White Paper

Understanding DSLAM and BRAS Access Devices

This whitepaper discusses the functional and architectural evolution of Next Generation Broadband Access Devices required to deliver triple play services.

Agilent Technologies

Introduction

DSL is the prevailing broadband Internet access technology today featuring over 100 million subscribers worldwide. [1]

The modern broadband market requires service providers to offer more than traditional data Internet access to DSL subscribers. Supporting delivery of combined voice, data and video services - or Triple Play over the same DSL connection has become an absolute necessity to stay competitive in the broadband access market.

Delivery of Triple Play services places a spectrum of new functional, performance and quality of service requirements on the broadband DSL networks and devices such as DSLAMs and B-RASs that make up DSL networks infrastructure. Multicast & IGMP Snooping support along with multi-megabit bandwidth per subscriber are required for IPTV service delivery. VoD service takes the bandwidth requirements even further. Fragile VoIP traffic is extremely sensitive to delay and jitter, while IPTV traffic is particularly sensitive to packet loss. Both video and VoIP traffic need to be prioritized against the data services with uneven and unpredictable bandwidth utilization. The list can go on and on.

Considering these trends, it is natural that testing and validating the performance and quality of service of Triple-Play capable DSL access networks and devices has become an issue of strategic importance for service providers and equipment manufacturers alike.

Responding to this trend, test equipment manufacturers develop new testing tools that generate realistic Triple-Play traffic and measure its performance characteristics as it is processed by the DSL network components. Benchmarking tools are complemented with the new comprehensive test methodologies that describe real-world Triple-Play traffic models, identify key Triple-Play performance metrics and analyze them under variety of stress and saturation network conditions.

Developing effective testing methodologies for Triple-Play DSL networks requires deep understanding of the DSL technology as well as functionality and architecture of the devices that form DSL network. The purpose of this whitepaper is to provide insight into architecture, functionality and performance characteristics of the DSLAM and B-RAS devices that make up DSL network infrastructure and deliver broadband access connectivity to DSL service subscriber.

DSL Technology Overview

DSL overview

Digital Subscriber Line (DSL) is a broadband access technology that enables high-speed data transmissions over the existing copper telephone wires (“local loops”) that connect subscriber’s homes or offices to their local telephone company Central Offices (COs). Contrary to the analog modem network access that uses up to 4kHz signal frequencies on the telephone wires and is limited to 56Kbps data rates, DSL is able to achieve up to 52Mbps data transmission rates by using advanced signal modulation technologies in the 25kHz and 1.1Mhz frequency range.

DSL flavors

There are a number of different DSL standards defined by American National Standards Institute (ANSI) and European Telecommunications Standards Institute (ETSI) and embraced by the industry. These DSL technology variants are typically characterized by different upstream and downstream data rates, maximum wire lengths and designated customer applications – residential, small office or business oriented. Collectively, the DSL standards are referred to as xDSL.

Roughly, xDSL standards can be divided into the following three groups:

I. Symmetric DSL – provides the same data rate for upstream and downstream transmissions and includes the following types:

DSL Variant	Max Up/ Downstream Rate	Max local loop wire length
HDSL - High data rate Digital Subscriber Line	1.5Mbps/1.5Mbps	3.7 km
SDSL - Symmetric Digital Subscriber Line	2.3Mbps/2.3Mbps	3 km
SHDSL - Symmetric High bit rate Digital Subscriber Line	4.6Mbps/4.6Mbps	5 km

II. Asymmetric DSL – provides higher downstream then upstream data transmission rates and includes the following types:

1. Windsor Oaks Group LLC. “Broadband Trends Report 1Q05: Global Broadband Subscribers Exceed 166 million”, June 2005

DSL Variant	Max Up/ Downstream Rate	Max local loop wire length
ADSL – Asymmetric Digital Subscriber Line	1Mbps/10Mbps	5.5km
ADSL Lite - Asymmetric Digital Subscriber Line Lite	384Kbps/1.5Mbps	5.5km
ADSL 2 - Asymmetric	1Mbps/12Mbps	5.5km
ADSL 2+ - Asymmetric Digital Subscriber Line 2+	1Mbps/20Mbps	5.5km
ADSL 2++ or ADSL 4 - Asymmetric Digital Subscriber Line 2++	52Mbps over short distances	Developing technology

III. Symmetric and Asymmetric DSL – can transmit data both symmetrically and asymmetrically and includes the following type:

DSL Variant	Max Up/ Downstream Rate	Max local loop wire length
VDSL – Very High bit rate Digital Subscriber Line	10Mbps/10Mbps symmetric 1.5Mbps/52Mbps	0.3km – 1.3km
VDSL 2 – Very High bit rate Digital Subscriber Line	100Mbps/ 100Mbps symmetric	.5 km

Asymmetric Digital Subscriber Line (ADSL) variants are by far the most popular DSL implementations mostly due to its suitability for Internet browsing applications that are heavily geared towards downstream data transmission (download):

Windsor Oaks Group LLC 2005 Broadband Trends Report [2]

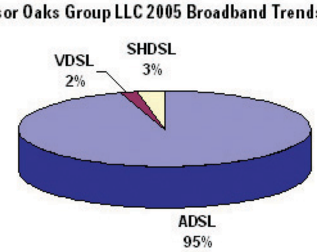

DSL alternatives

DSL is not the only broadband access technology on the market capable of delivering multi-megabit data rates to service subscribers. The notable alternatives are cable network access via television conduits, satellite network solutions like High Earth Orbit Satellite or Direct TV, other wireless technologies such as WiMax and of course the legacy T1/T3 leased lines.

