# Technical Report
# DSL Forum
# TR-092

## Broadband Remote Access Server (BRAS) Requirements Document

## August 2004

**Produced by:**
**The Architecture and Transport Working Group**

**Editor:**
Ed Shrum, BellSouth Communications

**Working Group Co-Chair: David Allan, Nortel Networks**
**Working Group Co-Chair: David Thorne, BT**

**Notice:**

The DSL Forum is a non-profit corporation organized to create guidelines for DSL network system development and deployment. This Technical Report has been approved by members of the Forum. This document is not binding on the DSL Forum, any of its members, or any developer or service provider involved in DSL. This document is subject to change, but only with approval of members of the Forum.

| Revision History | Date | Reason for Update |
|---|---|---|
| Version 1 | July, 2003 | Created new document based on dsl2003.188 |
| Version 2 | September, 2003 | Changes agreed to at the Boston meeting |
| Version 3 | February, 2004 | Inclusion of comments from the Paris meeting and conference calls that took place on 12/16/03, 1/6/04, 1/27/04, and 2/11/04 |
| Version 4 | March, 2004 | Changes agreed to at the Brussels meeting |
| Version 5 | June 1, 2004 | Changes agreed to at the Toronto meeting |
| Version 6 | June 23, 2004 | Changes agreed to during the June 23, 2004 conf call. |

**Table of Contents**

Table of Figures

# 1 Introduction and Purpose

The DSL Forum TR-59 'DSL Evolution – Architecture Requirements for the support for QoS-Enabled IP Services' presents an architecture for evolving DSL deployment and interconnection. It outlines a common methodology for delivering QoS-enabled applications to DSL subscribers from one or more Service Providers. The Broadband Remote Access Server (BRAS) is fundamental to supporting the concepts outlined in TR-59.

Figure1-1 depicts the concept of many-to-many access as a fundamental paradigm shift enabled by the capabilities of a BRAS. While these capabilities are possible in a pure ATM environment, a BRAS device provides greater flexibility and scalability.

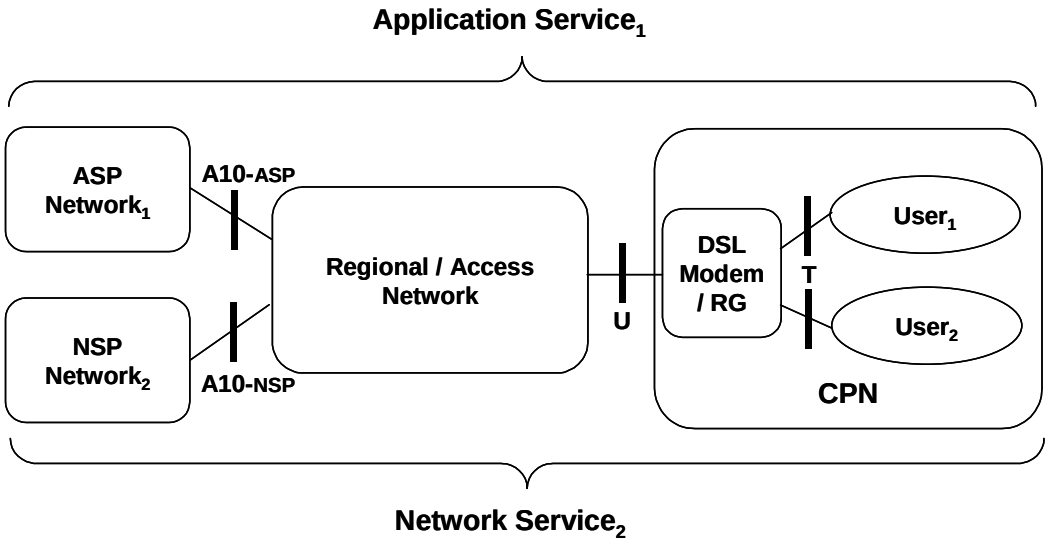


**Figure1-1 - Many-to-Many Access**

The BRAS can perform several logical functions (e.g. LAC, IP router, or a MPLS PE router) as it aggregates user sessions from the access network. The requirements included in this document should be applied broadly across all of these logical functions unless where explicitly stated. In addition to providing basic aggregation capabilities, the BRAS is also the injection point for providing policy management and IP QoS in the Regional and Access Networks. Figure 1-2 depicts the logical representation of where the BRAS is located in the Regional/Access Network. The BRAS is the last IP aware device between service providers (ASPs and NSPs) and the customer network, and as such is leveraged to manage the IP traffic through the layer 2 Access Network. To accomplish this, the BRAS will need to provide a congestion management function that will allow the synthesis of IP QoS through downstream elements that are not QoS aware, which enables DSL providers to support enhanced IP applications.

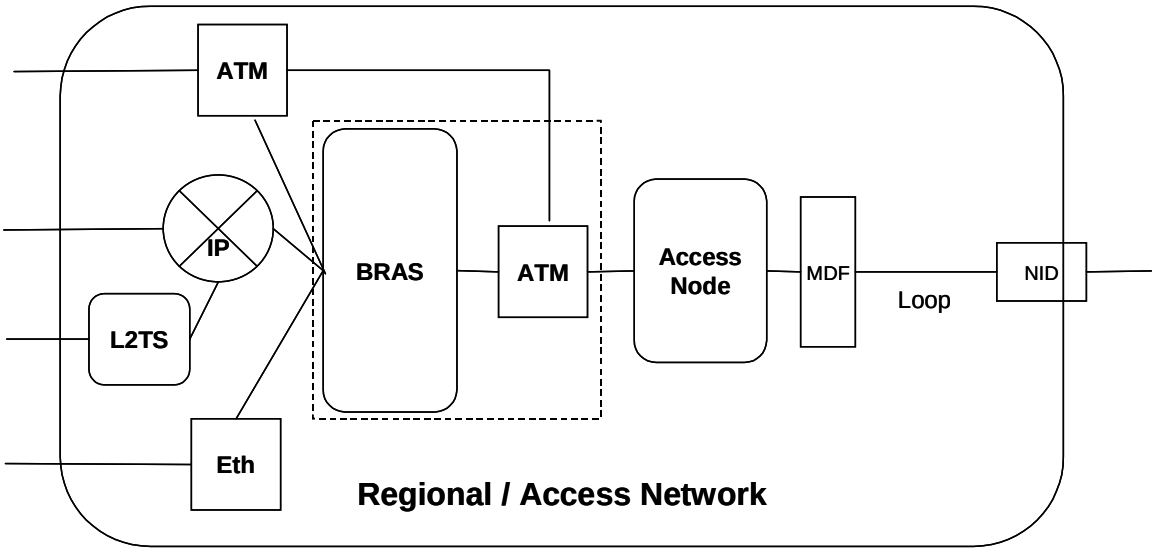


**Figure 1-2 - TR-59 Based Regional/Access Network**

## 1.1 Scope

The requirements included in this document are intended to meet the TR-59 phase 1 and 2 requirements. While the requirements in this document are not meant to be inclusive of all possible deployment scenarios, it is intended to include a broader scope than TR-59. Requirements for services not fully described in TR-59 are included as appendixes. Additionally, service providers may require different scale options for a BRAS device depending on their deployment architecture. The requirements contained within this document pertain to equipment appropriate for the more common large central office deployments. A large CO device will typically support between 64K to 128k subscriber sessions and an aggregate downstream bandwidth of 2.5 Gbps. While focused on a large CO device, this document does not preclude the development of smaller scale devices. The feature requirements contained in this document (with the exception of scaling attributes) apply to any size device.

## 1.2 Requirements

In this document, several words are used to signify the requirements of the specification. These words are often capitalized.

**MUST** This word, or the adjective “REQUIRED”, means that the definition is an absolute requirement of the specification

**MUST NOT** This phrase means that the definition is an absolute prohibition of the specification.

**SHOULD** This word, or the adjective “RECOMMENDED”, means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications must be understood and carefully weighted before choosing a different course.

**MAY** This word, or the adjective “OPTIONAL”, means that this item is one of an allowed set of alternatives. An implementation that does not include this option MUST be prepared to inter-operate with another implementation that does include the option.

## 1.3 Key Terminology

The following definitions apply for the purposes of this document:

| | |
|---|---|
| **Access Network** | The Access Network encompasses the elements of the DSL network from the NID at the customer premises to the BRAS. This network typically includes one or more types of Access Node and often an ATM switching function to aggregate them. |
| **Access Node** | The Access Node contains the ATU-C, which terminates the DSL signal, and physically can be a DSLAM, Next Generation DLC (NG-DLC), or a Remote Access Multiplexer (RAM). A DSLAM hub can be used in a central office to aggregate traffic from multiple remote physical devices, and is considered logically to be a part of the Access Node. When the term "DSLAM" is used in this document, it is intended to very specifically refer to a DSLAM, and not the more generic Access Node. The Access Node provides aggregation capabilities between the Access Network and the Regional Network. It is the first point in the network where traffic on multiple DSL lines will be aggregated onto a single network. |
| **Broadband Remote Access Server (BRAS)** | The BRAS is the aggregation point for the subscriber traffic. It provides aggregation capabilities (e.g. IP, PPP, ATM) between the Regional/Access Network and the NSP or ASP. Beyond aggregation, it is also the injection point for policy management and IP QoS in the Regional/Access Networks. |
| **Core Network** | The center core of the Regional Network. The functions contained herein are primarily transport oriented with associated switching or routing capabilities enabling the proper distribution of the data traffic. |
| **Downstream** | The direction of transmission from the ATU-C (Access Node) to the ATU-R (modem). |
| **Dropping** | The process of discarding packets/cells based on specified rules, which may be the result of for example, a policing action or policy decision. |
| **Edge Network** | The edge of the Regional Network. The Edge Network provides access to various layer 2 services and connects to the Regional Network core enabling the distribution of the data traffic between various edge devices. |
| **LNS Group** | A configured set of LNSs for a given NSP. This set of LNSs may be used for load-balancing, redundancy, etc. The LNS group is either configured locally, or returned via RADIUS Tunnel-Server-Endpoint attribute(s). |
| **Loop** | A metallic pair of wires running from the customer's premises to the Access Node. |
| **Many-to-Many Access Sessions** | The ability for multiple individual users or subscribers, within a single premises, to simultaneously connect to multiple NSPs and ASPs. |
| **Microflow** | A single instance of an application-to-application flow of packets, which may for example be classified by source address, source port, destination address, destination port and protocol id, or stateful means. |
| **Network Access Identifier (NAI)** | The user ID submitted by the client during PPP authentication as defined in RFC 2486. |

| | |
|---|---|
| **Policy** | A set of rules to administer, manage, and control access to network resources [RFC3198]. |
| **Profile** | A set of data that may also include policy rules. For example a profile may include data associated with a specific subscriber that could include billing address. The subscriber "profile" could also have a QoS policy rule that describes how their data should be treated. |
| **Regional Network** | The Regional Network interconnects between the Network Service Provider's network and the Access Network. A Regional Network for DSL connects to the BRAS, which is technically both in the Regional Network and in an Access Network. Typically more than once Access Network is connected to a common Regional Network. The function of the Regional Network in this document goes beyond traditional transport, and may include aggregation, routing, and switching. |
| **Regional/Access Network** | The Regional and Access Networks – grouped as and end-to-end QoS domain and often managed by a single provider. |
| **Routing Gateway** | A customer premises functional element that provides IP routing and QoS capabilities. It may be integrated into or be separate from the modem. |
| **Session** | A logically identifiable relationship formed between two (or more) communicating entities for exchanging control and data packets. An example of which would be a PPP session. |
| **Subscriber** | The client that is purchasing the DSL circuit from the Service Provider and is receiving the billing. |
| **Traffic Classification** | The process of selecting packets based on common criteria, such as the content of packet headers or session identification. |
| **Traffic Marking** | The process of setting packet header fields, such as DSCP, MPLS EXP or 802.1p/q COS field in a packet/frame/cell based on defined rules. Traffic marking may result from for example, a classification decision, a policing action, or a policy decision. |
| **Traffic Metering** | The process of measuring the rate and/or burst of a traffic stream selected by a classifier. The instantaneous state of this process may be used to affect the operation of a marker, shaper, or policer, and/or may be used for accounting and measurement purposes. |
| **Traffic Policing** | The process of dropping, marking or remarking packets/cells within a traffic stream in accordance with the state of a corresponding meter against a defined traffic profile, using mechanisms such as the token bucket scheme defined by [RFC2697]. |
| **Traffic Remarking** | The process of changing header fields, such as DSCP, MPLS EXP or 802.1p/q COS field in a packet/frame based on defined rules. |
| **Traffic Shaping** | The process of delaying packets/cells within a traffic stream to cause it to conform to some defined traffic profile. |
| **Traffic Stream** | a set of one or more microflows or sessions, which are selected by a particular classifier. |

| | |
|---|---|
| **Tunnel Group** | A named group of L2TP tunnels between a BRAS (LAC) and an LNS (or LTS) or set of LNS (or LTS). The tunnel group is used to represent a logical connection between a given NSP and a BRAS (or set of BRASs). The individual L2TP tunnels that are members of the tunnel group can be configured for load balancing of traffic and/or for redundancy. While the configuration of the tunnel group (number of member tunnels, load balancing rules, etc.) may differ from BRAS to BRAS, the tunnel group name is global and can be returned via RADIUS. Newly established PPP sessions may be directed to a tunnel group for an NSP independent of the tunnel group's configuration or state on a given BRAS. |
| **Tunnel Server Endpoint** | RADIUS attribute defined in section 3.1 of RFC2868. The Tunnel Server Endpoint attribute may be used to return the IP address of one or more LNSs (or LTSs) at an NSP. |
| **Upstream** | The direction of transmission from the ATU-R (modem) to the ATU-C (Access Node). |
| **User** | Typically, a member, employee or guest at the Subscriber's household or business using the DSL circuit capabilities. |

# 2 General Requirements

R-2-01 The device MUST have redundant Stratum 3 (or better) internal oscillators for node timing and meets or exceeds the synchronization requirements in Section 4.6, Issue 1 revision 2, of Telcordia GR-1110-Core.

R-2-02 The BRAS MUST be non blocking i.e. the forwarding capacity of the internal implementation (e.g. switch fabric, forwarding plane, etc) must be equal or exceed that of the incoming interfaces.

R-2-03 The device MUST not have 'head-of-line blocking' problems

R-2-04 Traffic overload on one port MUST not affect the normal behavior of other ports

## 2.1 Availability

R-2-05 The device SHOULD provide overall availability (hardware and software) excluding scheduled maintenance of 99.999%

R-2-06 All components MUST be hot swappable.

R-2-07 It MUST be possible to make all appropriate configuration changes and software upgrades on the running system, without affecting active users.

R-2-08 Automatic non-revertive switch over from a failed card to a redundant one MUST be supported. Manual switch over should be supported as an option.

R-2-09 All components of the device MUST be configurable to provide either 1:1, N:1, or a distributed equipment redundancy capability across all of the following system components, if present, including:

a) Switching Fabric

b) Packet Forwarding Engine

c) Fabric Interfaces

d) Control and Route Processors

e) Physical Interfaces Trunk/Line Cards including APS (e.g., SONET GR253) and other interface card redundancy capabilities. All ATM, POS, Ethernet, and channelized interface line modules up to OC-48 MUST support N:1 redundancy.

f) Management System/Mgmt Interfaces

g) Power Converters/Supplies

h) Fans

i) Power Feeds

R-2-10 The switch over from primary to secondary (redundant) power input MUST be automatic and cause no disruption of service.

