STRATEGIC WHITE PAPER

Optimizing the Broadband Aggregation Network for Triple Play Services

The proposed Triple Play Service Delivery Architecture provides a comprehensive approach to architecture design, which allows service providers to implement a service-rich and cost-optimized service delivery foundation that can seamlessly scale to support triple play and business service rollouts.

The service delivery architecture allows network operators to progressively integrate their HSI, voice, and video services (over both fiber and copper) within a unified and homogeneous Ethernet aggregation network environment. The key benefits of the proposed service infrastructure include cost optimization, reduced risk, and accelerated time-to-market for new services.

ALCATEL

Table of Contents

Introduction

The aggregation network needs to be optimized and versatile to meet the combined demands of voice, video, and high-speed Internet services.

The digitization of content and the convergence of voice, data, and video have compelled key players in previously non-competing industries at different levels in the value chain to fundamentally rethink their visions, strategies, and business models. Companies with historically different backgrounds and core competencies (content providers, cable companies, and incumbent service providers) will increasingly be in direct competition as they offer combined voice, video, and data services.

An improved regulatory environment has resulted in many incumbent service providers rapidly making considerable investments in order to add video to their existing voice and data services, thereby improving their competitive position in the market. Examples include SBC's Lightspeed project, with the stated goal of passing 18 million households by 2007, and Verizon's fiber to the premises (FTTP) project, with one million homes and businesses passed by the end of 2004.

All service providers competing in this market will be faced with the challenging task of offering a compelling value proposition for their service offerings in order to achieve and maintain a high average revenue per user (ARPU) and to sustain the long-term profitability of their services (see Figure 1).

For providers who successfully meet this challenge, benefits will include the ability to:

> Successfully sell new, higher-value services to their established customer bases
> Leverage their brand equity to offer high-margin bundles and minimize churn
> Progressively build their service infrastructures around profitability models, as opposed to network or infrastructure constraints
> Consolidate multiple specialized networks into a cost-optimized and efficient service infrastructure that will achieve economies of scale and maximize the profitability of services
> Streamline network operations across protocol layers to achieve economies of skill, and evolve the overall service infrastructure to a more homogeneous, scalable, and flexible architecture
> Streamline business operations to achieve significant economies of scope across service offerings (order fulfilment, support, billing, marketing)
> Future-proof their service infrastructures for the future layering of ultra-high bandwidth business and residential applications
> Minimize the overall financial and technological risks over the next 7 to 10 years

A comprehensive approach to triple play architecture design is required to enable service providers to achieve significant economies of scale, skill, and scope by evolving their second- and perhaps first-mile infrastructures to a homogeneous and cost-optimized service delivery foundation that can scale and support future ultra-high-speed service rollouts, such as 100 Mb/s to the home for multiple high-definition TV streams.

Figure 1 - Projected ARPU for Various Triple Play Services

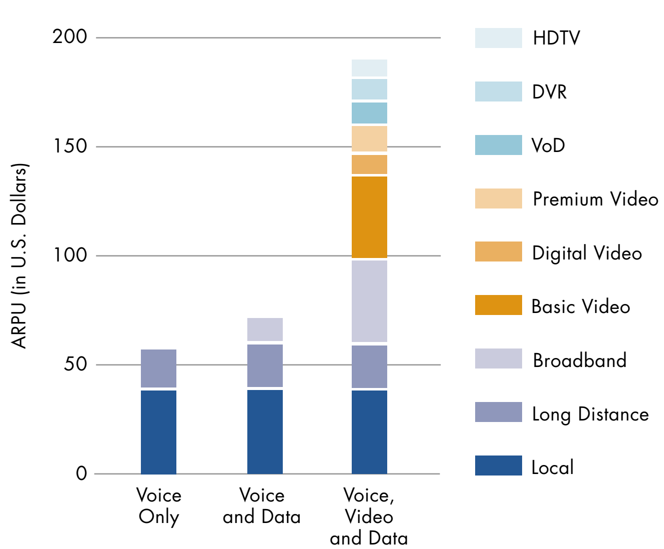

Source: The Yankee Group, 2004

Blueprint for Optimizing Triple Play Service Infrastructures

The Triple Play Service Delivery architecture described in this article allows network operators to progressively integrate their high speed internet (HSI), voice, and video services within a unified and homogeneous Ethernet-based aggregation network environment. The key benefits of the proposed service infrastructure include cost optimization, reduced risk, and accelerated time to market for new services.