Among the alternative broadband technologies cable networks & operators present the fiercest competition for DSL networks and service providers (which are traditionally Telcos). Cable access networks provide up to 10Mbps downstream bandwidth and are often Triple-Play-ready with their traditional broadcast video and Internet access services.

Although all alternative broadband technologies have their advantages, DSL is the most cost effective option due to it's ability to utilize nearly 700 million telephone lines installed worldwide for multi-megabit data access without extensive and expensive infrastructure upgrades.

Consequently DSL is the most popular and widespread broadband access technology to date, accounting for appr 66% of the worldwide 166.4 million subscriber base:

Windsor Oaks Group LLC Broadband Trends & 2005 Market Outlook Reports [3]

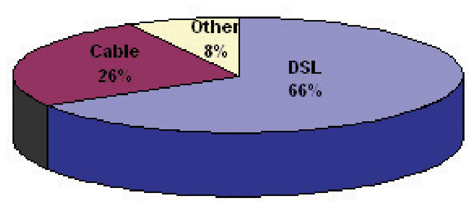

Future growth analysis by the same body forecasts the worldwide subscriber volume to reach 422 million by 2010, of which DSL is expected to account for 70%.

Competing for the enormous revenues in the broadband access market, DSL network operators are using bandwidth, performance and reliability of their networks as well as value added services such as VoIP, IPTV, VoD and online games (often bundled in attractive Triple-Play service packages) as key differentiating factors against their respective competitors – DSL and otherwise. It is thus critical for those operators to extensively test their network infrastructure and Triple-Play services to ensure the compliance with their marketing claims.

2&3. Windsor Oaks Group LLC. "Broadband Trends Report 1Q05: Global Broadband Subscribers Exceed 166 million", June 2005

DSL infrastructure building blocks – DSLAM and B-RAS devices

When digital data is sent from a DSL subscriber's premises, it travels from subscriber's computer or network through a DSL modem and on to the other end of the line at the phone company's Central Office (CO). At the CO end of the line (local loop) the data is received by the Digital Subscriber Line Access Multiplexer (DSLAM). The DSLAM aggregates the digital data streams coming from a number of subscribers onto a single high-capacity uplink (ATM or Gigabit Ethernet backhaul) to the Internet Service Provider. At the ISP the aggregated data from multiple subscribers is processed by the Broadband Remote Access Server (B-RAS) which authenticates the subscriber's credentials, validates the users access policies and routes the data to its respective destinations on the Internet.

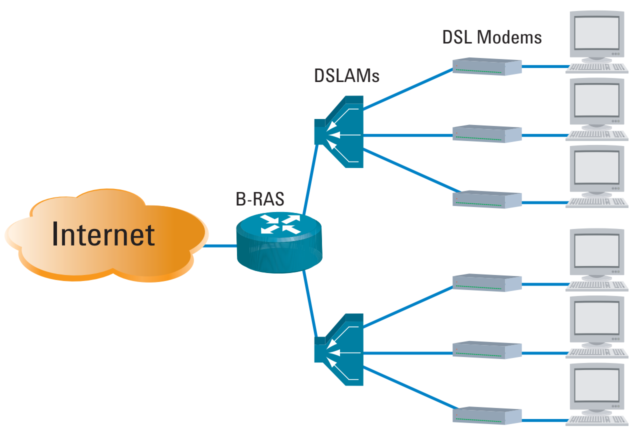

This is an extremely simplistic outline of the DSL access network flow but it carries the message that what really makes DSL happen are the DSL modems and DSLAM and B-RAS devices.

The following chapters will concentrate on the DSLAM and B-RAS architecture, functionality and classification as well as mention the performance and scalability challenges these devices face in modern Triple-Play networks.

DSLAM Architecture, Functionality and Performance

DSLAMs overview

The Digital Subscriber Line Access Multiplexer or DSLAM is the equipment that really allows the DSL to happen. The DSLAM handles the high-speed digital data streams coming from numerous subscribers' DSL modems and aggregates it onto a single high-capacity uplink – ATM or Gigabit Ethernet to the Internet Service Provider.

Contemporary DSLAMs typically support multiple DSL transmission types – ADSL, SDSL, etc as well as different protocol and modulation technologies within the same DSL type.

Responding to the requirements posed by broadband network evolution towards provision of value added services such as VoDSL and IPTV, modern DSLAMs, in addition to DSL aggregation functions, begin to provide advanced functionality such as traffic management, QoS, authentication via DHCP Relay, IGMP Snooping as well as in some cases IP routing and security enforcement.

Following rapid growth in DSL broadband network access popularity and subscriber base, revenues from DSLAM equipment sales are on the rise as well and have reached record $5 billion in 2004, according to Infonetics. Considering this growing market, DSLAM equipment development has become an extremely hot area in which many leading vendors compete for leadership and market share. Along with advanced functionality aspects, capacity, performance and scalability of a DSLAM have become key differentiating factors on which purchasing and deployment decisions are often cast.