R-2-11 Switch-over procedures for controller, fabric cards, if present, and power supply MUST function correctly when operated in redundant mode.

R-2-12 The device MUST support Ethernet redundancy with a recovery on a point-to-point Gigabit Ethernet connection within 800 ms.

R-2-13 The device MUST maintain operational state for VCs, PPP, or RFC 2684 sessions when it switches to a redundant control processor, and preserve all relevant functions or protocols bound to them.

R-2-14 [North America] The device MUST support APS 1+1 port protection on SONET ports as a configurable option.

R-2-15 The device MUST support a forwarding plane detection and recovery (e.g. switch fabric detection and switch over) switchover time within 60 ms.

R-2-16 The device MUST limit the duration of control plane outages to 2 seconds, specifically session establishment and session/policy updates.

R-2-17 The device MUST maintain all sessions (e.g. ATM, PPP or IP) in any single component failure scenario.

R-2-18 A fully configured device SHOULD be fully operational within 10 minutes from a cold start condition.

R-2-19 The device MUST continue forwarding when one control card is removed from or fails in the platform if a redundant card is present.

R-2-20 The device MUST support the re-insertion of a control card without disrupting forwarding of traffic

## 2.2 Chassis Requirements

R-2-21 Any interface card MUST be able to be plugged into any available slot not reserved for control or fabric modules without restrictions or limitations.

R-2-22 The cooling system MUST be redundant and hot swappable. Should a fan fail, the remaining fans MUST be able to cool a fully loaded unit.

R-2-23 Temperature sensors MUST be used to monitor the system's temperature. Should the programmed temperature thresholds be exceeded, an alarm indication MUST be generated and forwarded to the appropriate alarm/fault management system.

R-2-24 Visual indicators (LEDs) MUST be available for indication of status of the cooling system or temperature across the unit (i.e. Active/Failed Fan, Hot/Normal temperature)

R-2-25 The device SHOULD provide multiple or cost effective chassis designs for small, medium, and large central office locations. The requirements contained within this document pertain to equipment appropriate for large central office deployments. The feature requirements contained in this document (with the exception of scaling attributes) apply to any size device.

## 2.3 Power Requirements

R-2-26 The power supply (modules) MUST accept DC power.

R-2-27 An AC powered option SHOULD be available.

## 2.4 Scalability and Performance

The table below provides scalability metrics for the device. Those requirements that are on a per interface basis are required independent of the size of the device (large CO, medium, or small chassis designs), while the per device requirements are specific only to the large CO device.

| Scope | Per DS3 ATM | Per OC3c/STM1 | Per OC12c/STM4 | Per Large CO Device |
|---|---|---|---|---|
| Provisioned PVCs<br>MST /SHD | --- / 8,000 | --- / 24,000 | --- / 64,000 | --- / 256,000 |
| Active PVCs with at least one IP or PPP interface<br>MST / SHD | 2,000 / 4,000 | 8,000 / 12,000 | 16,000 / 32,000 | 64,000 / 128,000 |
| Bridged 2684 IP + PPP Sessions<br>(1.5 x Active PVCs) | --- | --- | --- | 96,000 / 192,000 |
| IP Interfaces<br>(1.5 x Active PVCs) | --- | --- | --- | 96,000 / 192,000 |
| Total # of Triggered RIP Updates | | | | 64,000 / 128,000 |
| Total # of RIPv2 Updates | --- | --- | --- | 5% of the total number of active VCs that the device can terminate |
| Total # of Triggered RIP or RIPv2 Route Adv. | --- | --- | --- | 50 per host |

R-2-28 The device MUST support RFC 2544, Benchmarking Methodology for Network Interconnect Devices, requirements.

R-2-29 The throughput capability of the forwarding engine function of the device, in terms of Packets Per Second (PPS) SHOULD equal to the sum of the rate of all the interface types that can exist in a valid configuration of the device.

R-2-30 PPS performance SHOULD be sufficient to fill all types of supported interfaces, with 64 byte IP packets at line rate.

## 2.5 Cable Management

R-2-31 The device unit SHOULD provide appropriate cable handling to the rear of the chassis that allows easy access and clean cable routing.

## 2.6 Software Status and Quality Process

R-2-32 Software updates MUST be available on-line.

R-2-33 The device SHOULD be certified to meet CMM Level 3for process for software development and quality assurance.

R-2-34 The device SHOULD be certified to meet ISO 9001 process for software development and quality assurance .

R-2-35 [North America] The engineering and manufacturing process of the device MUST be TL-9000 certified

# 3 Physical Interfaces

## 3.1 General

R-3-01 The device MUST have a craft access console (through RS-232 VT-100 type terminals).

R-3-02 The device SHOULD have redundant Building Integrated Timing Source (BITS) inputs for external node timing.

R-3-03 The device MUST be able to derive timing from channelized interfaces.

R-3-04 Any interface MUST be able to be used as either access or trunk without impacting the functionality of the interface.

R-3-05 The device MUST support physical loop back on a per port basis on all SONET and TDM interfaces.

## 3.2 SONET

R-3-06 [North America] The device ‘s SONET interfaces MUST support implementation of SONET 1+1 Automatic Protection Switching (APS) that is compliant with requirements in section 5.3.2.1, issue 2 of GR-253-CORE.

R-3-07 [North America] The device’s SONET interfaces MUST be compliant with GR-253-CORE including SONET physical layer management capabilities that are compliant with the requirements for memory administration, alarm surveillance, performance monitoring (near and far end, testing processes, and control features in section 6, of issue 2, of Telcordia GR 253-CORE.

R-3-08 [North America] The optical interfaces SHOULD support the use of exchangeable or tunable ITU-grid lasers, such that a given output may be configured to work directly with passive WDM transport gear.

R-3-09 [North America] All optical interfaces MUST support single-mode operation.

R-3-10 [North America] All optical interfaces SHOULD support multi-mode operation.

R-3-11 [North America] All optical interfaces SHOULD be available in short range, intermediate range, and long range laser configurations.

## 3.3 TDM

Requirements for what interface types will support which protocol instances is provided in the table below (MST=MUST; SHD=SHOULD; NA=not applicable):

| Interface Module | PPP | ATM | MPLS |
|---|---|---|---|
| DS3 | **MST** | **MST** | **NA** |
| Chan OC-3 | **MST** | **MST** | **MST** |
| OC-3c | **MST** | **MST** | **MST** |
| Chan OC-12 | **MST** | **MST** | **MST** |
| OC-12c | **MST** | **MST** | **MST** |
| OC-48c | **MST** | **SHD** | **MST** |
| STM-1 | **MST** | **MST** | **MST** |
| STM-4 | **MST** | **MST** | **MST** |
| STM-16 | **MST** | **MST** | **MST** |

## 3.4 Ethernet

R-3-12 The device MUST support 10/100 Mbps, and 1000 Mbps twisted pair and fiber optic requirements of "IEEE Std 802.3, 2002 Edition, Carrier sense multiple access with collision detection CSMA/CD) access method and physical layer specifications" including:

- 100Base-TX
- 1000Base-LX (also known as 1000Base-LH) using GBIC or SFP-GBIC
- 1000Base-SX using GBIC or SFP-GBIC
- 1000Base-ZX using GBIC or SFP-GBIC.

R-3-13 The device MUST support auto-negotiation on all copper 10/100 Ethernet interfaces.

R-3-14 Optical Gigabit interfaces line cards MUST support GigaBit Interface Converter or Small Form-factor Plugable modules.

R-3-15 The device MUST support IEEE 802.3ad link aggregation / load sharing functionality on physical ports.

# 4 Protocols

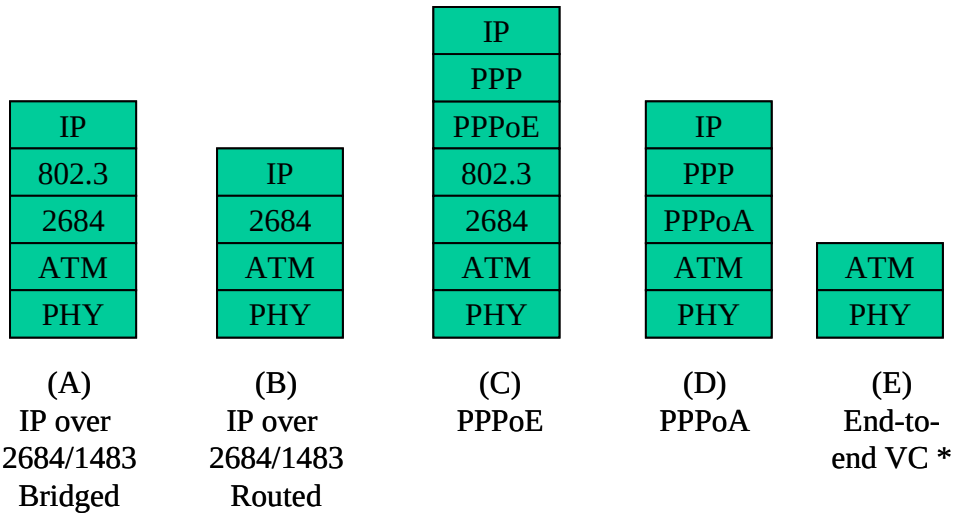


* Stacks E & X are required when the device supports the ATM cross connect function

**Figure 4-1 - Protocols Stacks**

R-4-01 Stacks A, through E illustrate the protocols that MUST be supported between the BRAS device and the end user across the U interface. Stacks Y and Z MUST be support between the BRAS and NSPs or BRAS to BRAS.

R-4-02 The device MUST be able to aggregate end users sessions using protocol stacks (A) through (D) of Figure 4-1 into any of the aggregation stacks (Y) or (Z) of Figure 4-1.

R-4-03 Multiple simultaneous sessions received from one end user across a single PVC MUST be able to be aggregated differently from one another using either stacks (Y) or (Z).

R-4-04 The device MUST support termination and routing of different protocol stacks coming from one PVC simultaneously. Specifically IP over 2684 and PPPoE over 2684 MUST be able to be dealt with when coming across a single PVC.

R-4-05 The device MUST simultaneously support the following aggregation services described in the following sections 4.4 PPP/PTA, 4.8 LAC, and 4.3 Bridged Ethernet.

## 4.1 IPv4

R-4-06 The device MUST support IPv4 as defined in RFC 791, Internet Protocol.

R-4-07 The device MUST support the reassembly of fragmented packets (RFC 791, Internet Protocol).

R-4-08 The device MUST support ICMP as defined in RFC 792.

## 4.2 ATM

R-4-09 The device MUST support ATM VP and VC aggregation and termination

R-4-10 The device MUST support ATM VP, VC cross connect capabilities

R-4-11 The device MUST preserve ATM class of service if providing a VC, VP cross connect function

R-4-12 The device SHOULD support ATM Forum User-Network Interwork Interface (UNI) Specification Version 3.1.

R-4-13 The device MUST support the ATM UBR service class, as defined by the ATM Forum.

R-4-14 The device SHOULD support the ATM UBR with MDCR service class, as defined by the ATM Forum Traffic Management Specification 4.1, af-tm-0150.000.

R-4-15 The device MUST support the ATM nrt-VBR service class, as defined by the ATM Forum Traffic Management Specification 4.1.

R-4-16 The device MUST support the ATM rt-VBR service class, as defined by the ATM Forum Traffic Management Specification 4.1.

R-4-17 The device MUST support the ATM CBR service class, as defined by the ATM Forum Traffic Management Specification 4.1.

R-4-18 The device SHOULD support ATM OAM VC Management per I.610.

R-4-19 The device MUST support the full range of VPI and VCI values. These ranges MUST be fully configurable.

R-4-20 The device MUST support ATM F4 and F5 OAM loop back, AIS , and RDI (remote defect indication) on all ATM interfaces per I.610 specification of the ITU-T.

### 4.2.1 ATM Scalability

R-4-21 The device MUST support at least 256 virtual paths.

R-4-22 The VC counts supported on the device (active and configured) MUST not be affected by the upper layer protocols provisioned on VC.

## 4.3 RFC 1483/2684

R-4-23 The device MUST be able to terminate end user sessions using protocol stack (A) and (B) of Figure 4-1 IP over bridged Ethernet.

R-4-24 The device MUST support RFC 2684 LLC/SNAP encapsulation of all bridged Ethernet frames (i.e., IEEE 802.3 and PPPoE) over ATM AAL5.

R-4-25 The device MUST be able to process IP over RFC 2684 (VC MUX) and route the IP packets across to any other interface type.

R-4-26 The device MUST be able to auto-sense PPP frames within RFC 2684 packets received from the customer.

## 4.4 PPP

R-4-27 The device MUST be able to terminate end user sessions using protocol stacks (C) and (D) PPPoE (RFC 2516) and PPPoA (RFC 2364).

R-4-28 The device MUST support RFC 1661, The Point-to-Point Protocol (PPP).

R-4-29 The device MUST support RFC 1570, PPP LCP Extensions.

R-4-30 The device MUST support LCP based MTU re-negotiation.

R-4-31 The device MUST support RFC 1332, PPP IP Control Protocol (IPCP).

R-4-32 The device MUST support PAP as defined in RFC 1334, PPP Authentication Protocols.

R-4-33 The device MUST support CHAP as defined in RFC 1334, PPP Authentication Protocols.

R-4-34 The device MUST support the following authentication sequences:

a. PAP only

b. CHAP only

c. CHAP then PAP

R-4-35 The device SHOULD support extensions to IPCP to include primary DNS, secondary DNS and NBNS IP addresses per RFC 1877 – PPP Internet Protocol Control Protocol Extensions for Name Server Addresses.

R-4-36 IP subnet masks MUST be communicated with IPCP using the PPP IPCP option with option code144, the length of the option being 6 and the mask being expressed as a32-bit mask (e.g. 0xFFFFFF80), not as a number indicating the consecutive number of 1s in the mask (from 0 to 32).

R-4-37 The device MUST be able to terminate PPP coming in using PPPoE and/or PPPoA and forward the IP packets across any interface (excluding management interfaces) such that there are no restrictions.

R-4-38 The device MUST be able to inter-work with PPPoE client implementations that support PPPoE RFC 2516.

R-4-39 The device MUST properly establish the PPP (LCP) link (as described in RFC 1661) prior to accepting user ID and authentication (this is standard operation for PPP).

R-4-40 The device MUST support a variety of delimiters for NAI parsing, including but not limited to / and @ when examining the userid passed during PPP authentication.

R-4-41 If a PPP session is terminated for any reason, the device MUST gracefully shutdown the associated PPPoE or PPPoA session.

R-4-42 The device MUST shutdown and clean-up all PPPoE or PPPoA sessions if the PVC is deleted for any reason, and update any information related to these sessions (such as RADIUS Accounting, routing, etc.).

R-4-43 The device MUST detect and drop any PPPoE packets if its session packets are found to be inconsistent with the MAC address recorded during the discovery phase.

R-4-44 The device MUST gracefully shutdown a PPPoE or PPPoA session upon receiving a PADT packet from the subscriber associated with the PPPoE or PPPoA session.

R-4-45 The device MUST drop (do nothing to process) PPPoE or PPPoA packets that refer to an invalid (ended or unrecognized) session ID.

R-4-46 The device MUST be able to alarm or log the condition under which a subscriber sends PPPoE or PPPoA messages with invalid session ID or MAC address.

