

# Broadband Access Architectures

**Point-to-Point Protocol (PPP) Comes of Age**

## White Paper

---

Internet access is not a luxury anymore. For millions of users, it is a necessity. The Internet facilitates many day-to-day activities in people's personal and professional lives. It allows people to instantly communicate with each other, "telecommute" to the office, shop, study, and provides entertainment. Now that Internet access is universally available, the next major requirement is bandwidth – and bigger is better! Slow dial-up speeds are no longer sufficient for downloading MP3 soundtracks, video clips, large corporate spreadsheets, and presentations or even surfing most graphics-enabled Web sites.

Many Internet users (not just the proverbial power users) are willing to pay a premium for high-speed broadband access. Network providers (especially the ILECs) are rapidly deploying broadband services to gain market share and capture some of this incremental revenue. The broadband market has continued to grow, even during the recession – albeit at a declining rate. Currently there are more than ten million broadband Internet subscribers in the United States, and that quantity is increasing daily; further incremental growth can be expected as the global economy improves. Accordingly, many service providers and network equipment manufacturers (NEMs) are building hardware and infrastructure to support this high-growth market.

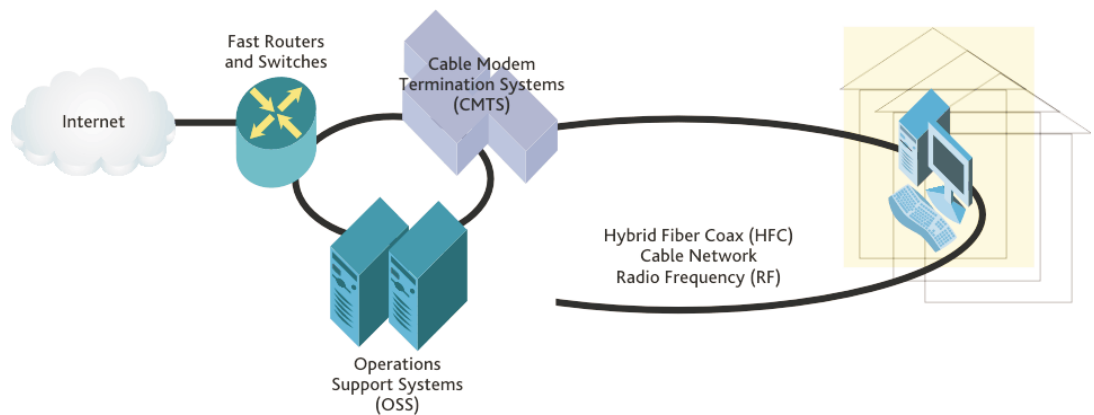
The two most common manifestations of broadband Internet access are ADSL and cable modems. Both provide high-speed connectivity and are readily available in most major metropolitan areas. Both also have similar challenges with regard to large-scale deployments. These challenges include ease of deployment (after all, sending installation technicians to all of the subscribers' homes is not a scalable solution); authentication; network security; and support for multiple PCs, user sessions, and all types of client hardware and software. Several network architectural alternatives are available to address these issues. Since a network's fundamental hardware infrastructure is typically dictated by external factors (cable plant, distance, fiber availability, quantity of subscribers, funding, etc.), the architectural flexibility really refers to the network protocols that reside on top of the physical infrastructure. These protocol options include: DHCP, L2TP, PPPoA, PPPoE, and static IP addresses. As discussed below, each of these alternatives have positive and negative implications.

## Infrastructure

The underlying infrastructure for ADSL or cable modem broadband edge networks is similar for both access methods. Many basic hardware components can be used in either type of implementation. The essence of the designs consists of inexpensively grouping together large quantities of users at the edge of the network and then forwarding aggregated traffic to the core of the Internet via ultra-high speed communication links. Basic user services such as authentication, address assignment, and security are pushed to the edge of the network, but these will be discussed in the next sections.