DSLAM functionality evolution

ATM DSLAMs

As the ATM was the main high-speed data backbone transport used in Telecommunications networks during the initial DSL rollout (1999-2001), the typical DSL network access architecture deployed at that time used ATM Permanent Virtual Circuits (PVCs) from the subscriber via DSLAM to B-RAS, at which point the PPP sessions were terminated and the traffic was routed to core network. In this architecture the first generation DSLAMs with ATM uplink port or ATM DSLAMs were designed as simple Layer-2 ATM multiplexers or concentrators that provided seamless integration of the "last mile" ATM over DSL links into the ATM access network.

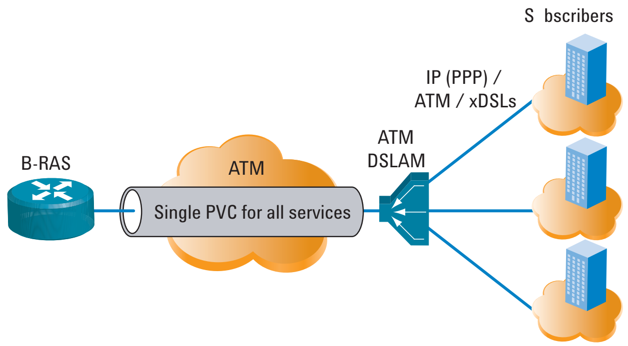

Second Generation ATM DSLAMs

The first generation ATM DSLAMs, were perfectly adequate for aggregating “best effort” services – typically Internet surfing, and used a single level of QoS over its PVC connection – usually Unspecified Bit Rate or UBR. As service providers expanded their DSL networks to business customers and began offering SLA-based value-added services such as FRoDSL (Frame Relay over DSL), VPN (Virtual Private Networks) and VoDSL (Voice over DSL), the single “best effort” QoS level of first generation ATM DSLAMs has become insufficient.

In response to this trend, the second generation ATM DSLAMs were built to incorporate the ATM switching fabric and fully utilize the Switched Virtual Circuits (SVCs) and all of the class of service, traffic shaping and traffic prioritization capabilities inherent with ATM. The ATM DSLAMs with ATM switching capability thus enabled service providers to offer business customers prioritized SVCs for voice traffic, frame relay or VPN services and low-priority “best effort” SVCs for Internet surfing to home users.

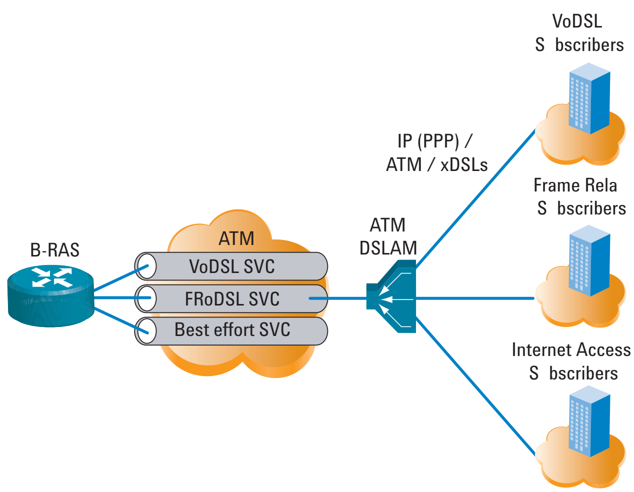

Ethernet or IP-DSLAMs

Further quest for more profitable value-added services such as VoIP, IPTV, VoD and HDTV in addition to high-speed Internet access (combination known as Triple-Play) has placed new bandwidth, scalability and QoS requirements before the DSL network providers. While existing ATM based networks had the required QoS capabilities, their high deployment and maintenance cost (cost of ownership) has caused the DSL network providers to look at Ethernet and IP-based architectures as an alternative to ATM backhaul. As Ethernet standards such as Metro Ethernet have evolved to provide the resilience and quality required for carrier network backbones, and with advent of Gigabit and 10-Gigabit Ethernet delivering the superior to ATM bandwidth, the Ethernet has become a transport of choice for both carrier backbone and access network segments.

Following this trend the new generation of DSLAMs has appeared that used Ethernet uplinks for DSL traffic aggregation. These devices have become known as Ethernet DSLAMs or IP-DSLAMs. In it’s simplest implementation IP-DSLAMs function as Layer-2 switches that backhaul subscriber traffic to Metro B-RASs or Broadband Network Gateways (BBNGs) using Ethernet VLANs in combination with Ethernet Multicast capability.

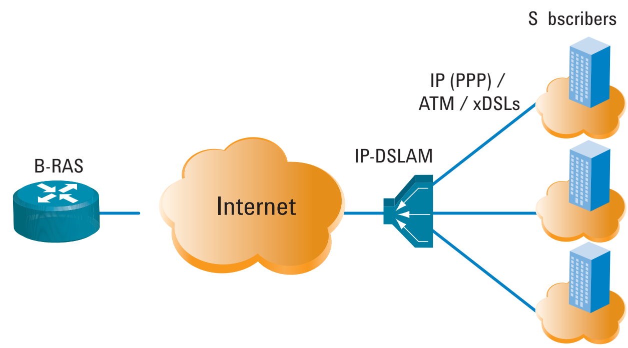

While already using Ethernet or Gigabit Ethernet uplinks, IP-DSLAMs still typically use ATM on the local loop as an intermediate layer between Ethernet and DSL – mostly for reasons of compatibility with existing xDSL modems.