At a high level, the service delivery architecture implements:

> *Ethernet-based service architecture:* Solves bandwidth bottlenecks and exponential capital expenditure and operating expenses issues in the second mile by leveraging the efficiency of this proven and widely deployed technology. A unified and homogeneous Ethernet environment across the second and first miles allows service providers to build a cost-optimized service infrastructure and achieve significant economies of scale, scope, and skill.

> *Multiple distributed service edges:* Allows service providers to achieve faster times to market for new services while retaining the existing broadband remote access servers (BRAS) using a point-to-point protocol over Ethernet (PPPoE) mode of operation for wholesale and retail HSI. It also allows network operators to leverage purpose-built platforms for each service to meet the specific requirements of each service.

> *Distributed multicasting functions in access and aggregation networks:* Enables service providers to optimize bandwidth and content delivery mechanisms, based on densities and penetration rates. It is also essential to subscriber and service scaling, and optimizes the bandwidth required in the aggregation network.

> *Carrier video and VoIP services using DHCP:* Enables service providers to introduce "plug-and-play" services delivered through set-top boxes and voice over internet protocol (VoIP) devices, which are designed for use with the dynamic host configuration protocol (DHCP). It also greatly simplifies and scales multicast video services and increases bandwidth efficiency.

> *Flexible deployment models:* The architecture allows data, video, and VoIP services to be rapidly rolled out without any lock-in to specific operational models. It allows service providers to maximize flexibility and minimize financial and technological risks by allowing all modes of operation, including:
> - copper (DSL/DSLAM) and fiber-based (FTTx) deployments in the first mile
> - single or multiple last mile circuits
> - bridged or routed home gateways
> - single or multiple IP address deployment models

Distributed Service Edges

The aggregation infrastructure for the service delivery architecture is based on two major network elements, optimized for their respective roles: the Broadband Service Aggregator (BSA) and Broadband Service Router (BSR). An important characteristic of BSAs and BSRs is that they effectively form a distributed "virtual node" with BSAs performing subscriber-specific functions where the various functions scale, and BSRs providing the routing intelligence where it is most cost-effective.

Network and service scaling are achieved by dividing the Layer 2 and Layer 3 functions between the BSA and BSR and by distributing key service delivery functions. BSAs are more distributed than BSRs, cost-effectively scaling per-subscriber policy enforcement since the required functionality is distributed to lower-cost Layer 2 aggregation interfaces instead of costly router or BRAS interfaces. The BSA incorporates Internet group management protocol (IGMP) proxy multicasting and wire speed security, per-subscriber service queuing, scheduling, accounting, and filtering. The BSA is a high-capacity Ethernet-centric aggregation device that supports hundreds of Gigabit Ethernet (GE) ports, tens of thousands of filter policies, and tens of thousands of queues.

BSAs are Layer 2 devices that forward traffic using Layer 2 mechanisms, but have the quality of service (QoS) and filtering intelligence to enforce higher-layer policies. Distribution of the QoS functionality on the BSA means that per-subscriber QoS is enforced accurately since it occurs after the last major downstream congestion point — the aggregation network. BSAs aggregate traffic for all services towards the BSR, which is an IP edge device optimized for DHCP-based video service delivery. It terminates the Layer 2 access and routes using IP over multiprotocol label switching (MPLS), with support for a full set of MPLS and IP routing protocols, including multicast routing (protocol independent multicast — sparse mode (PIM-SM) / IGMP). The BSR supports hundreds of GE and synchronous optical network (SONET) uplink ports (for large-scale deployments) and sophisticated QoS for per-service and per-content/source differentiation.

The connectivity between BSAs and BSRs is a Layer 2 forwarding model, shown in Figure 2 as a secure virtual private LAN service (VPLS) infrastructure, which refers to the fact that the BSA-BSR interconnections form a multipoint Ethernet network with security extensions to prevent unauthorized communication, denial of service, and theft of service. This approach supports all modes of operation, including multiple home gateway models, single or multiple network IP edges, and single or multiple circuits in the last mile. One of the advantages of using VPLS for this application is that VPLS instances can be automatically established over both hub-and-spoke and ring topologies, providing sub-50 ms resilience. Regardless of the fiber plant layout, VPLS enables a full mesh to be created between BSA and BSR nodes, ensuring efficient traffic distribution and resilience to node or fiber failure.

