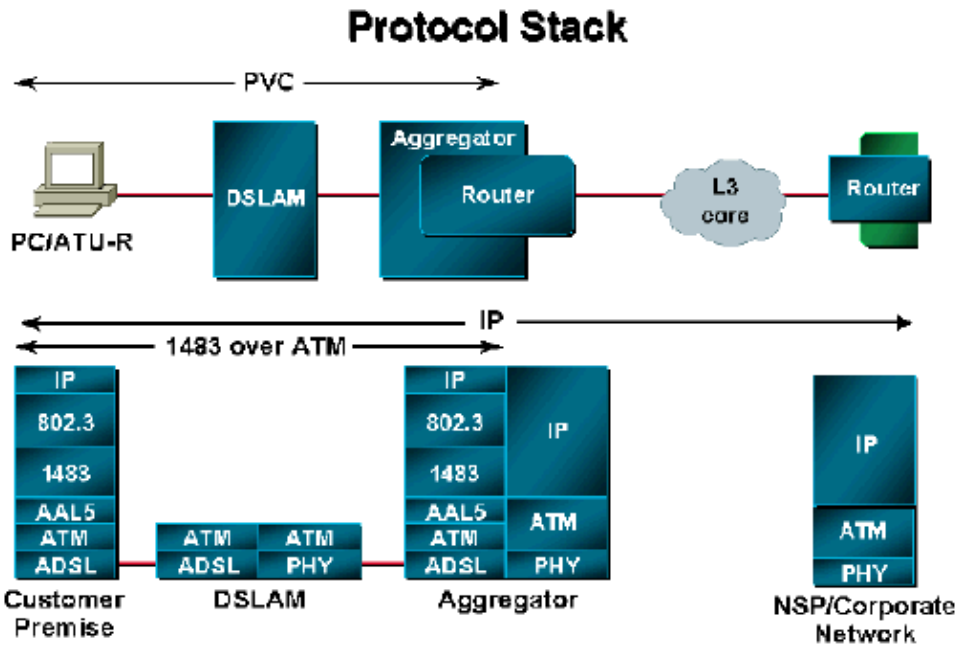
The baseline architecture is designed using the ADSL Forum Reference Architecture Model. The architecture covers different service offerings by the Network Access Provider (NAP) and different scenarios of how the subscriber traffic is forwarded to the Network Service Provider (NSP). In this architecture, RBE is the assumed encapsulation method used by the Cisco 6400. The content of this document is based on existing deployments, as well as some inhouse tests performed on the architecture. For enhanced features and modifications, refer to the release notes for the latest release of Cisco IOS® Software. Currently, RBE is supported on the Cisco 6400, Cisco 7200 and Cisco 7500 platforms. This document is limited to discussions of the Cisco 6400.

Technology Brief

From the network point of view, the ATM connection looks like a routed connection. Data traffic is received as RFC1483 packets, but they are RFC1483 Ethernet or IEEE 802.3 frames. Instead of bridging the Ethernet or IEEE 802.3 frame, as in the case of regular RFC1483 bridging, the router routes on the Layer 3 header. With the exception of some cursory checks, the bridge header is ignored. This is explained in detail in the next

section.

Operational Description



From an operational point of view, the router operates as if the routed-bridge interface were connected to an Ethernet LAN. The operation is described below in two ways: packets originating from the customer premises and packets destined for the customer premises.

For packets originating from the customer premises, the Ethernet header is skipped and the destination IP address is examined. If the destination IP address is in the route cache, the packet is fastswitched to the outbound interface. If the destination IP address is not in the route cache, the packet is queued for process switching. In the process switch mode, the outbound interface through which the packet must be routed is found by looking in the routing table. After the outbound interface is identified, the packet is routed via that interface. This occurs without the requirement for a bridge group or Bridge Group Virtual Interface (BVI).

For packets destined for the customer premises, the destination IP address of the packet is examined first. The destination interface is determined from the IP routing table. Next, the router checks the Address Resolution Protocol (ARP) table associated with that interface for a destination MAC address to place in the Ethernet header. If none is found, the router generates an ARP request for the destination IP address. The ARP request is forwarded to the destination interface only. This is in contrast to bridging, in which the ARP request is sent to all interfaces in the bridge group.