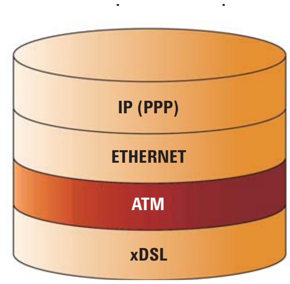

In recent times though, native IP (or IP over PPP) over Ethernet over DSL - or Ethernet in the First Mile implementations are increasingly being adopted by DSL service providers. Without the ATM layer overhead, local loop segments become more efficient and cost-effective.

Ethernet on First Mile (EMF) xDSL protocol stack

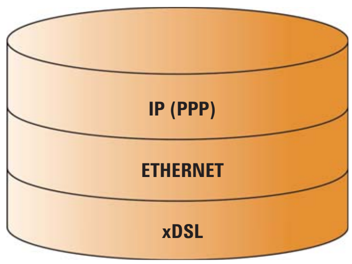

In some cases, advanced functionality ranging from local PPP and RFC1483 session termination (B-RAS off-load capability) to full IP routing, AAA, security, 802.1p prioritization and DiffServ QoS is incorporated into Ethernet DSLAMs, resulting in truly IP-enabled IP-DSLAMs.

However, as various industry surveys indicate (such as Heavy Reading 2005 Next-Generation DSL Equipment: The Path to Profitability Report), those truly IP-enabled IP-DSLAMs or IP-DSLAM/B-RAS hybrid devices are still a minority and most of IP-DSLAMs being deployed today are really Ethernet DSLAMs with basic multicast and IGMP Snooping or IGMP Proxy support.

The same surveys agree that the current killer configuration for advanced DSL service deployment consists of Ethernet DSLAMs with limited Layer-3 capability (i.e. mentioned multicast and IGMP support) and high-capacity Metro B-RASs.

Nevertheless it is important to note that the industry trend is definitely towards more advanced Layer-3 IP functionality on the DSLAMs and possibly towards the DSLAM/B-RAS convergence in the future high-capacity DSL network implementations.

RFC1483	Defines two encapsulation methods for carrying network traffic over ATM: routed protocol data units (PDUs) and bridged PDUs.
AAA	Authentication, Authorization and Accounting Services
802.1p	Layer 2 QoS protocol that provides for traffic prioritization and dynamic multicast filtering at MAC layer
DiffServ	Layer 3 IP QoS protocol that utilizes IP TOS packet field for carrying QoS information
IGMP Snooping	Technology that allows Layer 2 devices to examine IGMP messages (Query, Report & Leave) exchanged by hosts and multicast routers and configure relevant multicast forwarding table
IGMP Proxy	Issues IGMP host messages on behalf of hosts that are not directly connected to downstream multicast router
IGMP protocol	Protocol widely used in IPTV applications for establishing subscriber connections to TV broadcasting channels

DSLAM architecture

From the high-level perspective ATM DSLAMs, Ethernet DSLAMs and IP-DSLAMs architecture typically includes a number of xDSL line cards that terminate the subscriber local loops and one or more ATM OC-3/12/48 or Ethernet/ Gigabit Ethernet uplink cards for traffic backhaul. The line cards and uplink cards are interconnected by a high-capacity aggregation backplane that can take form of an ATM or Ethernet bridge or switch. Majority of modern DSLAMs are multiservice and support multiple DSL technologies – i.e. ADSL, ADSL2, ADSL2+, SDSL and VDSL, etc and therefore these devices accommodate for multiple xDSL line card types.

High-level DSLAM Architect re Diagram

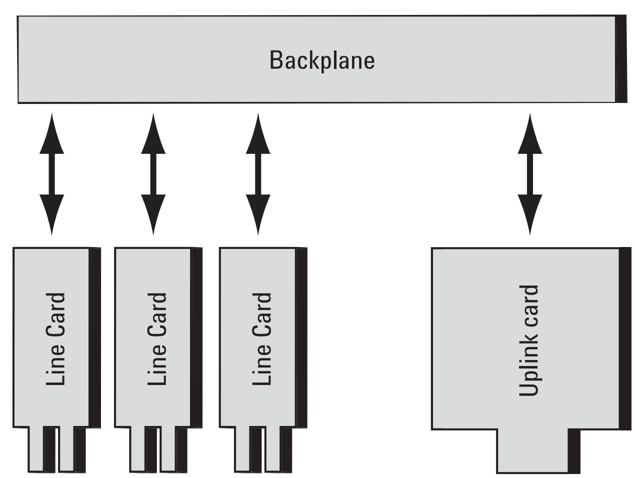

From the traffic processing perspective, two distinct architecture models have emerged – centralized and distributed.

In the centralized model all complex traffic processing (e.g. classification, filtering, QoS, etc) is performed on a single central uplink card. The line cards in centralized model are "dumb" and cheap and contain only the basic components required for traffic hand-off to the uplink card. Centralized architecture is considered best suited for high-density large-scale aggregation-centered DSLAMs with moderate complex traffic processing requirements. Example of centralized DSLAM implementations are products based on Intel IXP2400 NP design.