Service Differentiation, QoS Enablement

Alcatel's service delivery approach provides a model based on call admission for video and VoIP, with the need to guarantee delay/jitter/loss characteristics once the service connection is accepted. The architecture also meets the different QoS needs of HSI, namely per-subscriber bandwidth controls, including shaping and policing functions that have little or no value for video and VoIP services. In conjunction with the architecture's support for content differentiation, this enables differentiated service pricing within HSI.

The distribution of QoS policy and enforcement allows the service provider to implement meaningful per-subscriber service level controls. This is achieved by distributing the

Figure 2 - Triple Play Service Delivery Architecture

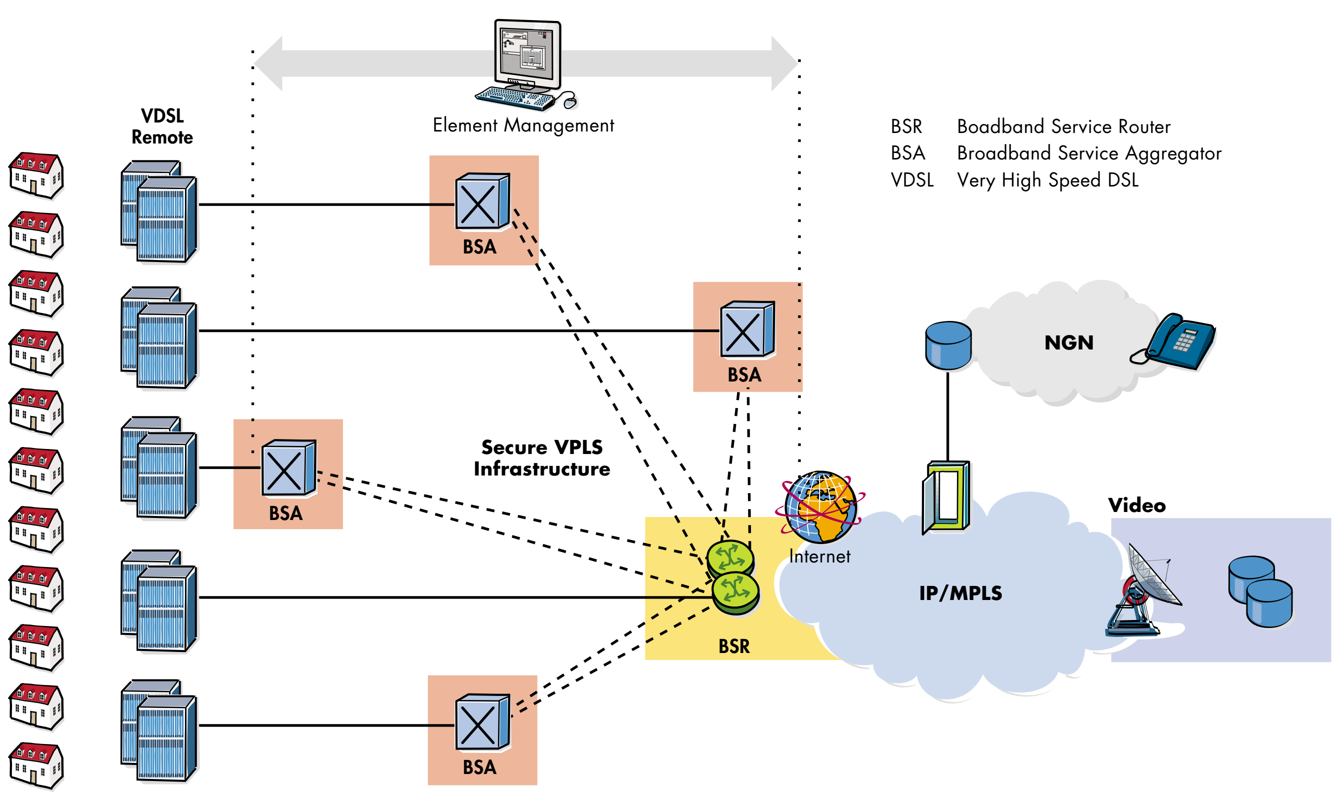

burden from hundreds of thousands of logical interfaces on high-speed edge router ports to hundreds of interfaces on lower-cost Ethernet distribution ports. Sophisticated and granular QoS feature support on the BSA allows the service provider to deliver truly differentiated IP services — differentiation based on the subscriber as well as on the content — which can be scaled cost-effectively.