R-4-47 The device MUST support the ability to control PPP session time limits based on NAI.

R-4-48 The device MUST support a configurable limit for the total number of PPP sessions on an access VC.

R-4-49 The device MUST support multiple PPPoE sessions on a single subscriber PVC and have the ability to encapsulate PPP packets from one PPPoE session onto a Layer 2 Tunneling protocol (L2TP) tunnel while terminating the other PPP sessions.

R-4-50 The device MUST support a PPPoE subscriber receiving service from L2TP and PTA NSPs simultaneously. One ADSL VC MUST support access to one or more L2TP NSPs and to one or more PTA NSPs.

R-4-51 The device MUST be able to limit the number of PPPoE sessions for each FQDN (domain/realm) on an access VC.

R-4-52 The device MUST support a PPP-based multiple destination selection service. Destinations are identified by the NAI of the PPP session.

R-4-53 The device MUST support the ability to assign access lists to subscribers that attach to specific FQDNs (domains/realms).

R-4-54 The device MUST support domain/realm/destination limiting/filtering. Each access VC is configurable with a list of valid destinations and realms. On that VC, subscriber that attempts to connect to a destination not in the list of valid destinations MUST be rejected.

R-4-55 The device MUST support using both the VPI/VCI of the ATM PVC that delivered the PPP session and/or the NAI provided by the user during the PPP authentication phase, to determine if a PPP session is to be L2TP tunneled.

R-4-56 The device MUST be able to determine the endpoint to which a PPP session is to be L2TP tunneled (a particular tunneling endpoint i.e., LNS) either by equating the VPI/VCI of the ATM PVC on which the PPP session was received or by using the NAI provided by the user during the PPP authentication phase.

R-4-57 If both a NAI and a VPI/VCI association are found, the VPI/VCI association MUST by default have precedence if the NAI points to different tunneled endpoints. However this behavior could be overridden by the VPI/VCI RADIUS authentication server, which could reject the session.

R-4-58 The device MUST support exception binding. PPP packets from all the PPPoE or PPPoA sessions coming in on a subscriber PVC are forced onto a given L2TP tunnel by default

unless the NAI matches a predefined character string (e.g. “myISP”). In that case the PPP session is terminated on the device and routed to the appropriate service provider.

R-4-59 The device SHOULD support the Exclusive PPP Session Feature (Enhanced security for corporate access).

When the Exclusive PPP Session feature is activated, access to a particular secure “exclusive” destination is allowed, but it is not possible to simultaneously access another destination over the same DSL VC.

- If a session to an exclusive destination is active, attempts to set up sessions to other destinations are rejected.
- If there are any active sessions to other destinations, attempts to set up a session to the exclusive destination are rejected.
- Multiple sessions to the exclusive destination are allowed, provided this is consistent with all other service specifications.

R-4-60 The device MUST support routing of IP datagrams recovered from terminated PPP sessions using Policy Routing. At a minimum, the product MUST allow a policy whereby the source IP address contained in the IP datagram is used to determine how to forward the IP datagram.

R-4-61 The device MUST be capable of supporting PPP termination and aggregation for sessions carrying IP packets that have private IP addresses and routing them to the correct NSP based on the NAI supplied during PPP authentication.

R-4-62 The device MUST be capable of supporting PPP termination and aggregation for sessions carrying IP packets that have private IP addresses and routing them to the correct NSP based on the incoming PVC’s VPI/VCI.

R-4-63 In PTA mode, the device MUST support per PPP session inactivity timers that can be configured to trigger the tear down of an PPP session that has been inactive for a configurable period of time.

R-4-64 The device MUST have the option to strip the domain name extension before forwarding the user name to a RADIUS servers.

### 4.4.1 PPP Scalability

R-4-65 The auto detection of access PPP sessions (PPPoE or PPPoA) SHOULD NOT affect the scalability metrics of the device.

R-4-66 The device MUST support establishing PPP sessions at a minimum rate of 100 PPPoE or PPPoA sessions per second (including IP address negotiation and route installment) with a directly attached RADIUS server, up to the maximum number of supported sessions on the device. Assuming that the RADIUS server is sized such that it is not a bottleneck.

R-4-67 The device MUST not contribute more that 300 msec to the PPP session setup time (including IP address negotiation and route installment).

## 4.5 Ethernet

R-4-68 The device MUST learn dynamically any MAC address supported on the RFC 1483/2684 bridged ATM PVC by means of ARP.

R-4-69 The device MUST only send unicast traffic to a port where is has learned its MAC address to keep subscriber MAC addresses from being snooped from other ports.

R-4-70 The device MUST support at least 1550 byte Ethernet payloads on physical Ethernet interfaces for placing 1500 byte IP packets into L2TP without requiring fragmentation.

R-4-71 The device MUST support at least 1550 byte Ethernet payloads on logical RFC 2684 Ethernet Interfaces.

R-4-72 The device MUST support IEEE 802.1p (Traffic Class Expediting).

R-4-73 The device MUST support VLANs as specified in “IEEE Std 802.1Q-1998, Virtual Bridged Local Area Networks” including classification, traffic management, and tagging.

R-4-74 The device SHOULD be able to use a 802.1Q VLAN interface as a logical IP interface (e.g. for an IP VPN).

R-4-75 The device MUST not permit the association of a particular Ethernet MAC address on more than one subscriber PVC at the same time.

R-4-76 The device MUST age MAC addresses that are dynamically learned and remove them from the MAC address table if the device does not send traffic to the MAC address or receive traffic from the MAC address for a certain period of time. This period of time is called the “MAC address timeout interval”.

R-4-77 The device MUST allow a configurable value for the MAC address timeout interval.

### 4.5.1 Ethernet Scalability

R-4-78 The device MUST support at least as many MAC addresses as active access VCs.

R-4-79 The device must provide a means to limit the number of MAC addresses learned on any given VLAN port.

R-4-80 The device MUST be able support the maximum number of available MAC addresses, per R-4-78,across any physical or virtual port.

R-4-81 The device MUST support an ARP table entries with at least as many entries as VCs.

R-4-82 The device MUST support 4,096 VLANS.

R-4-83 The device MUST support placing all subscriber traffic into a single VLAN on an uplink..

## 4.6 MPLS

R-4-84 The device MUST support RFC 3031, Multiprotocol Label Switching Architecture.

R-4-85 The device MUST support RFC 3032, MPLS Label Stack Encoding.

R-4-86 The device SHOULD support 3 levels of labels (label stacking).

R-4-87 The device MUST support RFC 3036, LDP Specification.

R-4-88 The device MUST support Multi-protocol Label Switching OAM mechanisms (e.g. “Detecting MPLS Data Plane Failures", draft-ietf-mpls-lsp-ping-05.txt).

R-4-89 The device MUST support fast re-route of LSPs capability (e.g. draft-ieft-mpls-rsvp-lsp-fastreroute.01.txt).

R-4-90 The device MUST use liberal label retention per RFC 3031.

R-4-91 The device MUST use independent label distribution control in RFC 3036.

R-4-92 The device MUST support Downstream Unsolicited label distribution per RFC 3036

R-4-93 The device MUST support E-LSPs per RFC 3270

R-4-94 The device SHOULD support L-LSPs per RFC 3270

R-4-95 The device MUST support Multi-protocol Extensions for BGP-4 per RFC 2283.

R-4-96 The device MUST support BGP Site of Origin.

R-4-97 The device MUST support BGP AS Override.

R-4-98 The device MUST support MPLS-BGP VPNs as specified in RFC2547 to provide virtual routing to multiple NSPs and ASPs.

R-4-99 The device MUST support LDP Filtering.

R-4-100 The device MUST support RFC 3209 - RSVP-TE: Extensions to RSVP for LSP Tunnels.

### 4.6.1 MPLS Scalability

R-4-101 The device MUST support at least 2,000 LSP instances

R-4-102 The device SHOULD support at least 10,000 LSP instances

R-4-103 The device MUST support at least 500 VPNs.

R-4-104 The device MUST support at least 2,000 customer sites participating in a 2 way dynamic routing protocol supported per VPN.

R-4-105 The device MUST support placing all active sessions (PPP, IP over RFC2684, L2TP) across multiple or in a single VPN or VRF.

R-4-106 The device MUST support minimum 1000 VRFs/PE.

## 4.7 L2TP

R-4-107 The device MUST support L2TP over the User Datagram Protocol (UDP) over Internet Protocol (IP) (UDP/IP) as defined in “Layer Two Tunneling Protocol ‘L2TP’,” RFC 2661.

R-4-108 The device MUST be able to function simultaneously as a LAC and LNS.

R-4-109 The device MUST support initiating and terminating L2TP tunnels on loop back addresses of the device.

R-4-110 The device MUST support multiple loop back addresses per virtual router.

R-4-111 The device SHOULD support L2TP Disconnect Cause Information (RFC 3145)

## 4.8 L2TP Access Concentrator (LAC) Requirements

R-4-112 Support LNS renegotiation - when BRAS acts as a LAC, it MUST allow the LNS to force LCP renegotiation. The B-RAS MUST support all LAC and LNS functions and messages related to “Proxy LCP and Authentication AVPs” as described in RFC 2661.

R-4-113 The device MUST terminate (clean up) L2TP tunnels if PVC transporting the tunnels is deleted.

R-4-114 The device MUST terminate (clean up) L2TP tunnels if the interface associated with the tunnel is deleted.

R-4-115 The device MUST terminate (clean up) L2TP tunnels if the SP (LNS) for the tunnels is deleted.

R-4-116 The device MUST terminate (clean up) L2TP tunnels if the virtual router associated with the tunnel is deleted.

R-4-117 The device MUST be able to rely on the PPP timeout mechanism or other method to clean up the call associated with an L2TP tunnel automatically if the associated PPPoE or PPPoA is terminated for any reason.

R-4-118 The device MUST support Static and Dynamic tunnel creation. Static tunnel creation means that the L2TP tunnel will be up after correct provisioning. Dynamic tunnel creation means that the L2TP tunnel will only be up when PPP sessions destined for that tunnel are present.

R-4-119 The device when acting as a LAC MUST be able to aggregate PPP sessions from PPPoA, PPPoE, and L2TP onto a single L2TP tunnel (RFC 2661).

R-4-120 The device MUST support L2TP tunneling of PPP frames from multiple PPP sessions.

R-4-121 The device MUST support multiple L2TP tunnels per physical or virtual interface.

R-4-122 The device MUST be able to support multiple L2TP tunnels to a single NSP.

R-4-123 The device MUST support tunnel groups where multiple L2TP tunnels can be created between a pair of LAC and LNS and be assigned to a tunnel group. The device MUST support the capability to configure a tunnel which can carry PPP sessions with multiple FQDNs (domains/realms).

R-4-124 After a PPPoE or PPPoA session is established between a subscriber and the device, the device MUST be able to select (or create) an L2TP tunnel to carry PPP traffic from the PPPoE or PPPoA session based on the domain name provided as part of the NAI presented during PPP authentication.

R-4-125 The device MUST not process or alter the NAI or authentication information other than examining the user name extension so that the PPP session can be properly routed.

R-4-126 The device MUST pass all PPP user ID and authentication information to the SP over the L2TP tunnel supporting the SP.

R-4-127 The device MUST not inhibit the use of IPCP between a NSP and a subscriber.

R-4-128 The device MUST establish a PPPoE or PPPoA session instance prior to establishing an associated L2TP session.

R-4-129 The device MUST terminate a PPPoE or PPPoA session instance upon termination of an associated L2TP session.

R-4-130 If a tunnel is not present the device MUST dynamically create a new one.

R-4-131 The device MUST refuse any new PPPoE or PPPoA sessions if the maximum number of PPP sessions and L2TP tunnels is reached (per port, device, etc).

R-4-132 The device MUST support the routing of a PPP session to a particular tunnel group based on RADIUS response fields.

R-4-133 The device MUST be able to terminate L2TP tunnels dynamically when L2TP tunnels are not carrying any PPP sessions.

R-4-134 The device MUST gracefully shutdown all PPP sessions associated with an L2TP tunnel if the L2TP tunnel is terminated for any reason.

R-4-135 The device MUST support load balancing of PPP sessions between L2TP tunnels in a tunnel group or LNS Group.

R-4-136 The device MUST support directing PPP sessions across multiple tunnels on a strict priority basis so that the first tunnel of the Tunnel Group or LNS Group fills, then the second tunnel of the Tunnel Group or LNS Group, etc. The device MUST support load balancing of PPP sessions across multiple tunnels on a weighted basis (e.g., 75% of session requests directed to tunnel 1, 25% to tunnel 2).

R-4-137 If a tunnel in a tunnel group is not available, then the PPP sessions MUST continue to load balance across the remaining tunnels.

R-4-138 The device MUST support the capability to add a tunnel to or delete a tunnel from a tunnel group, without disruption of the other tunnels in the group or the PPP sessions in those tunnels, and subsequently load balance over the resulting tunnels in the group.

R-4-139 The device MUST support the ability to cap the number of PPP sessions a L2TP tunnel can support.

R-4-140 The device MUST support fail over configuration between L2TP tunnels in a tunnel group or LNS Group.

R-4-141 The device MUST support the establishment of L2TP tunnels between loop back addresses.

R-4-142 The device MUST support the routing of a PPP session to an LNS or LNS Group based on the RADIUS Tunnel Server Endpoint attribute defined in section 3.1 of RFC2868.

R-4-143 If a LNS in an LNS Group is not available, then the PPP sessions MUST continue to load balance across the remaining LNSs

### 4.8.1 L2TP Scalability

R-4-144 The device MUST support at least 4000 L2TP tunnels.

R-4-145 The device MUST support at least 8,000 PPP sessions per L2TP tunnel (Non cumulative with R-4-144).

R-4-146 The device MUST support at least 1000 L2TP tunnels per virtual router up to the limits specified above.

R-4-147 The device MUST support allocating L2TP across any number of virtual routers or all configured in a single VR.

R-4-148 The device MUST support at least 8 tunnels per tunnel group.

## 4.9 IP Routing and Protocol Support

R-4-149 The device MUST support RFC 1812, Requirements for IP Version 4 Routers, and STD 005, Internet Protocol

R-4-150 The device MUST support RFC 1918, Address Allocation for Private Internet Space.

R-4-151 The device MUST support RFC 1519, Classless Inter-Domain Routing(CIDR): an Address Assignment and Aggregation Strategy.

R-4-152 The device MUST support ICMP echo request and reply, ICMP error handling, ICMP redirect, ICMP source address configuration (RFC 792, Internet Control Message Protocol).

R-4-153 The device MUST support ping and traceroute per VR.

R-4-154 The device MUST allow route advertisement Interval control.

R-4-155 The device MUST support Equal Cost Multipath across multiple default routes and multiple static routes.

R-4-156 The device MUST support Reverse Path Forwarding (Reverse Route Look-up).

R-4-157 The device MUST support disabling of directed domain broadcasting.

R-4-158 The device MUST support route summarization.

R-4-159 The device MUST support route redistribution between routing protocols.

R-4-160 The device MUST support the configuration of primary and secondary IP addresses on trunk interfaces

R-4-161 The device MUST support the use of unnumbered IP interfaces.

R-4-162 The device MUST support at least 50K entries in the RIB.

R-4-163 The device MUST support at least 50K entries in the FIB.

R-4-164 The device SHOULD support at least 250K entries in the RIB.

R-4-165 The device SHOULD support at least 250K entries in the FIB.

R-4-166 The device SHOULD support at least 10,000 IGP routes.