In the distributed model some or all complex traffic processing is off-loaded to the smart line cards based on programmable network processors (Linecard Traffic Processors or LTPs). The uplink card in such architecture can be as simple as an Ethernet switch in case of Ethernet backhaul, or still require a full-featured network processor for more complex scenarios such as IPoMPLS backhaul.

The distributed model is prevalent in DSLAMs with complex traffic processing capabilities – such as IP-DSLAMs with Layer-3 IP functionality, AAA, QoS and security enforcement.

The distributed DSLAM architecture model has a number of important advantages over the centralized model – such as local traffic processing capability on the line cards (local multicasting and local peer-to-peer traffic switching) and linear DSLAM shelf expansion costs (i.e. sparsely populated DSLAM shelves with inexpensive uplinks can provide low entry cost and its capacity can be expanded in "pay as you grow" fashion as additional line cards are installed).

DSLAM classification

Although no established formal classification exists, DSLAMs can be divided into distinct classes or types by a variety of criteria:

By supported xDSL service type – single service DSLAMs and multiservice DSLAMs. Majority of modern DSLAMs are multiservice. Single service DSLAMs are typically small devices deployed in apartment complexes or university campuses and serving relatively limited number of homogenous subscribers.

By architecture – centralized or driven by central processing unit typically embedded in the uplink module and distributed where some or all traffic processing is done on the line cards each equipped with powerful NP. As mentioned in the previous chapter, majority of high-end modern DSLAMs have distributed architecture.

By deployment location – CO (Central Office) deployed and OSP (Outside Plant) deployed. CO DSLAMs typically have distributed shelf architecture and a much larger subscriber capacity – of up to 10,000 lines, that is justified considering large number of served customers. They also tend to be more functionally sophisticated providing Layer-3 and other advanced functionality. OSP DSLAMs are deployed in remote locations closer to subscribers largely in order to shorten the local loop wire lengths and thus achieve higher bandwidths and service quality. OSP DSLAMs typically have smaller subscriber capacity (usually in dozens of lines), smaller footprint and are hardened for protection against elements. OSP DSLAMs can nevertheless have very high performance characteristics due to the high bandwidth xDSL technologies such as ADSL2++ that it may use on the local loop.

By hardware model and form factor – DSLAMs hardware implementations range from chassis based high-capacity DSLAMs with pluggable line cards and uplink modules to standalone pizza box devices with limited number of ports and a single T1 ATM or 10/100Base Ethernet uplinks.

DSLAM capacity, performance and scalability metrics

Although performance characteristics and performance analysis of DSLAMs are complex issues and a central topic of a dedicated Agilent whitepaper, it is still worthwhile to mention the key metrics that typically characterize DSLAM devices from their performance and scalability standpoints:

Subscriber capacity

As DSLAMs provide services for multiple subscribers via DSL ports the most basic DSLAM capacity metrics are line density, subscriber and session capacity. Depending on the type of DSLAM and functionality provided, these metrics can have one-to-one or one-to-many ratios. As mentioned in the architecture section, DSLAMs can have line, subscriber and session capacities ranging from single digits for standalone devices to tens of thousands for high-capacity CO DSLAMs.

Along with maximal capacity metrics, scalability and performance parameters such as line or session ramp-up and shutdown rates are extremely important for DSLAM sizing and performance assessment.

Throughput

As DSLAMs must provide committed to bandwidth to all served subscribers, bandwidth or throughput characteristics of a DSLAM play key role in sizing and performance analysis. Typically DSLAMs are measured for their aggregate sustainable throughput using all line interfaces using different types of traffic (packet sizes, session volumes) and in various subscriber traffic processing scenarios (with IGMP Snooping, QoS, AAA, etc – depending on the DSLAM capabilities).

Packet loss, latency and jitter

As modern Triple-Play services such as VoDSL, VoIP, IPTV and VoD supported by DSLAMs are extremely sensitive to packet loss, latency and jitter network parameters, these characteristics as exhibited by DSLAMs while processing subscriber traffic play extremely important role in DSLAM performance assessment. Typically loss, latency and jitter parameters are measured under traffic load at various degrees of device capacity saturation, with different types of traffic and in variety of traffic processing scenarios.

Functionality or service specific performance metrics

As modern DSLAMs (IP-DSLAMs) often provide capabilities beyond traffic aggregation – such as QoS, AAA, B-RAS off-loading (session termination), security, etc – performance characteristics specific to these capabilities (such as PPP session capacity, session establishment and termination rates, authentication rates or QoS quality under load) need to be measured and analyzed in course of relevant DSLAM type performance evaluation.

Additionally, specific types of provided service – such as VoDSL, VoIP and IPTV impose distinct and specific performance criteria on the devices along the path and have distinct performance metrics typical for these services – such as voice call clarity scores or TV channel switching latencies. In this context, DSLAMs supporting these services functionally or as a matter of being present on the delivery path are typically evaluated for those service specific performance metrics using specially designed test methodologies.

DSLAM market - major players

Considering the growing DSL broadband networks deployment and DSL hardware revenue statistics, DSLAM and IP-DSLAM equipment market is an extremely hot and competitive area, with dozens of North American, European and Asian vendors competing for technological superiority and market share.