### Cable Modems

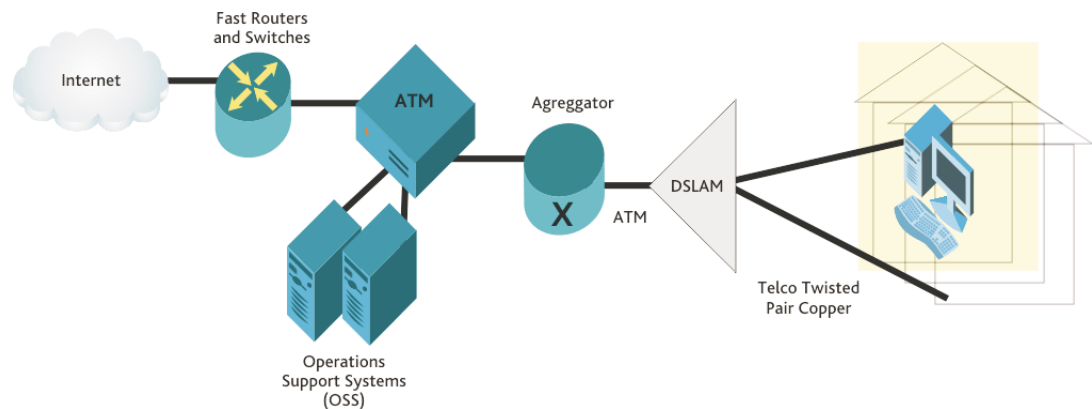
More than 100 million homes and businesses in North America have access to cable television. This means that all of these offices and residences have, at a minimum, some sort of a one-way coaxial cable termination. Over 60% of the North American cable-TV plant has been upgraded to support bi-directional communications, which is a fundamental requirement for broadband Internet access.



The figure above shows a typical cable modem infrastructure. Starting from the right, homes and business have connections to an RF-based broadband coaxial cable network. The physical connection to the network is via a cable modem that is interconnected to a workstation over an Ethernet cable at the user's home or business. At the remote end, these cables are terminated by an active electronic head-end called a Cable Modem Termination System (CMTS). The headends are then interconnected via a series of high-performance LAN and WAN routers and switches. These routers forward IP packets to and from the Internet. The current access bandwidth, according to the prevailing Data Over Cable System Interface Specifications (DOCSIS), can vary from 27 to 36 Mbps over a single 6 MHz cable channel. However, this bandwidth is shared among all of the households and businesses on any given coax segment. Typically, cable providers will try to deploy between 50 and 100 users per segment, resulting in unpredictable bandwidth availability for individual subscribers.

### ADSL

The traditional telephone companies (ILECs) also have an efficient means for delivering broadband Internet access. This is called Digital Subscriber Line (DSL) service. There are several different types of DSL, the most common of which is Asynchronous DSL, or ADSL. ADSL also has many different appearances. Some of these are based upon equipment availability, cable distances, or even economic factors. In short, ADSL can support data rates up to 7 Mbps, but telcos typically are offering "fractional T1" services at downstream rates varying from 384 Kbps to 1 Mbps. Like cable modems, DSL service is asynchronous with residential upstream return paths supporting speeds of 224 Kbps or less. Unlike cable modems, the ADSL cables are not shared with other users, so the subscriber will be able to use the full bandwidth without contention.



The basic ADSL reference model is illustrated above. Following the diagram from the right (homes and business) to the left (the big Internet), ADSL starts with a user's workstation attached to an ADSL modem called an ADSL Termination Unit – Remote (ATU-R). This modem is then connected to the end-user's existing twisted pair copper telephone cables. An inexpensive line splitter can be installed in the wall-jack at the user's location to permit concurrent voice and data support over the same twisted pair circuit (though some paradigms also provide for direct support of DSL-based telephone services). The twisted pairs from multiple locations are then grouped together at a telco Point or Presence (POP) with a device known as a DSL Access Multiplexer (DSLAM) which combines all of the individual user sessions onto an ATM trunk. Next, the traffic from multiple DSLAMs is further combined at an aggregation point, typically located in a central office. From there the traffic is forwarded to a router, which in turn connects to the Internet.