### 4.9.1 OSPF

R-4-167 The device MUST support OSPF version 2 as defined in RFC 2328.

R-4-168 The device MUST support RFC 2370, The OSPF Opaque LSA Option.

R-4-169 The device MUST support RFC 3137, OSPF Stub Router Advertisement.

R-4-170 The device MUST support RFC 1587, The OSPF Not-So-Stubby Area (NSSA) Option.

R-4-171 The device MUST be able to act as an ABR and an ASBR.

R-4-172 The device MUST support at least 100 OSPF adjacencies per instance of OSPF.

R-4-173 The device MUST support at least as many instances of OSPF on the platform as virtual routers supported by the device.

R-4-174 The device MUST support at least 5000 routes within a given OSPF area.

R-4-175 The device SHOULD support an OSPF graceful restart capability (e.g. draft-ietf-ospf-hitless-restart-08.txt).

R-4-176 The device SHOULD support OSPF sham links over 2547bis VPNs (e.g. draft-ietf-l3vpn-ospf-2547-00.txt).

R-4-177 The OSPF version-2 protocol, if used, MUST employ cryptographic authentication, as specified in RFC 2328.

### 4.9.2 BGP

R-4-178 The device MUST support BGP-4 as defined in RFC 1771 including the following BGP features/attributes: extended communities, VPN-IP addresses, route acceptance and announcement filters, multihop iBGP/eBGP, authentication, and OSPF or IS-IS LSA types in the extended community attribute.

R-4-179 The device MUST support BGP Outbound Route Filter (ORF) per Cooperative Route Filtering Capability for BGP-4 (e.g. draft-ietf-idr-route-filter-08.txt).

R-4-180 The device MUST support RFC 3065, Autonomous System Confederations for BGP.

R-4-181 The device MUST support BGP Policy-lists. This feature adds the capability for a network operator to group route map match clauses into named lists called policy lists. A policy list functions like a macro. When a policy list is referenced in a route map, all of the match clauses are evaluated and processed as if they had been configured directly in the route map.

R-4-182 The device MUST support RFC 1997, BGP Communities Attribute – these attributes MUST be settable by the device.

R-4-183 The device MUST support RFC 2439, BGP Route Flap Damping.

R-4-184 The device MUST support RFC 2918, Route Refresh Capability for BGP-4.

R-4-185 The device MUST support RFC 2796, BGP Route Reflection – An Alternative to Full Mesh IBGP.

R-4-186 The device MUST support capability negotiation per RFC 3392 (Capabilities Advertisement with BGP-4)

R-4-187 The device MUST support exact matches for BGP community attributes for ingress/egress route filtering/policies.

R-4-188 The device MUST support a minimum of 50 BGP Sessions per device.

R-4-189 The device SHOULD support BGP graceful restart (e.g. draft-ietf-idr-restart-08.txt).

R-4-190 The device MUST support RFC 2858, MBGP for VPN, multicast and IPv6.

R-4-191 The device SHOULD support for 4 byte AS numbers (draft-ietf-idr-as4bytes-07.txt).

R-4-192 The device MUST support RFC 2385, BGP MD5 authentication, and TTL scheme for security

### 4.9.3 ISIS

R-4-193 The device MUST support IS-IS routing protocol (ISO 10589).

R-4-194 The device MUST support configurable IS-IS Incremental SPF Algorithm.

R-4-195 The device MUST support IS-IS Administrative Tags (RT, ext community transparency)

R-4-196 The device MUST support at least 100 IS-IS adjacencies.

R-4-197 The device MUST support at least 5000 routes within a given ISIS area.

R-4-198 The device SHOULD support RFC 3567, ISIS hmac-md5 authentication.

R-4-199 The device SHOULD support draft-ietf-isis-restart-04.txt, ISIS graceful restart.

R-4-200 The device SHOULD support ISIS for multi-topology (e.g. draft-ietf-isis-wg-multi-topology-06.txt).

R-4-201 The device SHOULD support ISIS for point-to-point over LAN (draft-ietf-isis-igp-p2p-over-lan-03.txt).

R-4-202 The device MUST support RFC 2763, ISIS dynamic hostname.

R-4-203 The device MUST support ISIS Transient black hole avoidance (RFC 3277),

R-4-204 The device MUST support 3-way handshake for ISIS Point-to-Point Adjacency (RFC 3373).

### 4.9.4 RIP

R-4-205 The device MUST support RIP version 2 as defined in IETF STD 0056 for route advertisements from the device to customer CPE.

R-4-206 The device MUST support sending RIP updates to customer CPE without listening to updates from the CPE.

R-4-207 The device MUST support Triggered RIP as defined in IETF RFC 2091 for sending route advertisements from the BRAS to the customer CPE

R-4-208 The device MUST support sending Triggered RIP or RIPv2 updates to as at least as many hosts as specified in section 2.4 at a rate at least equivalent to the radius setup rate of the device.

R-4-209 The device MAY support additional mechanisms to send IP routing information to the CPE (e.g. TR-044).

## 4.10 IPv6

The intent of this section is to provide the high level requirements for IPv6 so that vendors will design platforms with the system resources capable of supporting IPv6 when the set of requirements are better understood. This section is not intended to be all encompassing of the requirements for supporting an IPv6 service.

R-4-210 The device SHOULD support RFC 2460, Internet Protocol, Version 6 (IPv6) Specification.

R-4-211 The device SHOULD support RFC 2373, IP Version 6 Addressing Architecture.

R-4-212 The device SHOULD support IPv6 w/ ISIS (e.g. draft-ietf-isis-ipv6-06.txt).

R-4-213 The device SHOULD support RFC2740, OSPFv3 for IPv6.

R-4-214 The device SHOULD support IPv6 NCP, RFC 2472, over PPPoA, PPPoE, and L2TP over IPv4.

R-4-215 The device SHOULD support DHCP-PD for IPv6 (RFC3315).

# 5 IP Services

## 5.1 IP Address Management

R-5-01 The device MUST support the assignment of an IP version 4 address to any routable interface.

R-5-02 The device MUST support the assignment of multiple IP version 4 addresses to a single routable interface.

R-5-03 The device MUST use the Address Resolution Protocol (ARP) to build ARP tables as defined in RFC 828.

R-5-04 The device MUST support anti-spoofing mechanism so that the B-RAS responds to subscriber ARP requests only when they originate with the proper IP source address and are received on the appropriate ATM VC (e.g. secure ARP, or Proxy ARP).

R-5-05 The device MUST allow network operators to add manual entries to ARP tables and to associate one or more IP address with an Ethernet MAC address.

R-5-06 The device MUST support RFC 3046 to perform a DHCP relay agent function in the assignment of IP addresses to end user CPE.

R-5-07 The DHCP relay agent in the product MUST inspect upstream packets to discover IP address and Ethernet MAC address and populate the ARP table.

R-5-08 The DHCP relay agent in the product MUST generate point-to-point DHCP requests on behalf of the client and the DHCP server.

R-5-09 The DHCP Relay agent in the product SHOULD follow the lease time and lease renewal negotiation, and be able to terminate any subscriber sessions based on lease time expired

R-5-10 The device MUST be able to filter traffic to ensure that a subscriber can not use IP address not assigned to them.

### 5.1.1 RADIUS

R-5-11 The device MUST support RADIUS as defined in RFC 2865.

R-5-12 The device MUST support RFC 2868, RADIUS attributes for tunnel protocol support.

R-5-13 The device MUST support RADIUS extensions as RFC 2869.

R-5-14 The device SHOULD support Dynamic Authorization Extensions to RADIUS as RFC 3576.

R-5-15 The device MUST be able to forward RADIUS accounting traffic to accounting servers that are different than the RADIUS authentication servers

R-5-16 Multiple RADIUS accounting servers MUST be able to be specified.

R-5-17 The device MUST support configurable IP addresses for RADIUS requests - The source IP address in the IP packets used to carry RADIUS request messages is configurable in the B-RAS. For example, a B-RAS loop back address can be used as the source IP address, even though the IP packets go through one of many IP interfaces (as in the case of load balancing).

R-5-18 The device MUST support a fail over mechanism for redundant RADIUS servers.

R-5-19 The device MUST be able to access at least 16 RADIUS Accounting and 16 Authorization servers per VR. Each RADIUS server MUST be configurable to act as a primary, secondary, tertiary, … servers.

R-5-20 The device MUST support load-balancing requests across multiple RADIUS servers.

R-5-21 The device MUST provide support for limiting the requests rate to a RADIUS server.

R-5-22 The device MUST drop RADIUS requests that exceed the configured maximum request rate.

R-5-23 RADIUS requests MUST be queued separately from other control traffic.

R-5-24 The device MUST support RADIUS responses indicating a profile or policy that the device then implements.

R-5-25 A single RADIUS client MUST work across multiple virtual routers.

R-5-26 The device MUST support multiple RADIUS clients that can be used to forward and receive authentication and IP address information to/from NSPs or Corporations. The determination as to which RADIUS client to use for a given PPP session MUST be based on the NAI or the incoming PVC.

R-5-27 If the VPI/VCI provides a mapping to a RADIUS client, the device MUST be configurable so that this mapping will take precedence over any NAI provided with the PPP session.

R-5-28 If the VPI/VCI mapping to a RADIUS client conflicts with the NAI mapping to a RADIUS client, the device MUST be configurable to reject the PPP session.

R-5-29 The device MUST support as the default behavior that the VPI/VCI mapping to a RADIUS client takes precedence over any NAI provided with the PPP session.

R-5-30 The listening UDP port of the RADIUS server MUST be configurable out side the standard range of ports.

R-5-31 Backup features such as the number of retries before the switch to the backup server, the time out value, the dead time value MUST be configurable.

### 5.1.2 Address Pools

R-5-32 The device MUST store pools of IP addresses for PPP distribution or DHCP distribution.

R-5-33 An IP address pool MUST be configurable with multiple CIDR blocks.

R-5-34 The device MUST support providing a particular subscriber with the same IP address assignment for every PPP session

R-5-35 The device MUST support limiting the number of IP addresses that can be assigned to a subscriber.

R-5-36 The device MUST prevent a subscriber from using statically assigned IP addresses outside of those assigned to the subscriber. That is, it should NAK any address requests with an IP address outside of the subscriber's authorized pool. Similarly, the device should reject any traffic with a source IP address outside of the subscriber's authorized pool.

R-5-37 The device must be configurable to support multiple statically assigned IP addresses per subscriber. The device must be able to support defining a base IP address and subnet representing an IP address pool for a subscriber.

R-5-38 The device MUST deny assigning an address to a subscriber if their pool is empty.

R-5-39 The device SHOULD support a default address pool

R-5-40 If the NSP specifies an un-named pool in their RADIUS response then the device MUST assign an address out of the default pool if it exists.

R-5-41 If the NSP specifies an un-named pool in their RADIUS response and assignment of an address out of the default pool has not been configured then the subscriber MUST NOT be assigned an address.

R-5-42 The device MUST support private addresses being assigned to a pool.

R-5-43 The device MUST support at least 20 address pools per VR or 2000 for the entire device.

# 6 Traffic Management

## 6.1 ATM traffic management

R-6-01 The device MUST support the over-subscription of the VBR service.

R-6-02 The device MUST support Per-VC queuing.

R-6-03 The device MUST shape cells on a per VC basis to conform with traffic descriptors associated with the given ATM PVC. These traffic descriptors are defined by the ATM service class that is used. They may include PCR, SCR, MBS.

R-6-04 The device MUST support cell shaping using a dual leaky bucket.

R-6-05 The device MUST not exhibit any performance impacts when traffic shaping is turned on.

R-6-06 The device MUST shape and police ATM traffic at both the VC and the VP levels concurrently for ATM traffic that traverses the BRAS (ATM cross connect).

R-6-07 The device MUST support early packet discard when ATM traffic is cross connected.

R-6-08 The device MUST shape ATM traffic at both the VC and the VP levels concurrently.

R-6-09 The device MUST provide support for partial packet discard on a per ATM VC basis. In this mode, when a cell on an ATM VC is dropped due to queue depth being exceeded, all subsequent cells in that AAL5 frame will be dropped.

### 6.1.1 Traffic Classification

R-6-10 The device MUST support per packet classification based upon the following fields:

- DSCP for IPv4 and IPv6
- Incoming port/interface/PPP Session using the FQDN/NAI
- Source IP address
- Destination IP address
- IP Protocol
- Source TCP/UDP port
- Destination TCP/UDP port
- Ether-type
- 802.1P
- 802.1Q
- Packet length

R-6-11 Based on the classification fields defined in R-6-10The device MUST have a means to mark the following values:

- DSCP
- 802.1P
- MPLS EXP bits

R-6-12 The device MUST have a means to collect metrics for example packet count, byte count and cell count, associated with an IP flow, based on any and all possible combinations of the layer 2, 3, and 4 headers including but not limited to the classification fields defined in R-6-09

R-6-13 The device MUST support a variety of delimiters for NAI parsing, including but not limited to / and @.

R-6-14 The parsing order for domain delimiters MUST be configurable.

R-6-15 The location of the domain string relative to the delimiter (before or after) MUST be configurable.

### 6.1.2 Virtual Routing

R-6-16 The device MUST allow the creation of a minimum of 500 virtual routers (i.e., contexts).

R-6-17 Each virtual router (VR) must have a separate routing information base (RIB), forwarding information base FIB, and management information base MIB from all other VRs.

R-6-18 The device MUST allow all sessions to terminate in one VR.

R-6-19 The device MUST allow all PVCs to terminate in one VR.

R-6-20 The device MUST allow dynamic session to VR binding based on configuration data retrieved from a policy repository or AAA response.

R-6-21 The device MUST allow static session to VR binding for Non-PPP based sessions.

R-6-22 The device MUST allow VR binding based on NAI, VC, VLAN, or any other interface type.

R-6-23 The device MUST be able to restrict the VR the customer can access based on NAI inspection.

R-6-24 The device MUST allow both static and dynamic configuration of a session's parameters within the virtual router. Some configuration data may be retrieved via a policy server or via AAA response.

R-6-25 The device MUST have a default VR (a VR to which a session can be bound by default).

R-6-26 The device MUST allow session parameters (rate, ACL, time out) to be applied to a VR or to each session in a VR.

R-6-27 The device MUST support PPP and RFC 2684 based traffic on the same VC where each type of traffic is bound to a different VR.

R-6-28 The device SHOULD support transparent virtual routing as a type of VR (Described in RFC 1812).

R-6-29 The device MUST support the configuration of a default virtual router for domains that are not bound to a specific virtual router.

R-6-30 The device MUST support the configuration of a default virtual router for NAIs that do not contain a domain.

### 6.1.3 QoS, Scheduling, Shaping, and Policing

R-6-31 The device MUST have a common QoS and scheduling subsystem across all protocol layers (layers 2, 3, and above)

R-6-32 The device MUST support at least 4 PPP/IP session per VC.

R-6-33 For one subscriber session on each VC, the device MUST support at least 6 independent subscriber level classes or PHBs.

R-6-34 The device MUST support traffic conditioning at all interfaces at line rate.

R-6-35 The device MUST support RFC 2475, An Architecture for Differentiated Services, including edge-conditioning functions such as Packet Classification, Policing, Shaping, Marking & Metering.

R-6-36 The device MUST support Assured Forwarding per hop behavior per RFC 2597.

R-6-37 The device MUST support Expedited Forwarding per hop behavior per RFC 3246.