For a scenario using unnumbered interfaces (where you may find two subscribers on the same subnet), the routed-bridge interface uses proxy ARP. For example, 192.168.1.2 (Host A) wants to communicate with 192.168.1.3 (Host B). However, Host A is on the same subnet as Host B.

Host A must learn the Host B MAC address by sending out an ARP broadcast to Host B. When the routed-bridge interface at the aggregation device receives this broadcast, it will send out a proxy ARP response with the MAC address of 192.168.1.1, Host A. It will take that MAC address, place it in its Ethernet header, and send the packet. When the router receives the packet, it discards the header and looks at the destination IP address, then routes it on the correct interface.

RBE Advantages

RBE was developed with the intention of addressing some of the issues faced by the RFC1483 bridging architecture. RBE retains the major advantages of RFC1483 bridging architecture, while eliminating most of its drawbacks.

- Minimal configuration at the customer premises equipment (CPE).

The service provider considers this important because it no longer requires a large number of truck rolls and no longer needs to invest heavily in personnel for the support of higher level protocols. The CPE in bridge mode acts as a very simple device. Minimal troubleshooting is involved at the CPE since everything that comes in from the Ethernet passes straight over to the WAN side.

- Easy to migrate from pure bridging architectures to RBE. There is no change required at the subscriber end.
- Avoids the IP hijacking and ARP spoofing challenges faced in typical pure bridging architectures. RBE also prevents broadcast storms by using point-to-point connections. Security is the major disadvantage in pure bridging architectures.
- Compared to pure bridging architectures, RBE provides superior performance because of the routing implementation at the aggregation device. Also, RBE is more scalable because it does not have bridge group limitations.
- Supports Layer 3 Web selection using the Cisco Service Selection Gateway (SSG).

Implementation Considerations


Some of the key points to consider before implementing this architecture are the same as mentioned in the RFC1483 Bridging Baseline Architecture paper.

RBE is recommended when:

- The scenarios are the same as in existing bridging architectures.
- The NAP wants only to perform minimal management of CPEs. The concept of a simple CPE requires minimal or no configuration after the CPE is deployed at the subscriber's location.
- The NAP does not want to install and maintain host clients on the hosts behind the bridged CPE. These installation and maintenance tasks increase deployment costs and maintenance, including the provision of help desk personnel with knowledge of the client software and the operating system on which the client is running.
- The NAP wants to deploy a scalable and secured bridged network using *existing* CPEs (which can only operate in RFC1483 bridging mode) and wants to offer service selection capabilities.

The next discussion explains how RBE architecture fits and scales to different business models.

Network Architecture



The RBE network architecture is similar to RFC1483 bridging architecture. As specified in that architecture, the aggregation device could be either in the NAP or at the NSP. If an end-to-end permanent virtual circuit (PVC) architecture is used, the NSP terminates the subscribers and configures RBE at the aggregation device. If the NAP prefers to provide wholesale services plus service selection, it may opt to terminate those subscribers and get IP addresses from a local Dynamic Host Configuration Protocol (DHCP) server. In the case of wholesale services, the NAP may opt to get the IP addresses from the NSP. These scenarios are covered in detail in the IP Management section of this document.

Design Considerations for RBE Architecture

RBE eliminates the major security risks involved with RFC1483 bridging architecture. Additionally, RBE provides better performance and is more scalable because the subinterfaces are being treated as routed interfaces.

This section explains some of the key points that must be considered before designing RBE architecture. For the subscriber side, the design principles remain the same as in RFC1483 bridging architecture.

In RBE, a single virtual circuit (VC) is allocated a route, a set of routes, or a classless interdomain routing (CIDR) subnet. Thus, the trusted environment is reduced to only the single customer premises represented by either the IP addresses in the set of routes or the CIDR block. The ISP also controls the addresses assigned to the user. This is done by configuring a subnet on the subinterface to that user. Therefore, if a user misconfigures equipment with an IP address outside the allocated address range (possibly causing ARP packets to flow up to the router), the router generates a "wrong cable" error and refuses to enter the erroneous IP to MAC address mapping into its ARP table.

RBE can be deployed using only point-to-point ATM subinterfaces. It cannot be deployed on multipoint subinterfaces. Even though the subscriber side is bridged, you do not need to define bridge groups or BVI interfaces because the subinterfaces are treated as routed interfaces.