## Challenges

Cable modems and ADSL both share the single greatest challenge – ease of deployment. That means service providers must make it very simple for the general public to adopt their service. An easy-to-use service will be readily installed by users and will result in greater customer satisfaction. This, in turn, will reduce costly service calls and will ultimately yield increased subscriber quantities.

Another major deployment challenge is address assignment – all of the subscribers require IP addresses to communicate over the Internet, but all individual addresses must be unique. These addresses must be configured on the users' PCs – a simple task for network engineers, but a potentially daunting challenge for non-technical people. Other challenges include network provisioning, security, and support for multiple devices. None of these issues can be addressed solely by the hardware infrastructures described previously. Instead, some additional "soft" layers must be added on top of the physical network.

## Alternatives

These challenges can be addressed by implementing an additional protocol between the end-user and the service provider. This protocol will not be carried to the Internet, instead it will facilitate some of the overhead functions necessary for user configuration and control. There are five basic options for configuration architectures, and these are: static IP addresses, Dynamic Host Configuration Protocol (DHCP), Layer Two Tunneling Protocol (L2TP), Point-to-Point Protocol over ATM (PPPoA), and the Point-to-Point Protocol over Ethernet (PPPoE). Each of these will have its own unique benefits and shortcomings.

## Static IP Addresses

Static IP addresses are not really a protocol but they do represent a valid architectural option. This is the most rudimentary configuration option. This simply consists of having the service provider assign an IP address to the user, who then manually configures the appropriate network parameters for his workstation. As the word "static" implies, this solution offers no flexibility and does not support multiple workstations or roving users. Static IP addresses are cumbersome for the service provider to administer, and they are equally difficult for the users to maintain.

## DHCP

The Dynamic Host Configuration Protocol (DHCP) was specifically designed to automate the IP configuration of PCs. This is very commonly implemented in office environments where users may come and go with their laptops. DHCP offers some significant advantages over static IP addresses. As the name indicates, DHCP is dynamic. This means that when a PC on a broadband network is activated, it will automatically be assigned its IP parameters by a central server. This architecture is quite flexible and can facilitate multiple workstations, as well as nomadic users. DHCP support is very easy to configure from the workstation (end-user) side and it is also equally straightforward for the service provider to administer a centrally located DHCP network server.

## L2TP

Another more recent alternative for broadband network access is the Layer Two Tunneling Protocol (L2TP). This is emerging as a key technology for the creation of Virtual Private Networks (VPNs). This architecture facilitates a "tunnel" through the Internet to a predetermined destination. This means the basic configuration and control parameters associated with any individual session can be assigned from a remote location across the Internet. Like DHCP, this protocol supports automatic user configuration and also supports transient users. However, the tradeoff for these benefits is a substantial increase in network complexity and overhead. L2TP also may have some serious scalability implications – maintaining thousands of active tunnels will be quite challenging for the service provider.

## PPP

The final alternatives are all based upon the well-established Point-to-Point Protocol (PPP). There are two different forms of this protocol commonly implemented in broadband access networks: PPP over Ethernet (PPPoE) and PPP over ATM (PPPoA). Additionally, there also is a combination of both of these protocols called PPP over Ethernet over ATM or PPPoEoA. PPP offers all of the flexibility and ease of use associated with the other architectures and adds a level of security as well. Hence, it is the protocol of choice for most broadband service providers.

PPP is ancient according to Internet technology – it was initially adopted as a standard in 1989 (RFCs 1171 and 1172). It has been updated several times since then and its most recent version is contained in RFC 1661 from July 1994. PPP was designed to support communications between devices (typically routers) over traditional telco leased-line circuits (i.e., T1s). Since line quality was a major concern in the late 1980s and early 1990s, there are several mechanisms intrinsic to PPP to help it tolerate circuit degradation. In order to facilitate this adaptation, PPP was designed to be a connection-oriented protocol. This would appear to be unnecessary since it was initially built for use on point-to-point circuits which can only have two unique end-points. In addition to link quality checks (the Link Control Protocol (LCP) is a subset of PPP which is used for this purpose), the fact that PPP is connection-oriented also enhances security by requiring that packets are forwarded only to specific, known addresses.