Although the complete list[4] of DSLAM manufacturers would be too long to include, the leading positions are clearly occupied by very few manufacturers who represent the de-facto industry standards. Specifically, Alcatel holds first place at appr. 32.5% of the market with Huawei being distant second at appr. 17% and Lucent and Siemens battling for the third place.

B-RAS Architecture, Functionality and Performance

B-RAS overview

The Broadband Remote Access Server (B-RAS) is a key component of DSL broadband access networks that serves as an aggregation point for subscriber traffic (IP, PPP and ATM) and provides session termination (PPPoX, RFC 1483) and subscriber management functions such as authentication, authorization, accounting (AAA), and IP address assignment.

Triggered by the new functional requirements that Triple-Play service delivery imposed on the DSL network infrastructure and devices, modern B-RASs started to provide advanced services beyond traffic aggregation – such as subscriber policy management (e.g. Web login or differentiated access management based on factors such as traffic volume or time of day), IP and Layer-2 QoS, security enforcement and VPN capabilities, as well as full IP routing and MPLS support.

As the complexity of B-RAS devices increased significantly, in order to streamline the B-RAS evolution and ensure inter-vendor interoperability, the DSL Forum has released the Technical Report 092 (TR-092), which defined functional requirements towards B-RAS devices in modern Triple-Play enabled DSL network environment (which in turn is analyzed in DSL Forum Technical Report 059 – TR-059). B-RAS vendors are now using the TR-092 and TR-059 as guidelines for product development and feature roadmap planning to ensure their platforms competitiveness and compatibility with the next-generation DSL broadband network architectures.

4. Windsor Oaks Group LLC. “Market Share Snapshot: DSL Port Shipments Reach Record Levels in 1Q05”, May 2005

As in case of DSLAMs, the exponential growth in DSL broadband access networks deployment for high-speed Internet access and various value-added services has caused B-RAS equipment sales to rise exponentially in recent years.

Although, due to the established trend of B-RAS functionality being incorporated into Edge Routers with a much broader range of applications (more on that trend in following sections), it is difficult to precisely estimate B-RAS market revenues. However, industry reports indicate that B-RAS applications make up 29% of the total IP Edge Router revenue (Infonetics), which is said to have reached $2.1B in 2004.[5]

B-RAS functionality evolution

B-RAS device heritage can be traced back to dial-up Remote Access Servers (RASs) that terminated user PPP sessions established over the analog phone lines, authenticated remote caller credentials and provided connectivity to the Internet. As the Internet market continued to explode, and with rapid growth of bandwidth-intensive network applications, dial-up access limited to 56Kbps speeds has become insufficient and given way to broadband network access technologies – primarily DSL. As PPP remained the protocol of choice for tunneling of subscriber connections over ATM and DSL lines, Broadband Remote Access Servers (B-RASs) have replaced RASs as PPP termination devices that authenticated user credentials and routed the subscriber traffic on to the Service Provider networks and the Internet.

As DSL broadband service delivery continued to evolve from high-speed Internet access to a wide variety of Triple-Play offerings, the fundamental functionality of B-RAS has also changed:

First Generation – software-based B-RASs:

As the DSL service was initially introduced in late 1990s, competition was limited and service was first offered to early adopters. Industry acceptance was slow and B-RAS vendors often tailored their products to each service provider's environment and requirements. The First Generation B-RASs were software-based and used general-purpose hardware platforms to allow for rapid customization and prototyping – which led to low performance specifications. Functionally, first generation B-RASs provided PPP session termination and subscriber management functions such as AAA and IP address assignment.

Second Generation – centralized architecture, hardware-based B-RASs:

During 2000 – 2003, DSL service was becoming a mainstream broadband access technology with standardized delivery architectures. With soaring competition and bandwidth serving as the only differentiator, DSL service began to commoditize. Consequently, providers started looking into value-added services as potential new sources of revenues - Triple-Play was born. Accommodating for these trends, Second generation B-RASs were implemented in hardware and had much higher performance specs. At the same time, functional scalability was limited as B-RAS devices were still designed around centralized, processor-based architectures and were optimized for single-service – Internet access. Attempts to add advanced features such as filtering or QoS still led to performance degradation.

Third Generation – Triple-Play and application-aware B-RASs:

In 2004 DSL deployment worldwide reached critical mass. Severe competition drove down basic service prices. As a result, DSL broadband access providers turned to alternative revenue sources by offering differentiated services. Advanced services were widely deployed that included FRoDSL, VPNs, VoIP, IPTV, VoD and interactive gaming. Triple-Play has become mainstream part of the broadband service portfolio. In order to comply with the new requirements, Third Generation of B-RAS devices has undergone a significant shift in functionality to support the complete range of high-bandwidth, multimedia-intensive Triple-Play services.