R-6-38 The device MUST support the default PHB (RFC 2474) for "best effort" traffic.

R-6-39 The device SHOULD support the Lower Effort PHB (RFC 3662) for “scavenger class” type services.

R-6-40 All traffic parameters for rate limiting, traffic shaping, associated with the PHBs listed above MUST be configurable.

R-6-41 The device MUST support per PPP session shaping for terminated and non-terminated sessions.

R-6-42 The device MUST support per L2TP tunnel shaping.

R-6-43 The device MUST support PPP and IP session level fairness. The device MUST support a configurable minimum through put per session to ensure that starvation below that level does not occur.

R-6-44 The device MUST implement rate shaping capability for each queue.

R-6-45 The device MUST support shaping towards the network at the VR and domain levels (i.e. the traffic aggregated over all PPP sessions and L2TP tunnels on the device for the given @domain is shaped to a given level).

R-6-46 The device MUST meet a 30 ms delay target for queuing and transmission of EF traffic towards the RG/DSL modem.

R-6-47 When required the device SHOULD be able to reduce the packet size in non-EF queues when packets are present in the EF queue to support low jitter traffic.

R-6-48 When fragmentation is required, the device MUST fragment all sessions on an access VC using MLPP interleaving (RFC 1990).

R-6-49 The device SHOULD support an EF window timer associated with fragmenting traffic using MLPP. The EF window timer is required to support real time applications that exhibit a more bursty nature (e.g. VoIP with silence suppression) so that LE, BE, and AF packets continue to be resized even when EF packets are not present.

R-6-50 When no packets are queued in the EF class for a duration longer that the EF window timer, the BE and AF packets MUST NOT be resized unless required for other reasons (negotiated MTU size).

R-6-51 The device MUST support PATH MTU negotiation when operating in an IP aware Mode.

R-6-52 The device MUST be able to maintain separate MTU sizes on different sessions.

R-6-53 The device must be able to enforce MTU size settings (packet lengths) for both IP aware and no IP aware sessions (i.e. layer 2 MTU). The device must be able to either discard or fragment packets that exceed MTU size defined for that session.

R-6-54 If multiple PVCs are provisioned per subscriber, the device MUST support the mapping between a Diffserv Code Point (DSCP) and a specific PVC

R-7-01 The device MUST be able to filter (silently discard) unsupported frames at the access interface.

R-7-02 The device MUST support IP layer rate limiting according to RFC 2697, A Single Rate Three Color Marker.

R-7-03 The device MUST support IP layer rate limiting according to RFC 2698, A Two Rate Three Color Marker.

#### 6.1.3.1 Hierarchical Scheduling and Policing

The BRAS will need to provide a congestion management function that will allow the synthesis of IP QoS through downstream elements that are not QoS aware. Accomplishing this is envisioned as a marriage of IP and ATM technologies with ATM and WFQ scheduling performed against diffserv and ATM queues. At a very high level, the queuing architecture desired for the BRAS can be described as IP DiffServ classification and queues mated to a slightly enhanced ATM scheduler. This results in emitting (shaping) ATM cells into the downstream network according to their VC contracts, ATM traffic engineering requirements, and so that no congestion occurs on the downstream links, systems, and topology. The result is that congestion queues in the BRAS, and eventual data discard occurs in packets being dropped from the DiffServ queues according to their precedence. Similarly in the upstream direction (from the ATU-r to the BRAS), the device must manage the access network using a hierarchical policing function to avoid congestion.

Multiple access sessions are supported in this model, however, all traffic is classified and scheduled in a monolithic system. Therefore, while it might appear at first that the Diffserv queuing and scheduling might apply only to IP-aware access – in fact all access, IP, Ethernet, PPP, and even ATM is managed by the same system that adheres to a combination of queuing disciplines taken from ATM and the Diffserv model. Note that the ATM disciplines are for backward compatibility, and don't otherwise interact with the Diffserv disciplines.

R-6-55 The device MUST support a Diffserv-aware hierarchical scheduler (per DSL Forum TR-059) that allows it to manage the network so that any potential congestion in the Access Network between the device and the RGs is avoided.

R-6-56 The hierarchical scheduler in the device MUST be able to model the congestion points in at least two subsequent ATM hops (corresponding to the daisy chaining of two ATM switching/multiplexing points in the Access Node); if the device does not include the ATM switching function, then the hierarchical scheduler in the device MUST be able to model the congestion point in yet a third, additional, ATM hop.

R-6-57 The device MUST support at least 5 layers of hierarchy (i.e. physical port, virtual path, VC group, VC scheduler, and session scheduler).

R-6-58 Hierarchical scheduling MUST be resource efficient in the sense that any traffic MUST be capable of using the unused bandwidth that has been allocated to other traffic classes.

R-6-59 The hierarchical scheduler MUST support allocating downstream bandwidth based on policy configuration across ATM, PPP, Ethernet, and IP technologies.

R-6-60 The Hierarchical scheduler in the device must support the modeling of bandwidth available to all involved downstream layer 2 device ports, including taking into account traffic that bypasses the device.

R-6-61 The Hierarchical scheduler in the device must be able to incorporate knowledge of individual ADSL synch rates and utilize this information when scheduling traffic for that particular VC.

R-6-62 The Hierarchical scheduler in the device must be able to shape and deliver downstream traffic based on the topological model such that downstream layer 2 ports are not congested.

R-7-04 The BRAS MUST be able to police upstream both for traffic aggregates and for sub-classes of the aggregate using the same topology information that exists for the hierarchical scheduler.

R-7-05 The BRAS SHOULD support random differential drop behavior for upstream traffic aggregates and sub-aggregates based on class.

## 6.2 MPLS Traffic Engineering

R-6-63 The device MUST support RFC 3209, RSVP-TE: Extensions to RSVP for LSP Tunnels.

R-6-64 The device MUST support RFC 3630, OSPF Traffic Engineering extensions.

R-6-65 The device SHOULD support IS-IS Traffic Engineering (e.g. draft-ietf-isis-traffic-04.txt).

# 7 Profile and Policy Management

R-7-06 The device MUST provide support for attaching policies to any interface, session, subscriber, or application flow.

R-7-07 The device MUST be capable of dynamically changing policing and shaping parameters at the ATM layer, IP layer, and PPP layers, based on information provided by an external policy server (i.e. policy decision point), without requiring session reinitialization

R-7-08 The device MUST provide a basic Policy Editor tool for defining policies. Such an editor should be able to export policies in a standard format for storage in external policy repositories (e.g. LDAP directories)

R-7-09 The device MUST provide a Policy Administration tool for associating policies with services, and indicating default policies for subscribers.

R-7-10 The device SHOULD provide a policy validation tool to identify policy conflicts.

R-7-11 The device MUST be able to access a repository for retrieving information such as user profiles, service definition parameters, RADIUS records and access, and service portal configurations.

R-7-12 The device MUST support a standard message format, i.e. data structure, based on PCIM for the retrieval of policy information such as user profiles, service definition parameters. There MUST be multiple bindings of this message format to transport protocols including RADIUS, COPS, SOAP/BEEP, NSIS, LDAP, SNMP.

R-7-13 The device MUST NOT require a session to be torn down and re-established for a new policy to be applied to or affect that session

R-7-14 The device MUST be able to receive a policy or policy indicator via a RADIUS response.

R-7-15 The device MUST support granular changes to policy rules, i.e. changing a single policy rule should only impact the subscribers within the scope of that policy rule.

R-7-16 The device MUST support installation of new policies at session establishment time.

R-7-17 The device MUST support the installation of new policies during an existing session.

R-7-18 The device MUST support Policy Accounting to capture packet level metrics and policy transactions.

R-7-19 The device MUST support and apply at least 100 policy transactions applications per second from an external directly attached uncongested server.

R-7-20 The device MUST support a minimum of 50 Policy Groups consisting of other Policy Groups and/or Policy Rules.

R-7-21 Policy rules are composed of policy conditions and policy actions. The device MUST support applying 10 policy rules per subscriber or a minimum of 10,000 policy rule applications.

R-7-22 The device MUST support forwarding based on the classification criteria in R-6-10 to MPLS LSP, VLAN, 802.1P, ATM VP/VC, or other traffic engineering capabilities in the Regional Network.

## 7.1 Policy Actions

R-7-23 The device MUST support the ability to police flows, which have been classified as per R-6-10.

R-7-24 The device MUST support the ability to filter flows, which have been classified as per R-6-10.

R-7-25 The device MUST support the policy route flows, which have been classified as per R-6-10.

R-7-26 The device MUST be able to associate a custom filter rule set with a given user profile, PVC or session.

R-7-27 The device MUST allow a NSP connection configured such that it can specify whether or not an end user is allowed to have simultaneous sessions with other NSPs while connected to it.

R-7-28 The device MUST be able to apply at least 4 ACLs, each containing multiple policies, per interface (any physical or logical interface) without affecting the device's performance.

# 8 Operations

R-8-01 The device MUST support hitless upgrades, i.e., the ability to upgrade operating system software without interruption of service.

R-8-02 The device MUST support rollback procedures, i.e., the ability to rollback operating system software to a previous version.

R-8-03 The device MUST support SNMPv1.

R-8-04 The device MUST support SNMPv2.

R-8-05 The device MUST support SNMPv3.

R-8-06 The device MUST support NTP.

R-8-07 The device MUST allow software images being downloaded from an EMS."

R-8-08 The device MUST support the configurable option to set management traffic to be physically or logically segregated from user traffic..

R-8-09 The device MUST continue to operate properly without interruption if an IP interface, PVC, port, tunnel, VPN, tunnel group, or any other provision-able interface type or logical grouping of sessions is deleted.

R-8-10 The management system MUST support the centralized provisioning of services that span multiple network elements (i.e. VPNs).

R-8-11 The management system MUST support the collection and reporting of service related statistics, including for example the configured VPNs and number of subscribers per VPN.

R-8-12 The device MUST support BGP v3 MIB per RFC 1269 or BGPv4 MIB (e.g. draft-ietf-idr-bgp4-mibv2-03.txt).

R-8-13 The device MUST track instances where the same MAC address is received from two or more PVCs. Both a counter and a syslog message is required. The syslog message MUST specify the PVCs involved and the MAC address seen.

R-8-14 The device SHOULD support Layer Two Tunneling Protocol 'L2TP' Management Information Base per RFC 3371.

R-8-15 If LDAP is supported, the device MUST support a fail over mechanism for redundant LDAP servers.

R-8-16 If COPS is supported, the device MUST support a fail over mechanism for redundant COPS servers.

R-8-17 The device MUST provide notification (e.g. alarms) to the EMS of changes in operational state and provide supporting information regarding the state and what caused the change.

R-8-18 The device MUST provide acknowledgement to requests by the EMS, e.g., take circuit pack out of service.

R-8-19 The device MUST provide an operations interface that is capable of being accessed remotely in the event the EMS is out of service.

R-8-20 The device MUST support backup images of software releases and configuration parameters that can be used as fallbacks.

## 8.1 EMS Interface Requirements

The device will be managed by an EMS. Except in occasional conditions when an EMS is down or when a device is initially being placed in service, all day-to-day operations functions for a device will typically be performed using an EMS. A detailed listing of all EMS requirements is outside the scope of this document. Some of the device requirements in support of an EMS are as follows:

R-8-21 The device SHOULD be fully manageable by an EMS that will allow technicians and northbound OS to remotely perform all FCAPS[1] functions. The device will provide standard northbound interfaces and APIs so that EMS. Can provide the following functions:

- EMS will provide full configuration, inventory, auto-discovery and provisioning capabilities via GUI menus.
- EMS will provide alarm management including collecting, thresholding, displaying, sorting, filtering, assigning and clearing of alarms.
- EMS will provide accounting and billing aggregation functions for CDRs produced by the device for collection by a northbound OS.
- EMS will collect, threshold, aggregate and display all capacity and performance measurements associated with the device.
- EMS will provide security administration for the device including authentication, access control and audit trail management capabilities.

R-8-22 The decision for setting which device, the NE or the EMS/NMS has the master copy MUST be set through the EMS/NMS

R-8-23 The capability MUST exist to synchronize the databases of the EMS and the device, either by applying the EMS database to the device, or by applying the device database to the EMS.

## 8.2 Provisioning

R-8-24 The device MUST be able to provision/limit the maximum number of sessions per tunnel.

R-8-25 The maximum number of sessions allowed from an end user across a single PVC MUST be configurable, at least in the range 1-8. This is a property of the PVC. In this context, “sessions” includes PPP sessions as well as DHCP assigned addresses.

R-8-26 Over the NMS-EMS interface, the operator MUST be able to specify an appropriate profile when provisioning a service.

R-8-27 The EMS northbound interface MUST support flow-through provisioning and alarm transmission. The CORBA protocol is desired. An API is preferred, to allow the operator’s NMS to specify functions including but not limited to:
- checking the availability of a port and VPI/VCI before provisioning,
- returning the result of the above check, and
- provisioning a PVC.

R-8-28 The device MUST allow activated line cards or other modules to go into service with pre-specified default configuration parameters.

---

[1] FCAPS (fault, configuration, accounting, performance and security)

R-8-29 The device MUST support the auto detection of access ATM traffic with VPI/VCI in pre-configured ranges, and automatically associates those VCs with a default “service profile”. Properties of the default service profile, such as ATM VC traffic parameters, PPPoE, PPPoA, or DHCP, are configurable.

## 8.3 Fault Management

R-8-30 End users, service providers, and the Operator MUST be able to use pings and traceroutes to aid in fault isolation.

R-8-31 The device’s MUST be able to send facility and equipment alarms from the device to EMS

R-8-32 The delay between a fault’s occurrence and when the corresponding alarm is reported to the EMS MUST be less than 5 seconds.

R-8-33 The device MUST support DS3/DS1 physical alarm reporting capabilities that are compliant with the applicable requirements of Telcordia GR-499.

R-8-34 The device MUST provide alarms when a RADIUS server is not responding.

R-8-35 The device MUST support alarms when any component, logical interface, or control service is not available.

R-8-36 The device MUST support a generic mechanism to monitor system resources (e.g. IP Address Pools) by creating and monitoring thresholds. The result of a threshold crossing is the generation of a Threshold Crossing Alarm (TCA) by the device

R-8-37 The device MUST provide sufficient identifying information in alarms to facilitate troubleshooting such as date/time stamps, severity, component identifiers and hardware, software and firmware versions.

R-8-38 The device’s severity attribute included in fault alarms MUST, at a minimum, categorize the alarm as Critical, Major, Minor and Informational.

R-8-39 The device MUST provide alarms to facilitate automatic clearing of alarms once trouble conditions have been corrected.

R-8-40 The device MUST have the capability to perform proactive maintenance monitoring such as audit routines, monitoring of memory usage, detection of "lost" processes or other software failures and take appropriate recovery actions.

R-8-41 The device MUST support self-diagnostic tests and provide diagnostic tools in case of failure. The component must provide fault recovery programs to detect and locate faulty hardware units and reconfigure the system with a minimum degradation of call processing.

R-8-42 The device MUST respond to requests from the EMS for routine diagnostics to detect hardware equipment failures and report results to the EMS.

R-8-43 The device MUST report relevant information from audits and diagnostic results in descriptive statements rather than in hex dumps or in some other format, which requires translation. Trouble codes may accompany the English descriptive messages.

R-8-44 The device MUST provide a dedicated physical interface (e.g. Ethernet) for access by technicians to perform FCAPS functions directly on the component in the event the EMS is down.