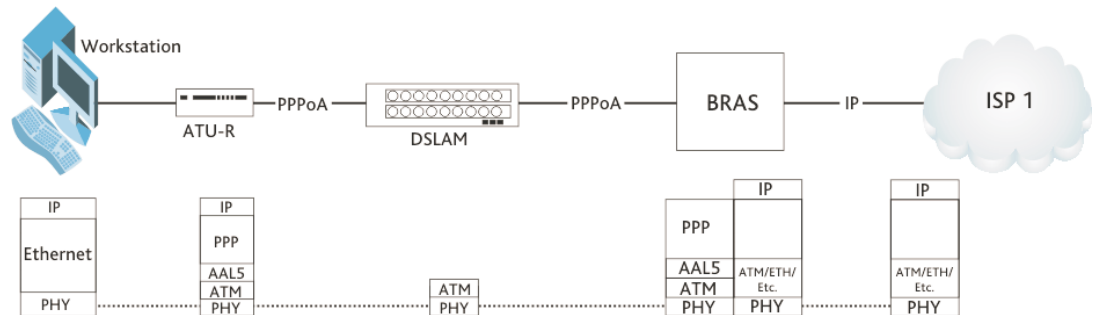
In the early 1990s, PPP was adapted to support dial-up links. Again, being connection-oriented, it helped secure the links. Furthermore, since a "handshake" procedure is required as the link is initiated, it was easy to insert an authentication algorithm into this process. In a dial-up environment, a device known as a Remote Access Server (RAS) will validate the users ID and password. PPP will then provide for the automatic configuration of a link, including supplying the remote device with an appropriate IP address. All of these capabilities have combined to make PPP the favorite protocol for dial-up connectivity – PPP is supported by virtually all ISPs. In fact, dial-up PPP is included in most commonly used PC operating systems.

**PPPoA**

All of the advantages of PPP in a dial-up environment are also extensible for higher speed broadband connectivity. However, the commonly used broadband topologies (ADSL and cable modems) are not based upon point-to-point circuits. Instead, they utilize some degree of network and/or component sharing to gain economies of scale. Hence PPP needed to adapt yet again to accommodate the evolving network architecture. Since ADSL is based largely on an ATM infrastructure, PPP over ATM (PPPoA) is used for this application. PPPoA is codified by RFC 2364, adopted in July 1998. This topology also lets service providers take advantage of ATM’s sophisticated quality of service facilities to provide service level agreements (SLAs) to their customers.

In a dial-up environment, PPP sessions are terminated on a Remote Access Server (RAS). Broadband networks use a similar device known as a Broadband Remote Access Server (BRAS) to terminate large quantities of PPPoA connections. This device is also commonly called an aggregation router or aggregator and is located in a central office facility of a service provider. The user end of the PPPoA connection can be terminated at the user’s workstation, but this will typically require the installation of additional complex software and an expensive ATM network interface card (NIC) in the computer. This complexity and expense are usually conveniently hidden from the user by installing the protocol stack in the ATU-R modem. Hence the PPPoA protocol usually operates between the DSL modem and the aggregator (a.k.a. BRAS).

The diagram below traces the PPPoA connections and protocol stack through the various pieces of network hardware. In essence, the protocol originates at the ATU-R at the user’s location and terminates in the BRAS at the service provider’s facility. An ATM permanent virtual circuit (PVC) carries the protocol (and the user data) between these two devices. The BRAS then aggregates all of the PVCs into an IP flow, which then is forwarded to upstream routers and the Internet.