The BSR performs service distribution routing based on guarantees required to deliver the service and associated content, rather than on individual subscribers. As illustrated in Figure 3, the BSR needs to classify content only based on the required forwarding class for a given BSA to ensure that each service's traffic receives the appropriate treatment towards the BSA.

In the BSR-to-BSA direction, IP services rely on IP layer classification of traffic from the network to queue traffic appropriately towards the BSA. Under extreme loading, which would be expected to occur only under network fault conditions, lower-priority data services or HSI traffic will be compromised in order to protect video and voice traffic. Classification of HSI traffic based on source network address or IEEE 802.1p marking allows the QoS information to be propagated to upstream or downstream nodes by network elements.

Figure 3 - Downstream QoS Enablement

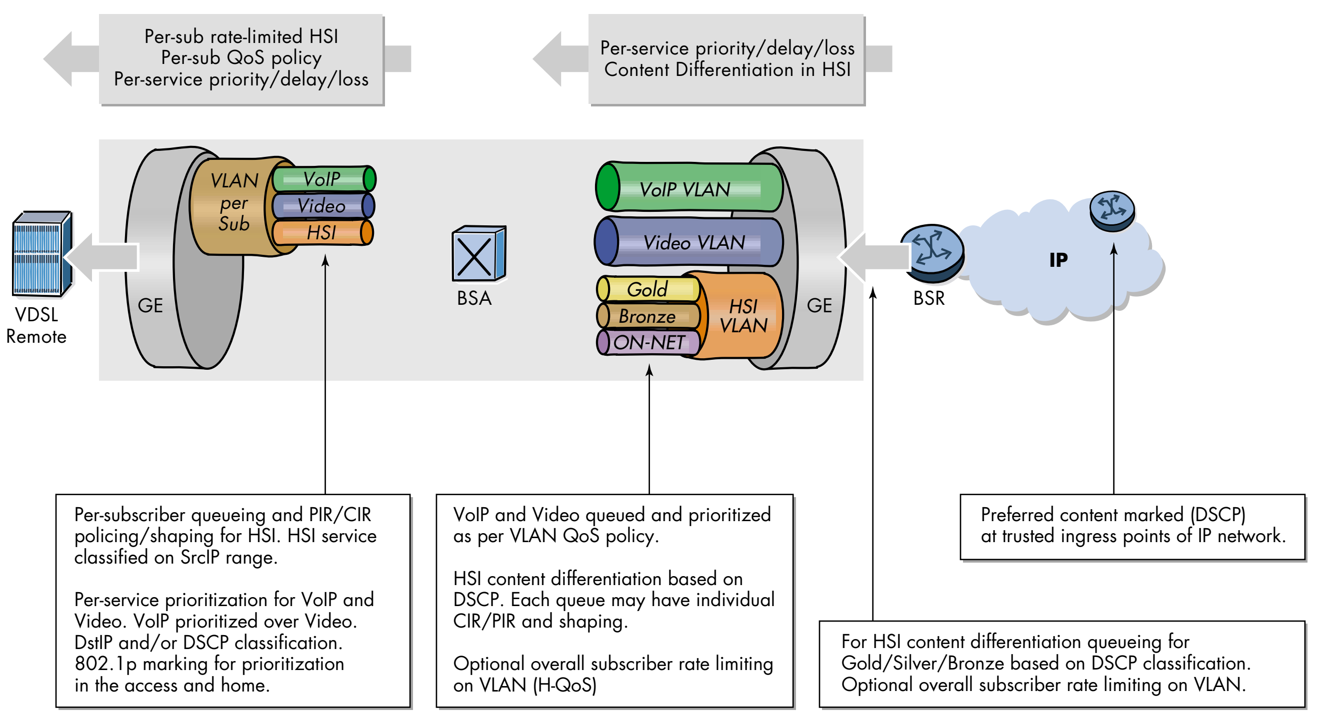

In the BSA-to-BSR upstream direction, traffic levels are substantially lower. Class-based queuing is used on the BSA network interface to ensure that video control traffic is propagated with a minimal and consistent delay, and that preferred data and HSI services receive better treatment for upstream/peering service traffic than the best effort Internet class of service (see Figure 4).