R-8-45 A recovery procedure MUST be provided in the event a new software load experiences a problem.

## 8.4 Configuration Management

Many Network Elements group sets of configuration parameters into “profiles”, where a specific profile includes specific values for each parameter of the profile. There may be several different types of profiles, where each type has a different set of parameters. For example, there could be a transmission profile, an ATM profile, or an IP filter profile.

The profile approach simplifies the management of configuration in some ways, which is desirable. On the other hand, if a limited number of profiles are available within each profile type, this can be limiting.

R-8-46 The capability MUST exist to pre-configure customers on the device. In particular, it MUST be possible to configure many customers with the same characteristics using profiles (i.e. support for bulk provisioning).

R-8-47 The device MUST support declaring individual PVCs administratively down without having to tear down and rebuild the PVC.

R-8-48 The device MUST send autonomous message(s) to the EMS informing it of any status changes of the device’s shelf or plug-ins.

R-8-49 The device MUST provide inventory of sub-components and configuration parameters.

R-8-50 The device MUST provide error notifications when operator entries are out-of-range or invalid.

R-8-51 The device MUST provide verification capabilities to allow operators to easily determine existing hardware and software configurations.

R-8-52 The device MUST provide a backup map or copy of all configurations and connection data in non-volatile storage for use in rapid restoration of the components or portions of the component, such as a port card, upon restart or faulty card replacement.

R-8-53 The device MUST support the ability to establish configuration parameters in ranges.

R-8-54 The device MUST support the ability of an EMS to perform bulk configuration of logical or physical entities via scripts. This requires the device to be able to provide acknowledgement messages to the EMS that a specific configuration request has been completed.

R-8-55 The device MUST support the ability to have software releases and patches be remotely installed.

R-8-56 The device MUST provide the ability to back out software generics or patches and restore the system to the previous generic or patch based on a request from the EMS.

R-8-57 Software updates MUST not negatively impact the continued, live operation of the components or affect the component logical assignments.

R-8-58 Software updates MUST be in modular format to allow for the addition of new features.

R-8-59 The device MUST support virtual partitioning capabilities so specific operators can only view and update specific logical or physical portions of the device.

## 8.5 Performance Monitoring

The remote network operators need to have the ability to monitor traffic on a client basis, PPPoE or PPPoA session basis, PVC basis, L2TP tunnel basis, and NSP basis. Basic information such as traffic rates, error rates, cell/packet drop rates is required.

R-8-60 The device MUST provide packet monitoring of received and transmitted packets on ATM PVCs, L2TP tunnels, PPPoE sessions, and PPPoA sessions. Measures of error packets are desirable, but per-user measures of “good” packets are essential to service activation procedures.

R-8-61 The device MUST provide SNMP MIB defined capacity management measurements that are accessible through the EMS. The measurements should at least include the following:

- Measure the number of simultaneous sessions per port, card, and system in 24-hour period.
- Measure the number of sessions rejected due to tunnel maximum reached per tunnel in a 15-minute period including a 5-minute peak measurement.
- Measure the interface utilization as a percentage of maximum octets in a 15-minute period including a 5-minute peak measurement.
- Measure discarded cells/packets due to congestion in a 15-minute period

R-8-62 The device MUST support DS3/DS1 physical layer performance monitoring (near end and far end) data collection that is compliant with the applicable requirements of Telcordia GR-499.

R-8-63 The device MUST support a reporting mechanism to identify the IP address or addresses being used by each subscriber.

R-8-64 The device MUST provide statistics for various IP packet sizes, for all types of supported interfaces.

R-8-65 The device MUST provide PPS performance measurement for all types of supported interfaces.

R-8-66 IP packets per second thresholds MUST be configurable on an interface basis and generate notifications if the thresholds are exceeded

R-8-67 The device MUST be able to log to a syslog server a message when a subscriber exceeds the maximum access session limit.

R-8-68 The device MUST be able to log to a syslog server a message when a subscriber fails to authenticate their PPPoE or PPPoA session with summary details of the authentication failure (invalid password, invalid domain, invalid username, etc…).

R-8-69 The device MUST support a command that would summarize the number of PVCs defined on the device. The command should provide a second option to the command to show the breakdown per encapsulation type.

R-8-70 There MUST be collection points available for ingress and egress data, as well as internal CPU utilization of the device, so statistics can be analyzed for traffic engineering purposes. Enabling these statistics does not impede the amount of PVCs, circuits, or any type of session that may be provisioned on the device.

R-8-71 The device MUST count and measure all data associated with determining capacity utilization and Quality of Service (QoS) metrics.

R-8-72 The device SHOULD retain traffic/performance data for a period of time as defined in the EMS.

R-8-73 The device MUST support configurable thresholds for IP address pool usage high and low water marks, and report threshold crossings to the management systems.

R-8-74 The device MUST support increasing of the IP address pool size, and decreasing of the IP address pool size in order to reclaim unused resources. In the case of tools implemented in external systems, the device MUST be able to process pool resize messages from those systems. When the architecture includes egress routers between the device and NSP, resized IP address pools MUST be reflected in the egress routers routing configuration.

R-8-75 The device MUST provide traffic mirroring feature at the subscriber session level, i.e. the ability to copy upstream and/or downstream packets to a specified destination (locally or preferably to a remote collection point) for at least 2% of the provisioned subscribers on the device.

## 8.6 Trouble Resolution

R-8-76 The device MUST allow a lookup of all sessions within a VC, VP, VR, per user login, per IP (if applicable), per MAC.

R-8-77 The device MUST support turning on debugging tools for an individual access side VC.

R-8-78 The device MUST support viewing L2TP status and statistics on the LAC side of L2TP tunnels.

R-8-79 The device MUST support a “show ip route” like command.

R-8-80 The device MUST support a “show ARP table” like command.

R-8-81 The device MUST support a command to display the PPP session information and the L2TP tunnel identifier the session is using. I.e. cross-reference the PPP session with an L2TP tunnel. One should be able to easily locate and identify individual end user PPP sessions within tunnels regardless of the number of tunnels present (e.g., multi tunnel environment).

# 9 Security

The requirements in this section are specific to capabilities needed for the protection of the device itself from external attacks.

R-9-01 The device MUST support the security requirements included in T1M1 T1.276.2003 requirements document. Where there are conflicting security requirements in this document, the more stringent requirement should take precedence.

R-9-02 The device MUST support a mechanism (e.g. configurable rate limit) for traffic accessing the control plane of the device to protect against (flooding) and other Denial of Service (DoS) attacks from external sources.

R-9-03 The device MUST support a mechanism (e.g. configurable rate limit) for subscriber traffic based on packet type (e.g. TCP SYN, ICMP etc.) to protect against distributed/flooding denial of service attacks from external sources.

R-9-04 The device MUST support a mechanism (e.g. configurable rate limit) for flows to/from users based on packet type (e.g. TCP SYN, ICMP etc.) to protect against flooding Denial of Service attacks from hacked/malicious subscriber hosts.

R-9-05 The device MUST be able to identify hacked/malicious subscriber hosts by using source address validating/filtering.

R-9-06 The device MUST support MAC layer ACLs to permit or deny based on source/destination MAC address or the classification parameters specified in R-6-10.

R-9-07 The device MUST be able to filter/block/log traffic based on any field or combination of fields of a packet header (layer 2 and above)

R-9-08 The device MUST support IP spoofing detection and blocking.

R-9-09 The device MUST support IP source route option detection and blocking.

R-9-10 The device SHOULD support port scan detection and blocking.

R-9-11 The device MUST support IP address sweep attack detection and blocking.

# Appendix A – Multicast Support

Figure 4 shows a TR-59 compliant network deployment model.

- The model of deployment is to have B-RAS functional to perform all of multicast operations in data plane and control plane.
- The multicast streams are delivered to B-RAS via PIM-SM protocol.
- IGMP is used on subscriber facing interfaces.
- An access control-list based mechanism allows or disallows for sets of multicast channels on a per subscriber basis. For example, a RADIUS attribute that specifies an access control list is applied to the subscriber's interface to filter out all IGMP messages except for allowed multicast channels.
- Support for source specific multicast as described below:

**Figure 1 TR-59 compliant Multicast operational model**

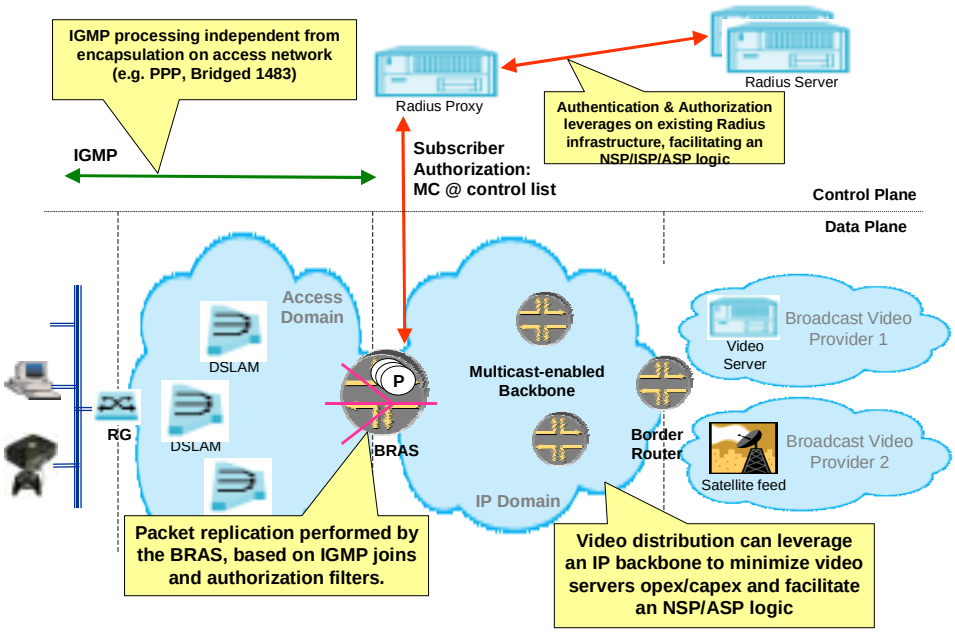


**Figure A-1 - Multicast Model**

Source-specific multicast (SSM) is a service model that identifies session traffic by both source and group address. SSM is ideal for one-to-many multicast services such as network entertainment channels. SSM builds shortest-path trees (SPTs) directly represented by (S,G) pairs. The "S" refers to the source's unicast IP address, and the "G" refers to the specific multicast group address. The SSM (S,G) pairs are called channels to differentiate them from any-source multicast (ASM) groups. While ASM supports both one-to-many and many-to-many communications, ASM's complexity is in its method of source discovery. For example, if you click on a link in a browser, the receiver is notified about the group information, but not the source information. With SSM, the client receives both source and group information.

To deploy SSM successfully, you need an end-to-end multicast-enabled network and applications that use an IGMPv3 stack.

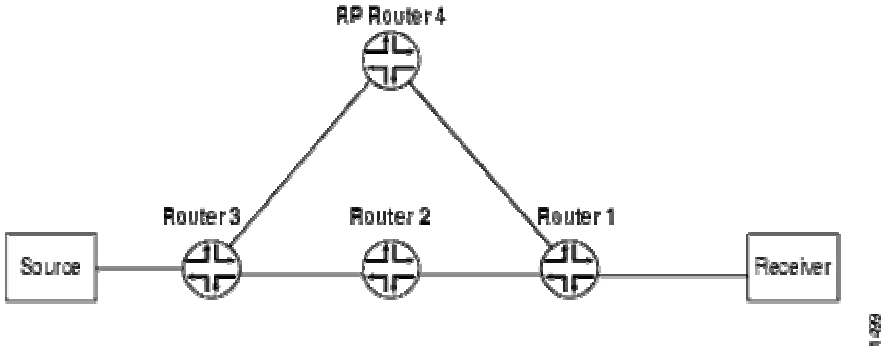


The picture above shows a PIM-SM network with both ASM and SSM. Router 4 serves as an RP only for ASM multicast groups. SSM channels do not need RP. Router 1 is the B-RAS, while the Receiver is one of the subscribers wanting to receive the multicast content. IGMPv3 is enabled on the Router 1 interface facing Receiver. All the routers in the distribution network connection to Source run PIM-SSM.

**Requirements**

| | |
|---|---|
| R-10-01 | The device MUST be configurable with the option to silently discard all subscriber upstream multicast traffic. |
| R-10-02 | The device MUST implement Host Extensions for IP Multicasting defined in RFC 1112. |
| R-10-03 | The device MUST implement Internet Group Management Protocol, Version 2 (IGMP v2) defined in RFC 2236. |
| R-10-04 | The device MUST implement Internet Group Management Protocol, Version 3 (IGMP v3) defined in RFC 3376. |
| R-10-05 | The device MUST Support RFC 2362, Protocol Independent Multicast-Sparse Mode (PIM-SM). |
| R-10-06 | The device SHOULD support Protocol Independent Multicast - Dense Mode (PIM-DM) (e.g. draft-ietf-pim-dm-new-v2-03.txt). |
| R-10-07 | The device MUST support Rendezvous Point auto-discovery in PIM-SM. |
| R-10-08 | The device MUST support, RFC 2365, Administratively Scoped IP Multicast |
| R-10-09 | The device SHOULD support Multicast Source Discovery Protocol (MSDP). |
| R-10-10 | The device MUST support Any-cast RP mechanism using PIM and MSDP. |
| R-10-11 | The device MUST support Source-Specific Multicast for IP. |
| R-10-12 | L3 features supported on the device MUST also be available for multicast, e.g. rate limiting and filtering. |
| R-10-13 | The device SHOULD support tunneling of PIM-SM. |
| R-10-14 | When operating in an IP-routed mode the device MUST provide multicast access controls including authentication and collect multicast usage information. |
| R-10-15 | The device SHOULD support a user authentication protocol for joining multicast streams as these standards mature i.e. IGAP: IGMP for User Authentication Protocol", IETF Internet Draft, draft-hayashi-igap-00.txt, |
| R-10-16 | The device MUST provide multicast functions without performance degradation. |
| R-10-17 | The device must support at least 1000 multicast groups. |

R-10-18 The device MUST support a mechanism to allow and disallow multicast groups via access control list on per subscriber basis, such as using RADIUS returned attributes.

R-10-19 The device MUST process IGMP join and leave messages within 100 milliseconds.

R-10-20 The device MUST support QoS features for multicast traffic without performance degradation.

R-10-21 The device SHOULD be able to implement Multicast VPNs as described by RFC 2547bis

R-10-22 The device MUST support multi-protocol BGP to enable BGP distribution of multicast routing information both within an AS and between ASs..

# Appendix B - IP VPN Services

While TR-059 describes access methods and ISP/ASP services, it does not explicitly detail how a DSL network could be used to offer VPN services. The proposed architecture in this document would facilitate the delivery of service provider initiated tunneled and terminated VPN services. These services could include layer 3 IP interfaces between the BRAS and NSP/ASP and support IP VPN service. With network based VPN services, geographically separated sites can establish secure communication channels through an IP network. This document describes a model for delivery of these services and requirements to support them. The requirements listed below are in addition to requirements already identified in the main body of this document.