**PPPoE**

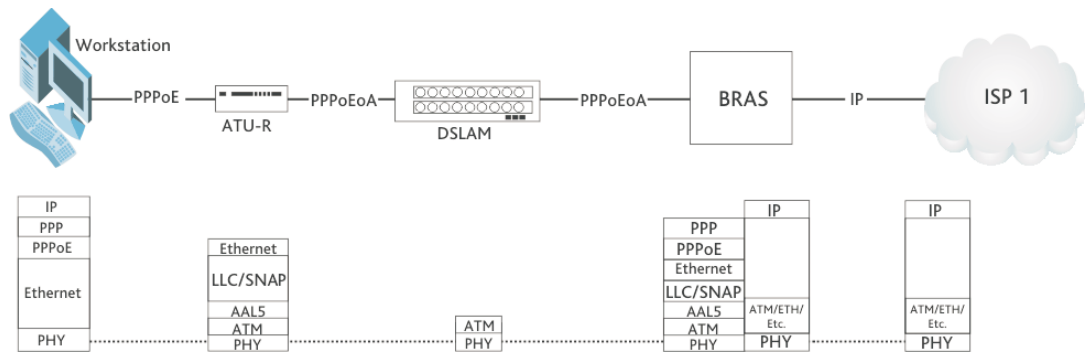
The diagram above indicates a typical ADSL broadband solution. Cable modems are supported by an even simpler architecture. Since the PC is connected to a cable modem via Ethernet, there is no need to translate multiple protocols. Instead, Ethernet can be extended across the cable network all the way to the cable modem equivalent of a central office. Therefore, it is logical to extend PPP the entire length of the connection thus increasing the range of the benefits associated with PPP and also eliminating the requirement for any protocol translation. Hence, it is feasible to run PPP on the workstation and continue that session all the way to an aggregation router at the ISPs facility. This requires another type of PPP known as PPP over Ethernet, or PPPoE. PPPoE is defined by RFC 2516 from February 1999.

The only drawback to a PPPoE solution is that it requires the addition of PPPoE software on the workstation. However, this shortcoming can also be turned around and used as an advantage for service providers. Service providers can distribute CDs with the appropriate PPPoE protocol stack (commonly called a “shim”) to their customers. Additionally, they can include other programs such as a GUI front-end for their service, advertising, and other management and diagnostic tools. Thus the PPPoE client software can be installed transparently, along with whatever accompanying utilities the service provider desires.

**PPPoEoA**

Returning to the ADSL model, it should be noted that an ATM PVC extends from the ATU-R to the BRAS. However, the user’s workstation is connected to the ATU-R via the ubiquitous Ethernet protocol. The ATU-R functions as a bridge between the two protocols. Since the PPP protocol stops at the ATU-R, the user’s Ethernet segment of the connection is not subject to the benefits of the PPP Link Control Protocol. Therefore, when the user’s PC is turned off, the service provider may still believe that a valid connection exists since the LCP keep-alive messages will continue to be answered from the powered-on ATU-R. This effects network security and also usurps network resources that could otherwise be used by active customers. This problem can be resolved within a DSL environment by extending PPP to the workstation in a manner comparable to the cable modem architecture. Initially this could only be accomplished by installing an expensive ATM NIC in the PC. However, a more progressive alternative is to invoke the PPP over Ethernet over ATM (PPPoEoA) protocol.

PPPoEoA is not an approved standard within any of the usual standards bodies. However, it is a de facto standard that is growing in popularity throughout the broadband service provider community. Simply put, PPPoE is enabled on the workstation. This is then encapsulated in ATM at the ATU-R for transit over the service provider’s network. The aggregation router then strips off the overhead bytes and is left with a valid IP datagram that can be forwarded to the Internet.



Diagrammatically, PPPoEoA appears to be very similar to PPPoA. However, the PPP protocol is extended to the workstation, and then encapsulated in both Ethernet and ATM. Though this adds a minor degree of complexity to the connection, it counterbalances that issue by applying all of the benefits of PPP over the entire length of the connection.