The ATM point-to-point subinterfaces can be numbered interfaces or unnumbered to some other interfaces.

By definition, a numbered interface is an interface which has a specific IP address assigned to it with a fixed subnet mask. For example:

```
Interface atm0/0/0.132 point-to-point
ip address 192.168.1.1 255.255.255.252
```

As shown in this example, when RBE is deployed with a numbered interface, there should be a separate subnet for each subscriber. The host at the subscriber end should be configured for 192.168.1.2. There is only one host at the subscriber end. If the requirement is to support more than one host, the subnet mask chosen should accommodate more hosts.

Numbered interfaces give the NAP control over the number of hosts the subscriber has connected behind the CPE. As explained above, this lack of control was a major problem in RFC1483 bridging architecture.

However, this methodology consumes too many IP addresses. You will need to allocate one subnet per subscriber, use one IP address for the ATM subinterface, and leave the broadcast address and all zero addresses unused. So, to have one host behind the CPE, you at least need to define a subnet mask of 255.255.255.252. Considering the scarcity of IP addresses, this may not be a feasible option unless the NAP/NSP is using private address space and performing Network Address Translation (NAT) to reach the outside world.

In order to conserve IP addresses, an alternative would be to use unnumbered interfaces. By definition, an unnumbered interface is an interface that uses another interface's IP address by using the **ip unnumbered** command. For example:

```
!  
interface loopback 0  
ip address 192.168.1.1 255.255.255.0  
!  
interface atm0/0/0.132 point-to-point  
ip unnumbered loopback 0  
!  
interface atm0/0/0.133 point-to-point  
ip unnumbered loopback 0
```

As shown in the example above, an IP address and subnet are only applied to the loopback interface. All ATM subinterfaces would be unnumbered to that loopback interface. In this scenario, all the subscribers being terminated on ATM subinterfaces (unnumbered to loopback 0) would be on the same subnet as that of loopback 0. This implies that subscribers would be on the same subnet, but would be coming in through different routed interfaces. In this situation, it becomes a problem for the router to identify which subscriber is behind which ATM subinterface. For Cisco IOS, 192.168.1.0 (in the IP Management diagram) is directly connected via interface loopback 0, and it is never going to send traffic destined to any of the host addresses on that subnet via any other interface. In order to resolve this issue, you need to explicitly configure static host routes. For example:

```
ip route 192.168.1.2 255.255.255.255 atm0/0/0.132  
ip route 192.168.1.3 255.255.255.255 atm0/0/0.133
```

As specified in this example, when the router needs to make a routing decision and needs to forward the traffic destined for 192.168.1.2, it will choose ATM 0/0/0.132 as the outgoing interface, and so on. Without specifying those static host routes, the router would choose the outgoing interface as loopback 0 and drop the packet.

Even though the unnumbered interface would conserve IP addresses, it requires an additional task of configuring static host routes on the Node Route Processor (NRP) for each subscriber. Note that if a subscriber has, for example, 14 hosts behind the CPE, it is not required to have static host routes for each host. A summarized route can be defined for the ATM subinterface.

Thus far, this explanation has assumed that the hosts behind the CPE will be configured for static IP addresses. This assumption is not true in real life designs. In the practical world, the NAP wants to perform minimal configuration and maintenance for the CPE and the hosts attached to it. In order to achieve that, the hosts should get their addresses dynamically using a DHCP server.

In order to get their IP addresses dynamically, hosts must be configured to get IP addresses from a DHCP server. When the host boots up, it sends out DHCP requests. These requests are then relayed to the appropriate DHCP server, which assigns an IP address to the host from one in its previously defined scope.

In order to forward the initial DHCP requests from the host to the appropriate DHCP server, you should apply the **ip helper-address** command to the interface which is receiving the broadcasts. After the broadcasts are received, the Cisco IOS looks at the configuration of the ip helper-address for that interface and forwards those requests in a unicast packet to the appropriate DHCP server whose IP address is specified in ip helper-address. After the DHCP server replies with the IP address, it sends the response to the interface on the router that originally forwarded the request. This is used as the outbound interface to send the DHCP server response to the host that originally requested the service. The router also automatically installs a host route for this address.

If RBE is enabled on a subinterface and is an IEEE 802.3 bridged protocol data unit (PDU), the Ethernet encapsulation is examined after ATM bridge encapsulation. If it is an IP/ARP packet, it is handled like any other IP/ARP packet. The IP packet is fastswitched. If it fails, it is queued for process switching.