It is worthy of note that the IP edge is no longer burdened with enforcing per-subscriber policies for hundreds of thousands of users. This function is now distributed to the BSA, and the per-subscriber policies can be implemented on the interfaces facing the access node.

The BSA must be capable of scheduling and queuing functions on a per-service, per-subscriber basis, in addition to performing packet classification and filtering based on both Layer 2 and Layer 3 fields.

Each subscriber interface must provide at least three dedicated queues. Alcatel's service delivery architecture makes it possible to configure these queues so that the forwarding classes defined for all services can all be mapped to one service VLAN upstream. In the BSA node, assuming hundreds of subscribers per GE interface, this translates to a thousand or more queues per port. This contrasts with a classical non-distributed solution in which the edge router would have to support some 100,000 to 200,000 queues per 10 Gigabit port.

In addition to per-service rate limiting for HSI services, each subscriber's service traffic can be rate-limited as an aggregate using a "bundled" service policy. This allows different subscribers to receive different service levels independently and simultaneously. It is also necessary for the combined bandwidth of all services to be scheduled to an overall rate limit to allow multicast traffic to be delivered to subscribers further downstream, and thus avoid further complex queuing and scheduling of traffic in the access node.

Figure 4 - Upstream QoS Enablement*

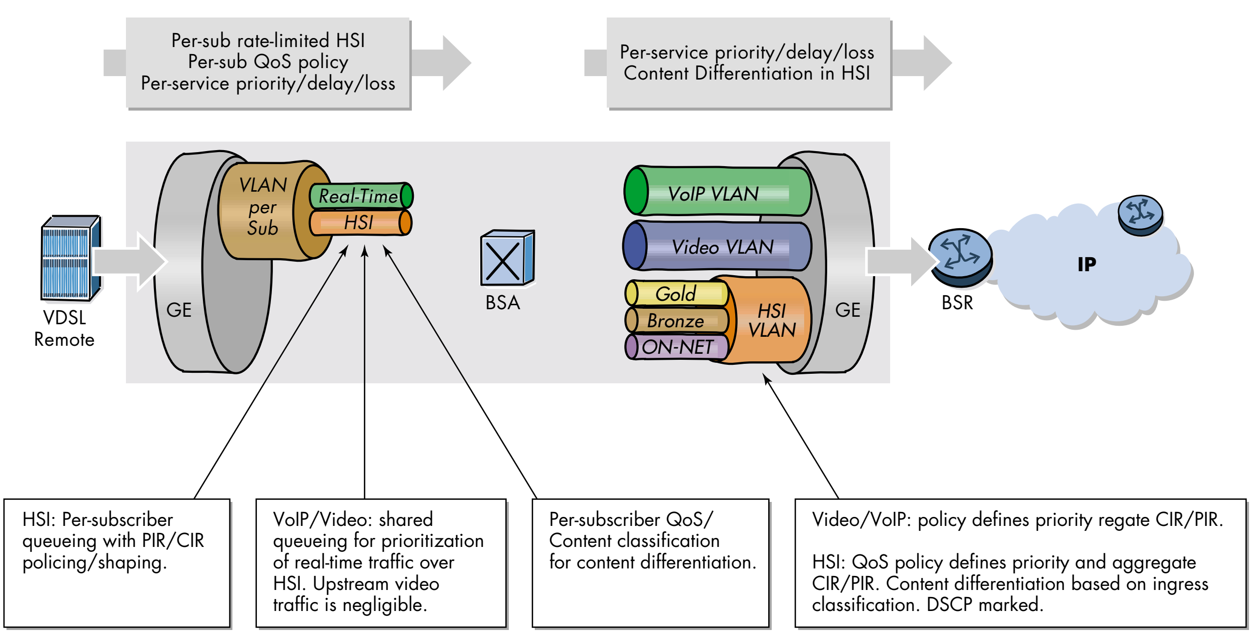

* A single virtual circuit per subscriber will be used to illustrate this section although the service delivery architecture also supports multiple virtual circuit-based models.

Distributed Multicasting

Today's predominant video service is broadcast TV, and it will likely remain so for some time. As video services are introduced, it makes sense to optimize investment by matching resources to the service model relevant at the time. Consequently, the objective of the service infrastructure should be to incorporate sufficient flexibility to optimize for broadcast TV in the short term, yet scale to support a full unicast (video-on-demand) model as video service offerings evolve.