## Summary of Alternatives

Several different architectures are available to support broadband Internet access. All of these alternatives are accompanied by their own unique benefits and drawbacks. These are summed up in the chart below. Broadband service providers have evaluated all of these alternatives and selected the choice that best fits their individual business models. Overwhelmingly, the service providers are selecting one of the manifestations of PPP. The network equipment manufacturers who are building very large, robust and scalable aggregation routers – boxes that can manage thousands of concurrent PPP sessions – have further ratified this decision.

	Static IP	DHCP	L2TP	PPPoA	PPPoE	PPPoEoA
Automatic IP Address Configuration	No	Yes	Yes	Yes	Yes	Yes
Cable Modem Support	Yes	Yes	Yes	No	Yes	No
DSL Support	Yes	Yes	Yes	Yes	Yes	Yes
Multiple User Sessions	No	No	Yes	Yes	Yes	Yes
Multiple Workstations	Yes	Yes	Yes	No	Yes	Yes
Network Selection	No	No	Yes	Yes	Yes	Yes
Quality of Service Support	No	No	No	Yes	No	Yes
Requires Additional Client Software	No	No	Yes	No	Yes	Yes
User Authentication	No	No	Yes	Yes	Yes	Yes

## PPP Testing Requirements

It appears inevitable PPP will be the protocol used for widespread broadband network deployment, so PPP testing will be a necessity for both service providers and network equipment manufacturers. Both types of organizations will need to test PPP conformance and performance. From a conformance perspective, this testing involves verifying adherence to the appropriate standards (though this will not be possible for PPPoEoA since no RFC exists). A major objective of conformance testing is to ensure multi-vendor interoperability. Another facet of conformance testing is the ability to perform “negative” testing – validating that the equipment operates correctly in response to induced errors. Performance testing involves stress testing the capacity of the equipment by building thousands of concurrent PPP connections. Quantitative benchmark assessments (i.e. call setup rates and packet throughput) are also significant parts of performance testing. And finally, both types of tests can be used together to verify that user data successfully traverses large-scale PPP networks.

Manufacturers of aggregation routers have one additional unique requirement. Since their business models typically require supporting all three types of PPP, they prefer test equipment that provides the flexibility to emulate PPPoA, PPPoE and PPPoEoA.

## Spirent’s Current Capabilities

Spirent Communications currently provides extensive capabilities for testing all forms of PPP. Together, the SmartBits product family and the AX/4000 product line support PPPoA, PPPoE, and PPPoEoA. With both product lines, it is possible to design and execute complex conformance and performance tests. In fact, together these are the industry’s premier products for testing PPP functionality.

Spirent Communications offers users the ability to test up to 32,000 concurrent PPP sessions per port. Link Control Protocol (LCP) “Hellos,” or “keepalives,” are fully supported. Data generation and verification can be performed over these PPP sessions. Session authentication (PAP and CHAP) and negative testing (inducing errors) can also be implemented. In sum, the Spirent product family is extremely valuable to PPP product developers and network architects.