Performance for RBE is a big win. Today's standard bridging code has the inherent problem of requiring two separate classifications for a packet before a forwarding decision can be made. A classification is defined as the process of examining (on the upstream) and modifying (on the downstream) the packet header for forwarding information, which is relatively expensive. A Layer 2 lookup is needed to determine whether the packet needs to be routed or bridged. Then, at Layer 3, a lookup is needed to identify the location to which the packet should be routed. This classification is done in the upstream as well as downstream directions, which has an impact on performance.

For RBE, it is predetermined by configuration that the packet is to be routed in the upstream direction. Therefore, it is not necessary to go through the bridging forwarding path, which was necessary in the case of standard bridging.

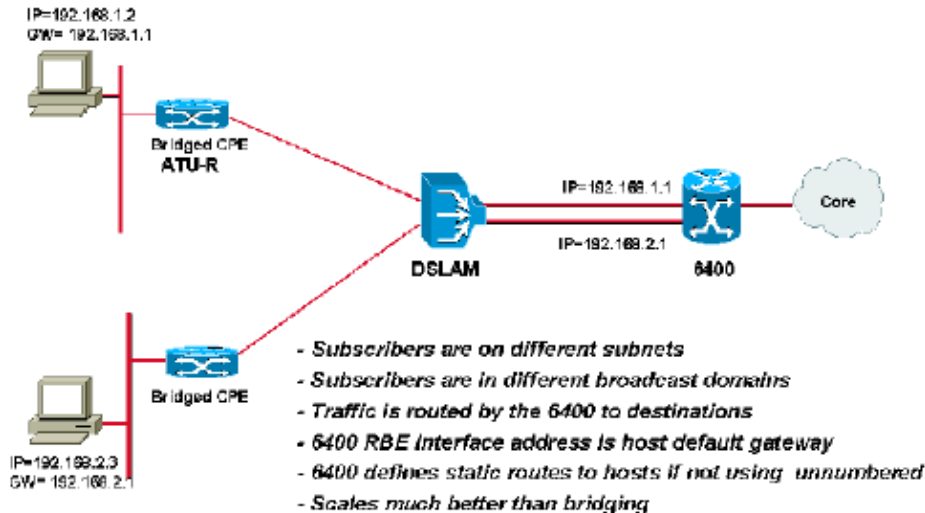
Key Points of RBE

CPE

The CPE configuration remains the same as in standard bridging. No changes to the CPE are required to deploy RBE.

IP Management

Numbered Interfaces



While deploying the numbered interfaces for RBE, the IP address allocation to the host behind the bridged CPE is usually handled via a DHCP server. As mentioned earlier, the DHCP server can reside in the NAP or in the NSP. For either case, the numbered ATM subinterface should be configured with the **ip helper-address** command. If the DHCP server is going to be located at the NSP, the NAP aggregation device must have a route to reach that server. The only scenario in which a NAP would use its own DHCP server and IP address range is when it wants to offer service selection capabilities to the subscribers, and those subscribers are LAN attached to the NAP.

If the NAP wants to use the IP address space of the NSP, one of the IP addresses for each subnet should be allocated to the ATM subinterface. Also, there should be some mutual agreement between the NAP and NSP so that the NAP configures the correct address. When the NSP's DHCP server assigns IP addresses, this agreement should be in place to ensure that the server provides the correct default gateway information to the host. The NAP can then either summarize a static route for all those addresses assigned to subscribers, or it can choose to run a routing protocol with the NSP to advertise those routes. In most scenarios, both the NAP and NSP would prefer not to use a routing protocol. Providing a static route is a good option.

This is the basic configuration required on the NRP for deploying RBE with numbered interfaces:

```
!  
interface ATM0/0/0.132 point-to-point  
ip address 192.168.1.1 255.255.255.0  
ip helper-address 192.168.3.1  
no ip directed-broadcast  
atm route-bridged ip  
pvc 1/32  
encapsulation aal5snap  
!  
interface ATM0/0/0.133 point-to-point  
ip address 192.168.2.1 255.255.255.0  
ip helper-address 192.168.3.1  
no ip directed-broadcast  
atm route-bridged ip  
pvc 1/33  
encapsulation aal5snap
```

Using unnumbered interfaces is the best way to conserve IP addresses. As explained earlier, when unnumbered interfaces are used with DHCP, host routes are dynamically installed. This may be the best approach to deploying RBE. The DHCP server can then be located at either the NAP or the NSP, as for numbered interfaces.