Optimizing for broadcast TV means implementing multicast packet replication throughout the network. Multicasting improves the efficiency of the network by reducing the bandwidth and fiber needed to deliver broadcast channels to the subscriber. A multicasting node can receive a single copy of a broadcast channel and replicate it to any downstream nodes that require it, thus substantially reducing the required network resources. This efficiency becomes increasingly important closer to the subscriber. Multicast should therefore be performed at each or either of the access, aggregation, and video edge nodes.

Multicasting as close as possible to the subscriber has other benefits since it enables a large number of users to view the content concurrently. The challenges of video services are often encountered in the boundary cases, such as live sports events and breaking news, for which virtually all the subscribers may be watching just a few channels. These exceptional cases generally involve live content, which is thus true broadcast content. Multicasting throughout the network makes it possible to deliver content under these circumstances while simplifying the engineering of the network.

If PPPoE is used for video services, it greatly affects the efficiency with which broadcast TV can be delivered. With PPPoE, multicasting can only be implemented by the device terminating the PPP session, and not by intermediate nodes. Consequently, instead of multicasting throughout the network to minimize the use of fiber and bandwidth in the access network, all multicast replication must take place in the BRAS. This greatly increases the bandwidth requirements in the aggregation and access networks. An alternative approach would be to distribute many BRAS nodes in the access network. However, this raises severe cost and operational complexity issues.

Efficient multicasting requires the distribution of functions throughout the access and the aggregation network to avoid overloading the network capacity with unnecessary traffic. The service delivery architecture realizes efficient multicasting by implementing IGMP snooping in the remotes, IGMP proxying in the BSA, and multicast routing in the BSR (see Figure 5).

Figure 5 - Distributed Multicasting

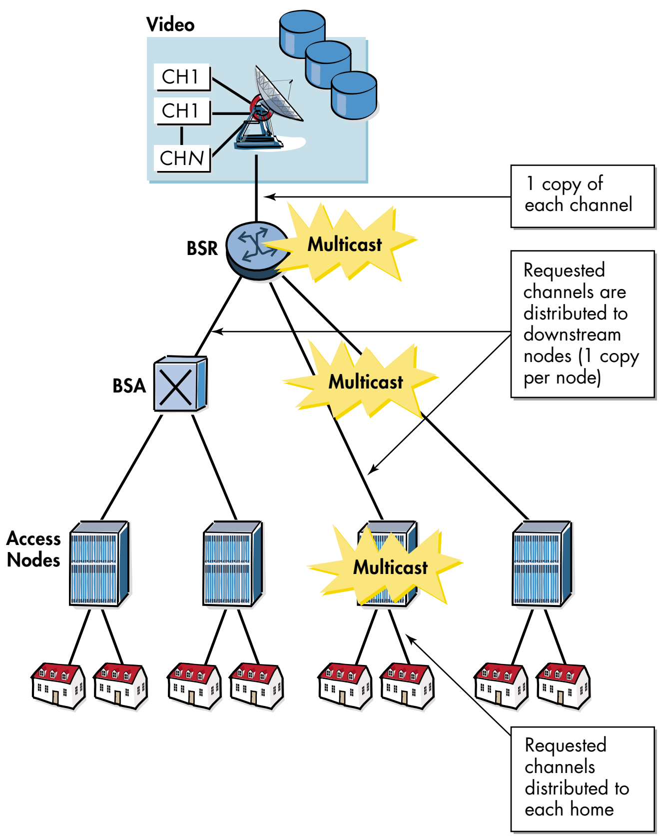

DHCP-Based Video and Voice Services

Alcatel's architecture uses DHCP for video and VoIP to provide "plug-and-play" service capabilities delivered through set-top boxes and VoIP devices optimized for that protocol. DHCP also simplifies and scales multicast video services, as indicated earlier, and increases bandwidth efficiency as a result of lower encapsulation overhead. This does not rule out the use of PPPoE for HSI. Continuing to use PPPoE for HSI, at least in the near term, will generally help minimize the impact on operations and speed up the time to market for video services (see Figure 6).

Using PPPoE for voice and video services would result in the home gateway having to support PPPoE client functionality with support for multiple PPP sessions to the network. It would also constrain multiple IP and multiple edge deployment models, which would require separate session terminations in the network.