In third generation of B-RASs centralized hardware model gave way to modular, highly scalable distributed architecture that allowed service providers to deliver the session capacity and throughput required to support advanced broadband service delivery. At this stage the B-RAS functionality has also begun to be integrated into Service Edge Routers that provided the following capabilities in one unified platform:

- ATM and Ethernet aggregation
- Session termination – ATM PVC, PPP
- AAA - authentication, authorization, accounting using PAP/CHAP, RADIUS, DHCP Option 82
- Comprehensive IP routing – BGP, OSPF, RIP
- IP address management – DHCP server, relay, proxy services

5. Infonetics. "Service Provider Routers and Switches: Quarterly Report", August 2005

- Integrated layer 2 switching – ATM, Ethernet, MPLS
- Policy management and dynamic per session QoS
- IP multicast routing – PIM, MBGP, IGMP

In addition to these capabilities, high-end third generation B-RASs or Service Edge Routers typically integrate Firewall-grade security enforcement, intrusion detection and prevention mechanisms, and content filtering as well as fault tolerance provisions such as switch fabric redundancy and hot swappable modules to ensure high availability and near zero downtime.

PAP/CHAP	Authentication protocols defined in RFC1334 and used by Point-to-Point Protocol (PPP).
RADIUS Remote Authentication Dial In User Service	Authentication, authorization and accounting protocol defined in RFC2865.
DHCP option 82 – Relay Agent Information Option	An extension to DHCP protocol defined in RFC3046. DHCP option 82, in DSL environments can be used for subscriber line authentication.

B-RAS architecture and deployment classification

From the architecture model standpoint, B-RASs fall into three distinct categories:

Software-based B-RAS architecture, using general-purpose hardware platforms was typical for early first-generation B-RASs.

Centralized hardware-based B-RAS architecture with majority of intelligence placed in the central processing unit. Centralized B-RASs were designed and optimized for High-Speed Internet (HSI) services, had high performance characteristics but did not scale well.

Distributed or modular hardware based B-RAS architecture with functional intelligence distributed between ASIC-driven line cards of variety of technologies (Gigabit Ethernet, OC-12, OC-48) and routing or switching modules (Edge Service Processors). Distributed B-RAS architecture provides for exceptional flexibility, scalability and power and allows service providers to deploy the number and types of interfaces as well as functionality portfolio that fit their requirements and budget, while retaining the ability to add capacity and capabilities as the needs grow. The majority of modern B-RASs or Service Edge Routers have a modular distributed architecture.

By deployment scenario, B-RASs can be roughly divided into the following classes:

Centrally deployed – in PoPs (Service Provider Points of Presence) merged with Edge Router functionality and acting as an aggregation router for multiple DSLAMs located in the COs (Central Offices). Centralized deployment model advantages include operational efficiency (a few high-end Metro B-RASs can serve very large numbers of DSLAMs/subscribers with relative ease of management), scalability and versatility (PoP-located B-RASs can also serve non-DSL customers). As a result, the centralized B-RAS deployment model is considered to be more cost effective and better suited for providing advanced broadband services and thus the overwhelming majority of B-RASs today are centrally deployed. [6]

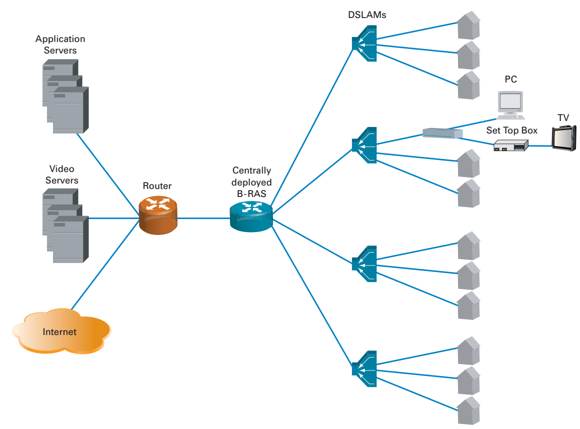

Distributed to COs, co-located with or within IP-DSLAMs. The distributed B-RAS deployment model has several advantages: Uplinks from the COs, which become the IP B-RAS uplinks, can benefit from IP DiffServ QoS as well the early filtering of the unauthorized traffic. At the same time, the distributed deployment model has significant drawbacks including potential instability of the resulting large routing network and potential RADUIS server overload (since each B-RAS is typically a RADUIS authentication client).

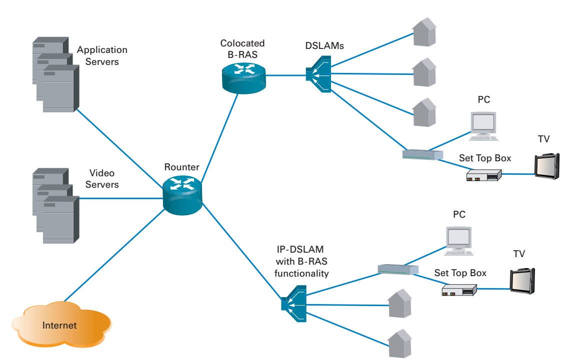

6. Heavy-Reading. "Next Generation DSL Equipment: The Path to Profitability", November 2003

B-RAS capacity, performance and scalability metrics

The B-RASs and Service Edge Routers with integrated B-RAS capabilities are intelligent network devices with a wide spectrum of advanced network functionality. Thus comprehensive analysis of the capacity and performance of these devices is a complex task that requires sophisticated test methodologies and state-of-the-art benchmarking tools.