Two types of VPNs are described in this document - terminated/MPLS-based and tunneled/L2TP-based. Both are in reference to what takes place at the BRAS. With tunneled VPNs, the traffic is encapsulated and transported over a virtual dedicated connection between communicating end devices. In this way, user traffic is tunneled from a site connected to a BRAS to a secondary site (which may be connected to a BRAS or some other PE device). With terminated VPNs, the PPP session from a user is terminated at the BRAS and the IP packets are encapsulated in a MPLS VPN to be forwarded to the BRAS or PE *connecting* other site.

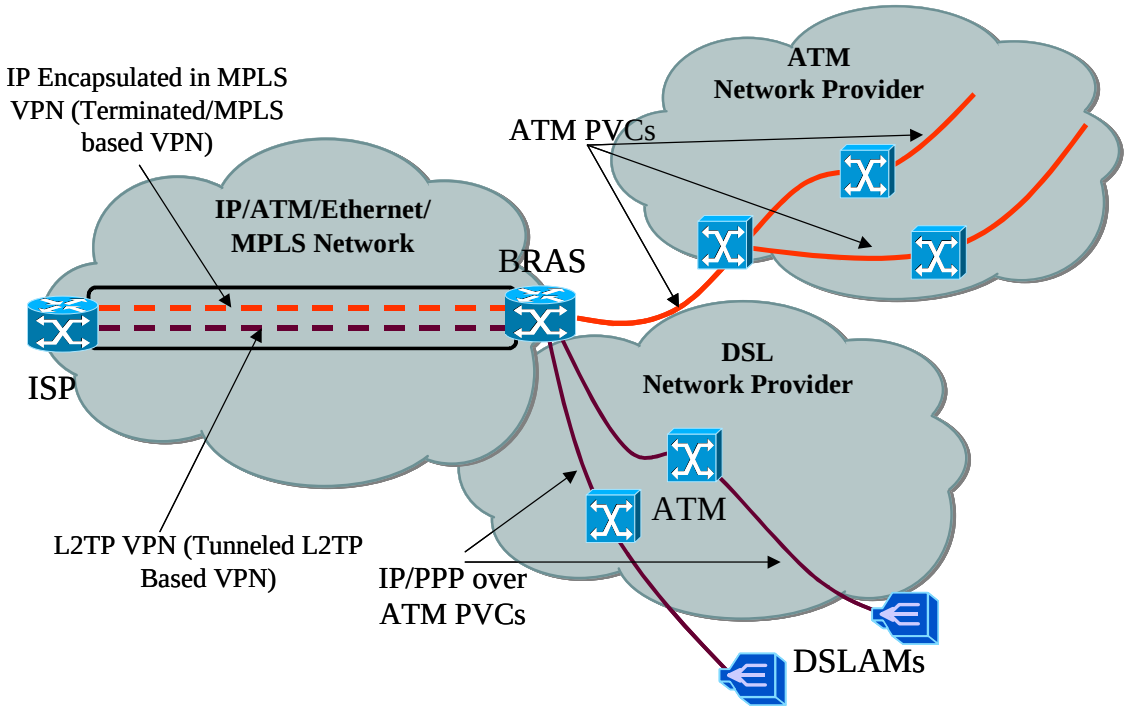


**Figure B-1 - DSL Virtual Private Network**

**Virtual Private Network Service model:**

IP VPN is a Layer 3 VPN offering that provides customers with multi-site, extended-area IP services to securely connect geographically dispersed sites over the service provider's network. A VPN provides a means for routing customer IP traffic across a service provider's network. A customer may retain current IP addressing whereby VPN addressing is utilized only between VPN gateway devices, such as the BRAS. An IP VPN also hides the complexity of WAN protocols from the end-user, by connecting LANs in such a way that the end-to-end WAN services supporting the network (e.g. ATM or Ethernet) are not apparent to the customer. The IP VPN service can support any type of traffic over the IP Layer 3 protocol

that the subscriber chooses to use. In addition, the IP VPN can be served over many WAN backbones (i.e. Frame Relay, ATM, Switched Ethernet, IP/MPLS). IP VPNs may be between two points or among a number of end points. As a network service offering, IP VPNs can be added to a transparent LAN services, over the same connection to the customer's LANs.

As is shown in Figure B-1 - DSL Virtual Private Network, the IP VPN starts at the edge of the service provider's network (PE) at the BRAS and flows across the network end to end at the IP Layer (3). At the BRAS, the subscriber's connection can be virtually connected/routed to other virtual connections within the same VPN. The BRAS has capabilities to maintain the Layer 3 connection for the VPN and determine where traffic is forwarded. Each IP VPN may be identified by its ID allowing many virtual VPNs to operate on the BRAS. In a typical DSL application, the subscriber access is via ATM PVC that connects the subscriber CPE to the BRAS. Over the ATM PVC from the customer's end device, the IP VPN is carried over the PPPoE connection to the BRAS.

Figure B-2 shows the simplified[2] protocol stacks and layout of a MPLS-based VPN. This method is used when an MPLS-enabled core is available. PPP traffic from the originating user is terminated on the BRAS and the retrieved IP packets are tagged with MPLS VPN labels for forwarding into the network core. The user may choose to use IPSec to encrypt the actual data; this is transparent to the VPN as the IPSec would be tunneled inside the IP layer implicitly in the figure. Using MPLS terms, the BRAS is a Provider Edge (PE) router and the inner core router is a Provider router (P).

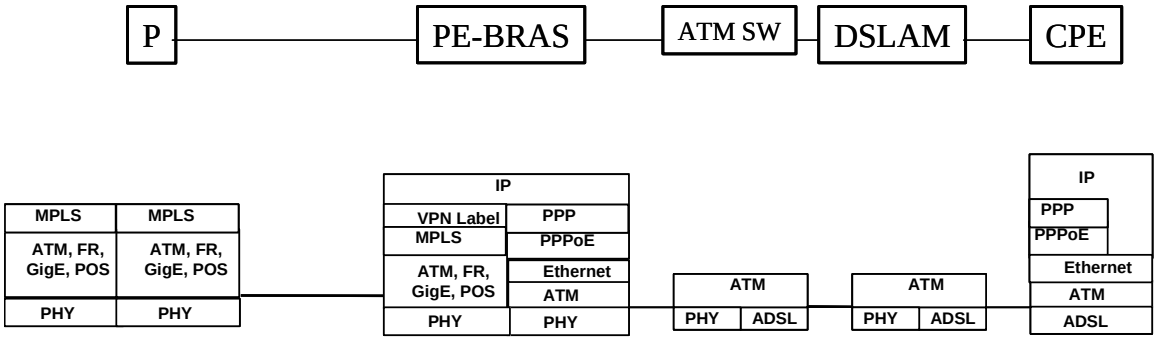


**Figure B-2 - MPLS-based VPN**

Figure B-3 shows an alternative VPN option where the BRAS tunnels the PPP traffic to the endpoint using L2TP, instead of recovering the IP packets. At the BRAS, the MPLS capability is not available nor is the user's PPP traffic terminated at the BRAS. At the LNS, the L2TP tunnel can then be mapped to an MPLS LSP or carried over other types of transport into the core.

[2] In both figures the full protocol stacks are not shown, only the portions relevant to the discussion are drawn.

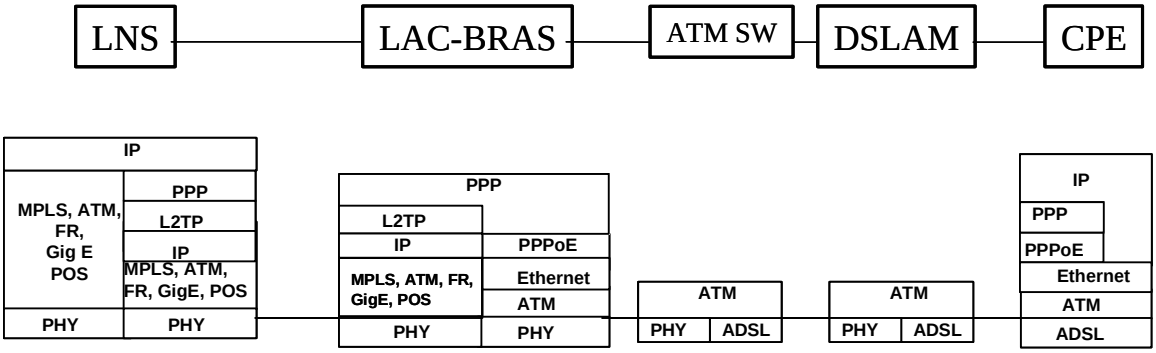


**Figure B-3 - L2TP-based VPN**

Figure B-4 and Figure B-5 show PPPoA and RFC 2684 based VPS, respectively.

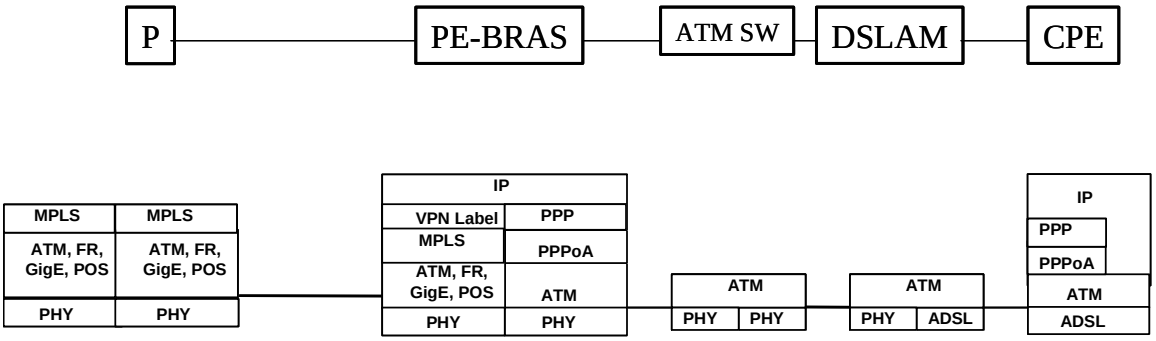


**Figure B-4 - PPPoA Based VPN**

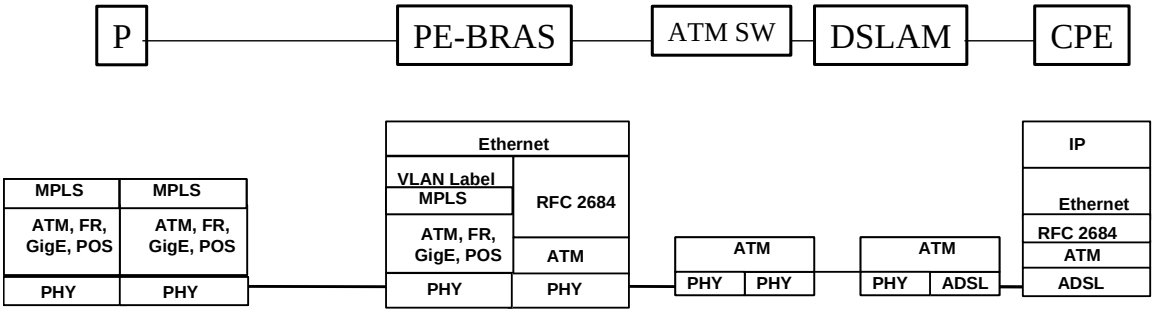


**Figure B-5 - RFC 2684 Based VPN**

**Requirements:**

The following requirements support VPN capability on the BRAS

R-11-01 The device MUST provide a level of privacy and security for VPNs that is equivalent to that obtainable from a pure Layer 2 infrastructure, such as Frame Relay or ATM. This includes IP routing information and routing protocol separation.

R-11-02 The device MUST be able run its own routing protocol for VPNs.

R-11-03 The device MUST be capable of IP layer separation for VPNs (IP routing information and routing protocol separation) such that there is no IP layer connectivity to other VPNs, or to Service Provider's network internals.

R-11-04 The device MUST be configurable to support limited and controlled exchange of IP layer traffic between different VPNs (extranets).

R-11-05 The device MUST be configurable to support customers who require combined VPN and Internet access services.

R-11-06 The device SHOULD support RFC 3022, Traditional IP Network Address Translator (Traditional NAT). This includes IP address and TCP/UDP port translation.

R-11-07 When Network Based IP VPNs are provided using 2547, the device MUST support the following protocols on the PE-CE link:

- BGP
- OSPF
- RIP
- Static routes

R-11-08 When Network Based IP VPNs are provided using 2547, the device SHOULD support ISIS on the PE-CE link.

R-11-09 The device MUST support BGP4/MPLS VPNs per Internet RFC 2547 bis.

R-11-10 The device MUST be able to support VPNs with overlapping address spaces.

R-11-11 The device MUST maintain a separate forwarding information per VPN.

R-11-12 The device MUST allow the viewing of VPN specific route tables.

R-11-13 The device MUST allow routes that are learned on a particular interface that is associated with a VPN to be placed into the forwarding table supporting that VPN.

R-11-14 The device MUST allow a default route to be configured on a per VPN basis.

R-11-15 The device MUST allow the association of a packet to a VPN to be based on the ATM PVC on which it arrived.

R-11-16 The device MUST allow the association of a packet to a VPN to be based on the physical interface on which it arrived.

R-11-17 The device MUST function simultaneously as a Service Provider Edge (PE) router and a Service Provider Backbone (P) Router as defined in RFC 2547.

R-11-18 The device MUST support customer routers that are dual attached to the IP-VPN network for reliability reasons. (Site of origin? Is so delete)

R-11-19 The device MUST use a single (e.g., global) BGP instance to peer with other Autonomous Systems (AS) and then filters the route information on a per VPN basis.

R-11-20 The device MUST allow each interface used to peer with an external network (different AS) to be configured as part of a specific VPN.

R-11-21 The device MUST support per VPN DHCP

R-11-22 The device MUST support RADIUS clients per VPN

R-11-23 The device MUST allow the association of a packet to a VPN to be based on the VLAN tag associated with the Ethernet frame carrying the IP.

R-11-24 The device MUST support limiting the number of routes learned from a CE.

R-11-25 The device MUST support warning/notification if number of routes learned from a CE exceeds the maximum limit.

R-11-26 The device MUST support per-VRF Loop back interfaces.

R-11-27 The device MUST support the propagation of vendor proprietary extended BGP communities (starting from x8000), associated with routes learned using EBGP session with Customer Edge (CE) device, to its IBGP peers.

R-11-28 The device MUST not drop a route learned from an IBGP peer which contains vendor proprietary extended BGP communities and propagate routes with or without (should be configurable) vendor proprietary extended communities to CE device over EBGP session.

# Appendix C - Transparent Virtual LAN Services

Transparent LAN Service (TLS) is a specialized Layer 2 VPN (Virtual Private Network) offering that provides customers with multi-site, extended-area LAN services to connect geographically dispersed sites over the service provider's network. Transparent LAN Services give the appearance that multiple sites are all connected to the same LAN segment. A TLS provides a means for transparently bridging and transporting customer Ethernet data across a service provider's network. TLS allows a customer to retain current IP addressing practices. TLS also hides the complexity of WAN protocols from the end-user, by connecting LANs in such a way that the WAN services supporting the network (e.g. ATM or Ethernet) are not apparent to the customer. TLS/VPLS can be served over many WAN backbones (i.e. ATM, Switched Ethernet, IP/MPLS). These virtual LANs may be point to point, point to multi-point or multi-point to multi-point.