Multiple PPP sessions are effectively equivalent to multiple circuits in the last mile with service separation taking place at the home gateway (via VLANs or virtual circuits).

Additionally, the use of PPPoE is not consistent with the use of a bridged home network, as addresses are assigned by the network via DHCP to devices behind the gateway. In other words, any use of PPPoE and PPP with bridge control protocol (BCP) adds overhead with little direct benefit. There are important differences to consider between the behavior of PPPoE and of DHCP when implementing redundancy. Unlike PPPoE, DHCP can support non-service-affecting failover without the home network being aware of the issue.

Additionally, and in the event of an edge failure or restart, all PPPoE sessions would be lost and would need to be re-established by the client with the backup edge device. PPPoE implementations generally handle these failure cases with warm redundancy models, which disrupt service. This can be exacerbated by the fact that many clients automatically attempt to re-initiate previously active sessions, thus placing a high load on the PPPoE edge platform. The processing capabilities required to react adequately to such events (for tens of thousands of subscribers) can significantly increase the cost of PPPoE-based solutions.

Figure 6 - Maintaining Present Mode of Operation for HSI

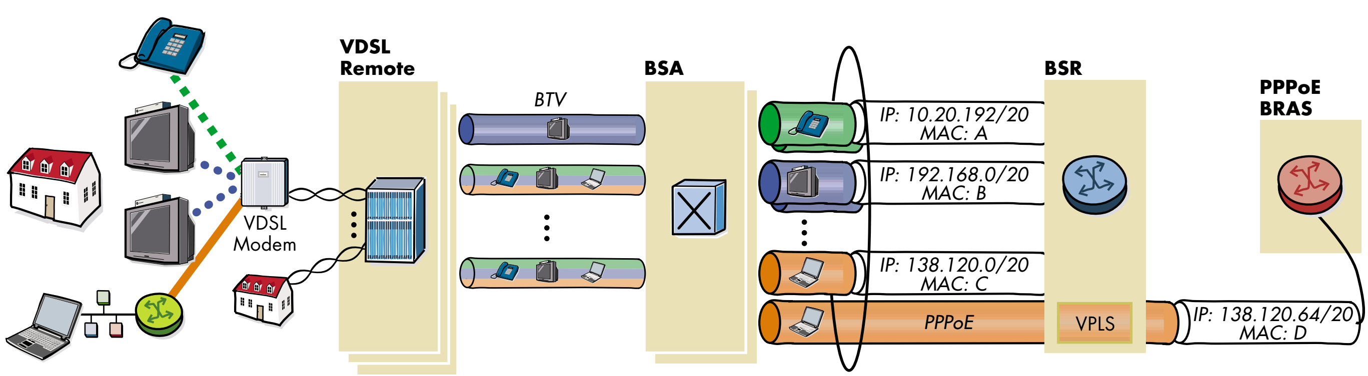

In contrast, DHCP-based services are not session-based and therefore do not suffer from these issues. In the context of the service delivery model, the virtual router redundancy protocol (VRRP) provides a proven and standard mechanism to achieve dynamic restoration (see Figure 7). In contrast to the PPPoE scenario, no device in the home network needs to be aware of a failure at the edge since the redundant node will take over from the failed node automatically and transparently.