Broadband subscriber session capacity and session handling performance

As the B-RAS's primary role is subscriber session aggregation and termination, PPPoX session capacity and session-handling performance (such as session ramp-up and tear-down rate) metrics are critical for B-RAS performance analysis and sizing. Typically, session capacity and session handling performance are measured with various subscriber management functions (such as RADIUS authentication or DHCP address allocation) enabled for each session to determine realistic performance limits.

Data plane performance – throughput

To accommodate modern bandwidth intensive Triple-Play services, B-RAS devices are typically required to deliver wire-rate throughput on all aggregation ports with all traffic services (IP routing, Layer-2 switching, QoS, security, etc) enabled on every port and for every session. Consequently, throughput and throughput-under-load measurements a play critical role in B-RAS performance analysis and are performed with different traffic types, using the full ranges of supported interfaces and all supported traffic processing features.

Real World Traffic performance

Modern B-RASs are complex devices that perform protocol- and transaction-dependent processing (user authentication, QoS, security enforcement and URL filtering). Therefore, the traffic structure or profile of traffic being handled by the B-RAS has a significant impact on device performance watermarks (such as aggregate goodput, voice or video quality and other application-layer metrics). In this context, in order to determine real B-RAS performance limits, Real World performance tests must be performed using realistic traffic profiles containing various duration PPPoX sessions, real application transactions (HTTP, VoIP, streaming video) and in various traffic processing scenarios (including advanced services such as security enforcement or data encryption).

Packet loss, latency and jitter

As with DSLAMs, the packet loss, latency and jitter traffic introduced by B-RASs during subscriber traffic processing play an extremely important role in B-RAS performance assessment - especially in the context of Triple-Play service delivery. Typically these performance parameters are measured under traffic load at various degrees of device capacity saturation, with different types of traffic and in a variety of traffic-processing scenarios.

Functionality or service-specific performance metrics

B-RASs or Service Edge Routers perform a wide variety of network functions from subscriber management to routing and MPLS switching. Therefore it is necessary to measure and analyze performance characteristics specific to these capabilities (such as authentication rates using different authentication schemes, DHCP maximum session rate and routing protocol performance) as part of comprehensive B-RAS or Service Edge Router performance evaluation.

Additionally, advanced services – such as IPSec VPN, VoIP and IPTV - impose distinct and specific performance criteria on the B-RAS devices. These services require the B-RAS to perform complex traffic processing operations (such as data encryption, QoS and multicast forwarding). Consequently, the B-RAS performance characteristics specific to these services – such as encrypted data throughput, voice call clarity scores and TV channel switching latencies need to be measured using specially designed test methodologies.

B-RAS and Service Edge Router market - major players

The majority of modern B-RAS implementations have been realized within Service Edge Routers, a single platform that provides DSL traffic aggregation, session termination, subscriber management, QoS, Internet routing, MPLS signaling and other advanced services. As with the DSLAM market, the B-RAS and Service Edge Router market is so diverse and competitive that a complete list of vendors would likely require a separate volume. Like the DSLAM market, the Service Edge Router market is dominated by a very limited number of vendors – Juniper, Cisco, Redback and Nortel. Together, these vendors are responsible for over 90% of the revenues and units shipped. [7]

7. Gartner. "Market Share: Service Provider Routers, Worldwide, 3Q05", December 2005

Summary

Triple Play service deployment has imposed a plethora of new functional, performance and capacity requirements on the key devices that make up DSL networks – IP-DSLAMs and B-RASs. This has compelled us, as a test equipment manufacturer, to develop state-of-the-art testing tools and test methodologies allowing our customers to ensure that their devices comply with requirements and deliver the desired performance and quality of service. Understanding the architecture and functionality of the DSL network devices is the first and most critical step in recognizing the test needs and comprehending the new test methodologies.

This page intentionally left blank.

This page intentionally left blank.

This page intentionally left blank.

www.agilent.com/comms/n2x

Sales, Service and Support

United States:
Agilent Technologies
Test and Measurement Call Center
P.O. Box 4026
Englewood, CO 80155-4026
1-800-452-4844

Canada:
Agilent Technologies Canada Inc.
2660 Matheson Blvd. E
Mississauga, Ontario
L4W 5M2
1-877-894-4414

Europe:
Agilent Technologies
European Marketing Organisation
P.O. Box 999
1180 AZ Amstelveen
The Netherlands
(31 20) 547-2323

United Kingdom
07004 666666

Japan:
Agilent Technologies Japan Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi,
Tokyo 192-8510, Japan
Tel: (81) 426-56-7832
Fax: (81) 426-56-7840

Latin America:
Agilent Technologies
Latin American Region Headquarters
5200 Blue Lagoon Drive, Suite #950
Miami, Florida 33126
U.S.A.
Tel: (305) 269-7500
Fax: (305) 267-4286

Asia Pacific:
Agilent Technologies
19/F, Cityplaza One, 1111 King's Road,
Taikoo Shing, Hong Kong, SAR
Tel: (852) 3197-7777
Fax: (852) 2506-9233

Australia/New Zealand:
Agilent Technologies Australia Pty Ltd
347 Burwood Highway
Forest Hill, Victoria 3131
Tel: 1-800-629-485 (Australia)
Fax: (61-3) 9272-0749
Tel: 0-800-738-378 (New Zealand)
Fax: (64-4) 802-6881

Printed in USA February 2, 2006
5989-4766EN

Agilent Technologies