This is the basic configuration required on the NRP for deploying RBE with unnumbered interfaces:

```
interface Loopback0
ip address 192.168.1.1 255.255.255.0
no ip directed-broadcast
!
interface ATM0/0/0.132 point-to-point
ip unnumbered Loopback0
no ip directed-broadcast
ATM route-bridged ip
pvc 1/32
encapsulation aal5snap
!
interface ATM0/0/0.133 point-to-point
ip unnumbered Loopback0
no ip directed-broadcast
ATM route-bridged ip
pvc 1/33
encapsulation aal5snap
```

How a Service Destination is Reached

So far, this document has discussed the basic access technology using RBE as the encapsulation method. However, using this architecture, the NAP/NSP can also offer various services and different options for where the NAP can forward the subscriber traffic to the NSP. These concepts are explained in the next sections.

Providing Internet Access

In this scenario, the primary function for the NSP is to provide high speed Internet access to the end subscribers. Because the NSP is going to provide the final service, IP address management becomes the responsibility of the NSP. It can assign public IP addresses to its end subscribers using a DHCP server, or it can opt to provide private IP addresses to the subscribers and then perform NAT to reach the outside world.

Wholesale Services

If the NAP wants to offer wholesale services to other ISPs, it can do so. In this scenario, the NAP usually does not prefer to handle IP addresses for all the subscribers for different NSPs. The NAP makes some arrangement with the ISP to provide IP addresses to those subscribers. This can be achieved by the NAP forwarding the DHCP requests coming from the subscribers to the DHCP servers at the NSPs. The NAP has to configure its ATM subinterfaces with one of the IP addresses from that range, and it needs to advertise those routes to the NSP. The route advertisement could be in the form of either a static route or some routing protocol between the NAP and NSP. Static route is the preferable method for the NAP, as well as the NSP.

Corporate Access

Corporate access usually requires Virtual Private Network (VPN) services. This means that the corporation will not provide any IP addresses to the NAP and does not allow the NAP to advertise the corporate IP address space in the NAP's IP core, as it could result in a security breach. Corporations usually prefer to apply their own IP addresses to their clients, or they will allow access via some secured means such as Multiprotocol Label Switching/Virtual Private Network (MPLS/VPN) or Layer 2 Tunneling Protocol (L2TP).

The other approach to providing secured corporate access is where the NAP provides the initial IP addresses to those subscribers. Therefore, the subscribers becomes LAN-attached to the NAP. After the subscribers have initial IP addresses, they can initiate a tunnel to the corporation through L2TP client software running on the host. In turn, the corporation will authenticate this subscriber and provide an IP address from its IP address space. This IP address is used by the L2TP VPN adapter. This way, the subscribers have the option to either

connect to their ISP for Internet connection or gain access to their corporation through a secured L2TP tunnel access. However, this requires the corporation to provide the tunnel destination IP address to the subscriber, which should be routable through the NAP's IP core.

Service Selection Capabilities

The NAP could offer various service selection capabilities using Cisco SSG's functionality. The SSG offers two methods for providing service selection: via Layer 2 (which is known as PTA–MD) and Layer 3 Web selection. With RBE, only the Layer 3 Web selection method can be used. This requires the subscribers to be LAN–attached to the NAP; that is, the NAP provides the initial IP address to the subscriber and provides access to the Cisco Service Selection Dashboard (SSD).

In the case of RBE architecture, the Cisco SSG's Web selection method is a good way to account for subscriber traffic.

Conclusion

RBE provides better performance and is more scalable than standard bridging. It also overcomes all security issues faced in standard bridging. RBE eliminates the broadcast storm problems of standard bridging. RBE provides a robust architecture for the NAP that wants to avoid the maintenance of client host software, bridging–related issues, and wants lower deployment costs. With RBE, all this is possible while using the existing bridging architecture.

Related Information

- [Cisco ADSL Product Support Information](#)
- [Technical Support & Documentation – Cisco Systems](#)

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

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Jun 07, 2005

Document ID: 12917