As shown in figures 1 and 2 below, the BRAS acts as the layer 2 bridge device for the Transparent LAN. The subscriber's virtual bridge port is located there as well as additional multiple virtual connections which may interface into the same bridging instance across multiple BRASs or Customer sites. This Bridging instance is responsible for the traditional Layer 2 MAC based learning bridge function to determine where traffic should be forwarded. Each virtual bridge may be identified by a VLAN id allowing many virtual bridges to operate on the BRAS. Subscriber access is via ATM encapsulated Ethernet frames which connects subscriber CPE to the BRAS. These frames are then link layer bridged and layer 2 forwarded to the appropriate port for delivery over the regional/access network. Frames are forwarded across and within the TLS based upon the link layer addresses (i.e. MAC addresses) associated with the individual hosts. The regional Broadband network can use any transport protocols (i.e. ATM, Switched Ethernet, IP/MPLS etc.) to support this forwarding provided the BRAS can encapsulate the Ethernet traffic accordingly. Scalability of this service model is of prime concern and as such limitations should be made on the amount of MAC addresses that can be learned per bridging instance as well as the amount of interconnected sites that can be supported per VLAN.

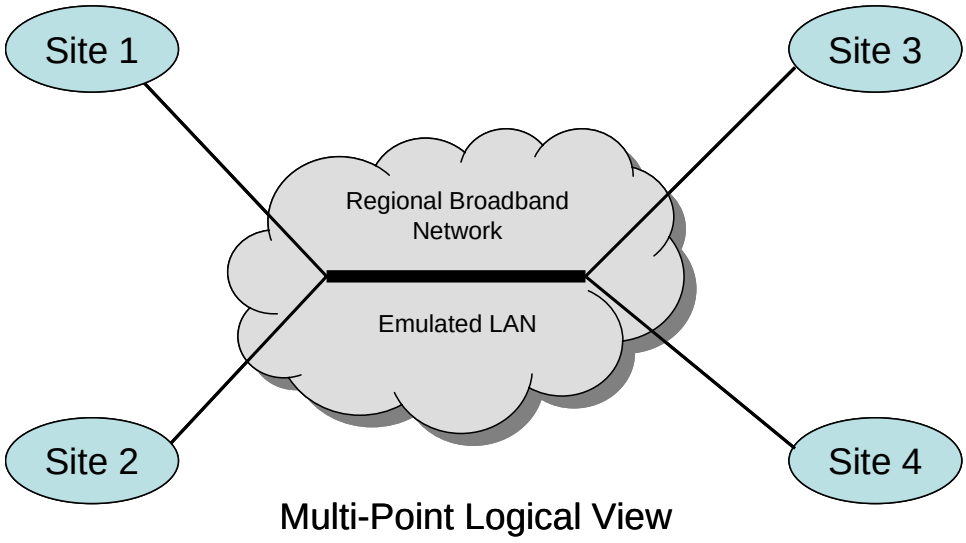


**Figure C-1 - Multi-Point Logical View**

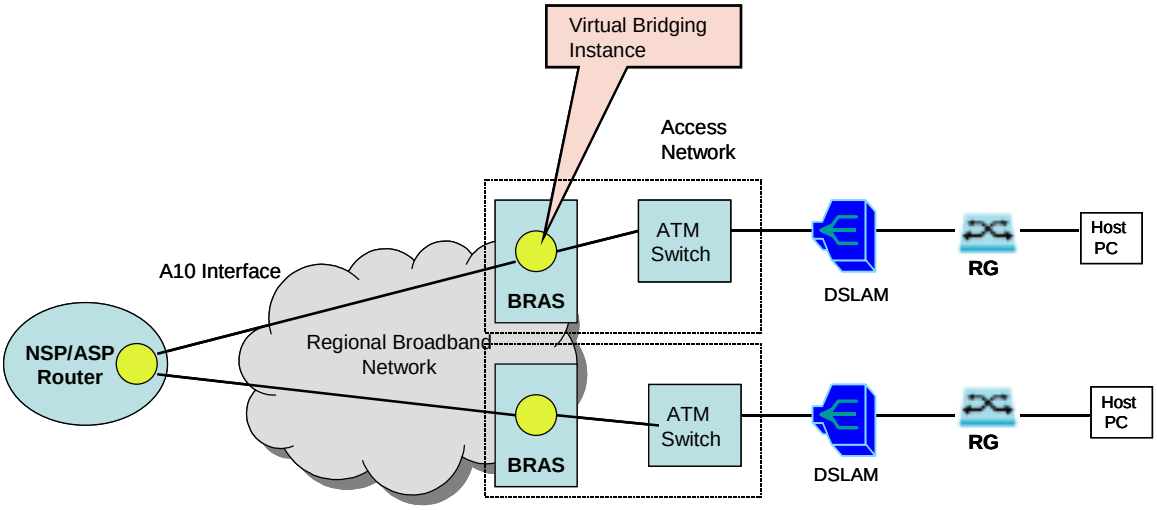


**Figure C-2 - Multiple Point to Point**

| IP |
|---|
| 802.3 |
| (MPLS) option |
| ATM, Ethernet, etc. |
| PHY |

**Figure C-3 - Protocol Stack for connection to and from NSPs and ASPs**

**Requirements:**

The following is minimum list of requirements in support of TLS and VPLS capability on the BRAS and are in addition to requirements already identified in the main body of this document in section 4.5:

R-12-01 The device MUST support multi-port Ethernet bridge functionality (802.1D)

R-12-02 The device MUST support tunneling of Ethernet control frames (e.g., Spanning Tree Protocol).

R-12-03 The device MUST support Jumbo Ethernet frames on physical Ethernet Interfaces.

R-12-04 The device MUST support VLAN stacking (the ability to put customer 802.1q VLANs inside a provider's VLAN). The IEEE standard MUST be implemented when finalized in IEEE 802.1AD working group.

R-12-05 The device MUST support a configurable use of the ethertype field when performing VLAN stacking until such time that the IEEE standard is finalized.

R-12-06 The device MUST comply with the 802.3ad Gigabit Ethernet Link Aggregation Trunking Standard. The device must support this functionality on physical ports.

R-12-07 The device SHOULD support VPLS (Virtual Private LAN Service) multipoint Ethernet frame transport over MPLS when standardized (RFC status).

R-12-08 The device SHOULD be able to map VLANs IDs to VPLS bridge groups.

R-12-09 The device SHOULD be able to rewrite IEEE 802.1q VLAN ID when connecting multiple TLS islands using Layer 2 VPLS.

R-12-10 The device MUST maintain a static VLAN to PVC mapping and a dynamically learned ATM PVC to MAC address mapping.

R-12-11 The device MUST allow the timeout interval on dynamically learned MAC address to ATM PVC mappings to be adjusted.

R-12-12 The device MUST support the mapping of multiple MAC addresses to a single ATM PVC.

R-12-13 The device MUST support bridge groups.

R-12-14 The device MUST support isolation between services.

R-12-15 The device MUST have ability to block all direct communication between two users in the same bridge group (either based on L2 or L3 mechanisms).

R-12-16 The device MUST support MAC layer ACLs to permit or deny based on Ethertype if the implementation allows customer traffic to be bridged between different subscribers.

R-12-17 The device MUST allow the Operator to clear a single user from the bridge table without clearing the entire table.

R-12-18 The device MUST allow lookup of all sessions within a VLAN by: user Login, IP Address (if applicable), and MAC Address

# Appendix D – RADIUS Attributes

## Introduction

RADIUS protocol is defined in RFC 2865 [1] and RFC 2866[2] with extensions defined in RFC 2869 [5]. Attributes for support of tunneling protocol are defined in RFC 2867[3] and RFC 2868[4].

The purpose of this appendix is to tabulate standardized RADIUS attributes that are defined in RFCs above that should be supported by the BRAS.

## Standard Attributes

Following table lists RFC 2865, 2866, 2867, 2688 and 2869 standard attributes that are used in BRAS environment.

| Attribute Name | Description | RFC | Attribute Number |
|---|---|---|---|
| User-Name | This attribute indicates the name of the user to be authenticated. | 2865 | 1 |
| User-Password | This attribute indicates the password of the user to be authenticated, or the user's input following an Access-Challenge. | 2865 | 2 |
| CHAP-Password | This attribute indicates the response value provided by a PPP Challenge-Handshake Authentication Protocol (CHAP) user in response to the challenge. | 2865 | 3 |
| NAS-IP-Address | This attribute indicates the identifying IP Address of the NAS which is requesting authentication of the user, and SHOULD be unique to the NAS within the scope of the RADIUS server. | 2865 | 4 |
| NAS-Port | This attribute indicates the physical port number of the NAS which is authenticating the user. | 2865 | 5 |
| Service-Type | This attribute indicates the type of service the user has requested, or the type of service to be provided | 2865 | 6 |
| Framed-Protocol | This attribute indicates the framing to be used for framed access. | 2865 | 7 |
| Framed-IP-Address | This attribute indicates the address to be configured for the user. | 2865 | 8 |
| Framed-IP-Netmask | This attribute indicates the IP netmask to be configured for the user when the user is a router to a network. | 2865 | 9 |
| Filter-Id | This attribute indicates the name of the filter list for this user. | 2865 | 11 |
| Framed-MTU | This attribute indicates the Maximum Transmission Unit to be configured for the user, when it is not negotiated by some other means (such as PPP). | 2865 | 12 |
| Framed-Compression | This attribute indicates a compression protocol used for the link. | 2865 | 13 |

| | | | |
|---|---|---|---|
| Framed-Route | This attribute provides routing information to be configured for the user on the NAS. | 2865 | 22 |
| Class | This attribute is available to be sent by the server to the client in an Access-Accept and SHOULD be sent unmodified by the client to the accounting server as part of the Accounting-Request packet if accounting is supported | 2865 | 25 |
| Session-Timeout | This attribute sets the maximum number of seconds of service to be provided to the user before termination of the session. | 2865 | 27 |
| Idle-Timeout | This attribute sets the maximum number of consecutive seconds of idle connection allowed to the user before termination of the session. | 2865 | 28 |
| NAS-Identifier | This attribute contains a string identifying the NAS originating the Access-Request. | 2865 | 32 |
| Acct-Status-Type | This attribute indicates whether this Accounting-Request marks the beginning of the user service (Start) or the end (Stop). | 2866 | 40 |
| Acct-Input-Octets | This attribute indicates how many octets have been received from the port over the course of this service being provided. | 2866 | 42 |
| Acct-Output-Octets | This attribute indicates how many octets have been sent to the port in the course of delivering this service. | 2866 | 43 |
| Acct-Session-Id | This attribute is a unique Accounting ID to make it easy to match start and stop records in a log file. | 2866 | 44 |
| Acct-Session-Time | This attribute indicates how many seconds the user has received service for. | 2866 | 46 |
| Acct-Input-Packets | This attribute indicates how many packets have been received from the port over the course of this service being provided to a Framed User. | 2866 | 47 |
| Acct-Output-Packets | This attribute indicates how many packets have been sent to the port in the course of delivering this service to a Framed User. | 2866 | 48 |
| Acct-Terminate-Cause | This attribute indicates how the session was terminated. | 2866 | 49 |
| CHAP-Challenge | This attribute contains the CHAP Challenge sent by the NAS to a PPP Challenge-Handshake Authentication Protocol (CHAP) user. | 2865 | 60 |
| NAS-Port-Type | This attribute indicates the type of the physical port of the NAS which is authenticating the user. | 2865 | 61 |
| Tunnel-Type | This attribute indicates the tunnel protocol(s) used. The attribute must be set to 3 – L2TP as this is the only supported method specified by DSLF. | 2868 | 64 |
| Tunnel-Medium-Type | This attribute indicates the transport medium type to be used to create a tunnel. The attribute must beset to 1 – IP as this is the only supported method | 2868 | 65 |

| | | | |
|---|---|---|---|
| | specified by DSLF. | | |
| Tunnel-Client-Endpoint | This attribute contains the address of the initiator end of the tunnel. | 2868 | 66 |
| Tunnel-Server-Endpoint | This attribute indicates the address of the server end of the tunnel. | 2868 | 67 |
| Acct-Tunnel-Connection-ID | This attribute indicates the identifier assigned to the tunnel session. | 2867 | 68 |
| Tunnel-Password | This attribute may contain a password to be used to authenticate to a remote server. | 2868 | 69 |
| Tunnel-Assignment-ID | This attribute is used to indicate to the tunnel initiator the particular tunnel to which a session is to be assigned. | 2868 | 82 |
| Tunnel-Preference | This attribute indicates the relative preference assigned to each tunnel. | 2868 | 83 |
| NAS-Port-Id | This attribute contains the text string which identifies the port of NAS which is authenticating the user. | 2869 | 87 |
| Framed-Pool | This attribute indicates the name of an assigned address pool. | 2869 | 88 |
| Tunnel-Client-Auth-ID | This attribute specifies the name used by the tunnel initiator during the authentication phase of tunnel establishment. | 2868 | 90 |
| Tunnel-Server-Auth-ID | This attribute specifies the name used by the tunnel terminator during the authentication phase of tunnel establishment. | 2868 | 91 |

## VSA Attribute List

Attributes are partitioned into two categories; per User and per Access VC. The attribute Session-Limit is applicable to per tunnel or per domain basis depending on LAC or PTA mode.

The per user attributes are used for specifying parameters for users as well as for domains. When user-name plus FQDN are specified, then the attribute defines the setting for a particular user. When only FQDN is specified, then the attribute defines the default setting for a domain.

| Per User Attributes | |
|---|---|
| Attribute Name | Description |
| IP-Pool-Name | This attribute specifies the name of an IP address pool. When configured, sessions associated with the user will always have IP addresses assigned from the named pool.<br>Definition of the IP address pools can be done through separate VSA and is usually done on a per domain basis. |
| Local-Forwarding | This attribute defines a flag to enable/disable subscriber-to-subscriber forwarding. When enabled, traffic between subscribers within the same domain is switched locally.<br>The feature is ideal for peer-to-peer applications such as gaming, IP telephony etc. |
| QoS-Policy-Name | This attribute specifies the name of a QoS policy. Possible uses include specifying tiered services such as Gold/Sliver/Bronze. |
| DNS-Servers | This attribute specifies the IP addresses of primary and secondary DNS to be used at IPCP negotiation. |
| WINS-Servers | This attribute specifies the primary and secondary WINS to be used. Specification can be done either by name or by IP address. |
| Service-Profile | This attribute identifies the pre-configured service profile to be used. |
| VPN-Association | This attribute identifies a VPN to which the user belongs. |
| Per Access-VC Attributes | |
| Attribute Name | Description |
| Allowed-Access-List | This attribute specifies a list of names. Only names on the list are allowed to set up session using the particular access VC.<br>Names can be the name of subscribers or the name of domains. |
| Session-Limit | This attribute specifies the maximum number of sessions allowed on the particular access VC. |

| Per Tunnel or per domain | |
|---|---|
| Attribute Name | Description |
| Session-Limit | This attribute specifies the limit for number of sessions.<br>In LAC mode, this attribute specifies the limit for number of sessions per L2TP tunnel.<br>In PTA mode, this attribute specifies the limit for number of sessions per domain. |

# Appendix D References

[1] RFC 2865 - Remote Authentication Dial In User Service (RADIUS). (IETF, June 2000)
[2] RFC 2866 – RADIUS Accounting (IETF, June 2000)
[3] RFC 2867 – RADIUS Accounting Modifications for Tunnel Protocol Support (IETF, June 2000)
[4] RFC 2868 – RADIUS Attributes for Tunnel Protocol Support (IETF, June 2000)
[5] RFC 2869 – RADIUS Extensions. (IETF, June 2000)