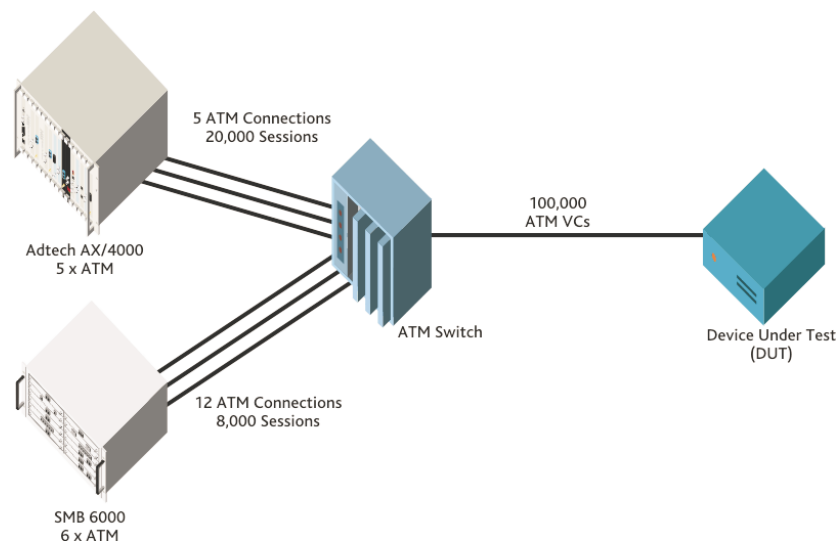
## Massive Scalability

Service providers are all anxious to sell broadband Internet services. However, like all service provider business models, broadband services will only be financially successful if they can be deployed and accepted on a very large scale. Thousands of customers must be serviced at a single POP and expensive networking equipment must be amortized over tens of thousands (maybe even millions) of subscribers.

Network equipment manufacturers are consistently developing larger and more scalable equipment. Manufacturers are "leapfrogging" each other in a game of numbers – each claiming to support more subscribers on a single port or greater density within a single chassis. In fact, some vendors are now claiming to support more than 100,000 simultaneous PPP sessions on a single ATM port.

These new scalability requirements create significant challenges from a testing perspective. As a result of these requirements, an increase in the capacity of test equipment by a full order of magnitude is necessary. Fortunately, the distributed architecture of Spirent Communications' sophisticated test systems was specifically engineered for this purpose.

The diagram below shows a hypothetical test for a typical network equipment manufacturer. In this example, the manufacturer wants to test approximately 100,000 concurrent PPPoA sessions on a single physical port. He wants to run data and LCP Hellos (a.k.a. keep-alives) simultaneously on all of the sessions. Spirent can offer two options for this scenario – an economical solution based upon the SMB-6000 product line, or a solution with advanced features (including negative testing) based upon the AX/4000 product line. Both options use multiple ATM ports on the test equipment. These ports are connected to a local ATM switch in the test lab, which then multiplexes the traffic together so that an aggregate stream of up to 100,000 PVCs is delivered to the device under test. In this manner, infinite scalability becomes possible.



## Summary

The broadband Internet access market is highly lucrative and will continue to be profitable for many years to come. Service providers from all parts of the world are swarming towards this money-making technology. However, there are no templates for building a broadband network. One cannot order a broadband network out of a catalog. Instead, there are many different design and product options that must be considered. These options all have technological, operational, financial, and political implications.

The myriad of choices associated with broadband deployment makes it necessary to test all proposed configurations – interoperability should always be verified. It also is equally essential for network equipment manufacturers to fully test their hardware as part of the development cycle and production process. Spirent Communications provides the equipment to enable these tests.

Though there are many practical choices to be made when building a broadband network, one individual protocol has substantial momentum with the current service providers and equipment manufacturers. That protocol is PPP, and it is gaining popularity due to its many features and its ease of use from both the service provider's and the user's perspectives. Because PPP has its own complexities (three separate protocols for broadband support) and scalability requirements, it necessitates both conformance and performance testing. Again, Spirent Communications has the unique set of tools that will facilitate this process.

## For More Information

- Black, Uyles, PPP and L2TP: Remote Access Communications, Prentice Hall, 1999
- Carlson, James, PPP Design, Implementation and Debugging, Addison Wesley Professional, 2001

## Internet Requests for Comments (RFCs)

- RFC 1661 – The Point-to-Point Protocol (PPP), July 1994
- RFC 2364 – PPP Over AAL5, July 1998
- RFC 2516 – A Method for Transmitting PPP Over Ethernet (PPPoE), February 1999

### Adtech Division

3465 Waialae Avenue  
Suite 200  
Honolulu, Hawaii  
96816 USA  
Tel: 808-734-3300  
Fax: 808-734-7100  
Sales: 800-927-2660  
E-mail: productinfo  
@spirentcom.com

### Sales and Information:

North America  
800-927-2660

Great Britain and Ireland  
+44 (0) 1293 767979

Continental Europe,  
Middle East and Africa  
+33 (0) 1 39 44 09 36

Asia  
spirentasia@spirentcom.com

[www.spirentcom.com](http://www.spirentcom.com)