Figure 7 - DHCP-Based Voice and Video Implementation

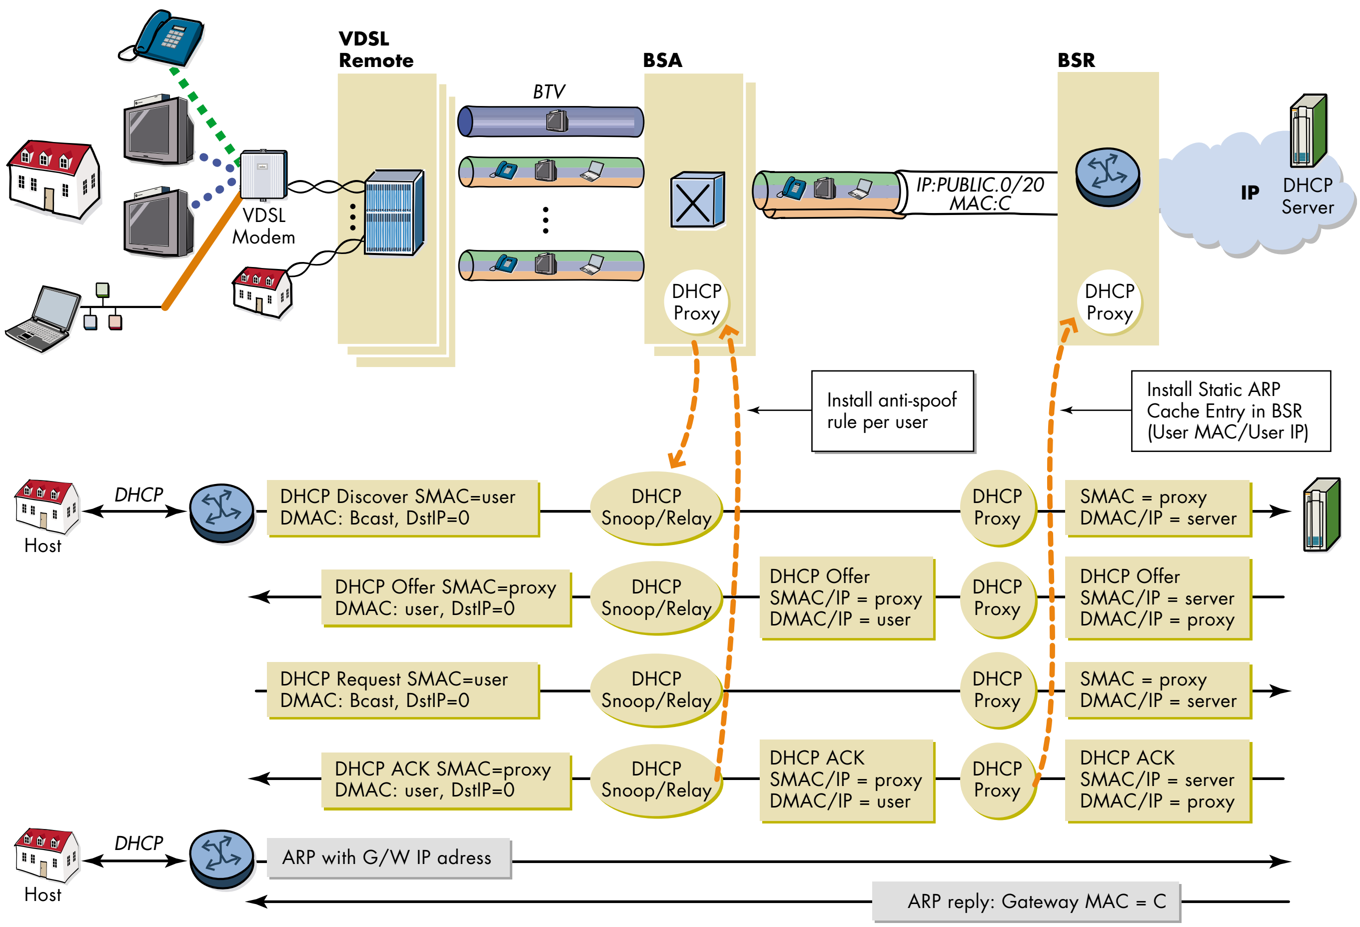

ACK Acknowledgment
ARP Address Resolution Protocol
DMAC Destination MAC Address
SMAC Source MAC Address

Implementing the Service Delivery Architecture Using the Alcatel Product Portfolio

As shown in Figure 8, Alcatel offers a complete triple play solution based on an industry-leading portfolio of products, including:

Alcatel 7450 ESS: Carrier-class Ethernet switch with a service-oriented architecture, supporting more services with greater reliability and scalability. The 7450 ESS supports hundreds of GE and 10 GE ports at wire-speed with advanced services, such as VPLS, as well as per-subscriber QoS, filtering and logging, and robust multicast support.

Alcatel 7750 SR: The industry's first service router, designed and optimized for the delivery of high-performance, highly available carrier data, voice and video services. The Alcatel 7750 SR functions as both a powerful router and a flexible service delivery platform. It is particularly suitable for triple play service deployments as a result of its ability to scale services across all dimensions without compromising performance. Features include high-performance unicast and multicast routing, VPLS support and leading QoS and traffic engineering capabilities.

Figure 8 - Alcatel's Triple Play Service Delivery Architecture

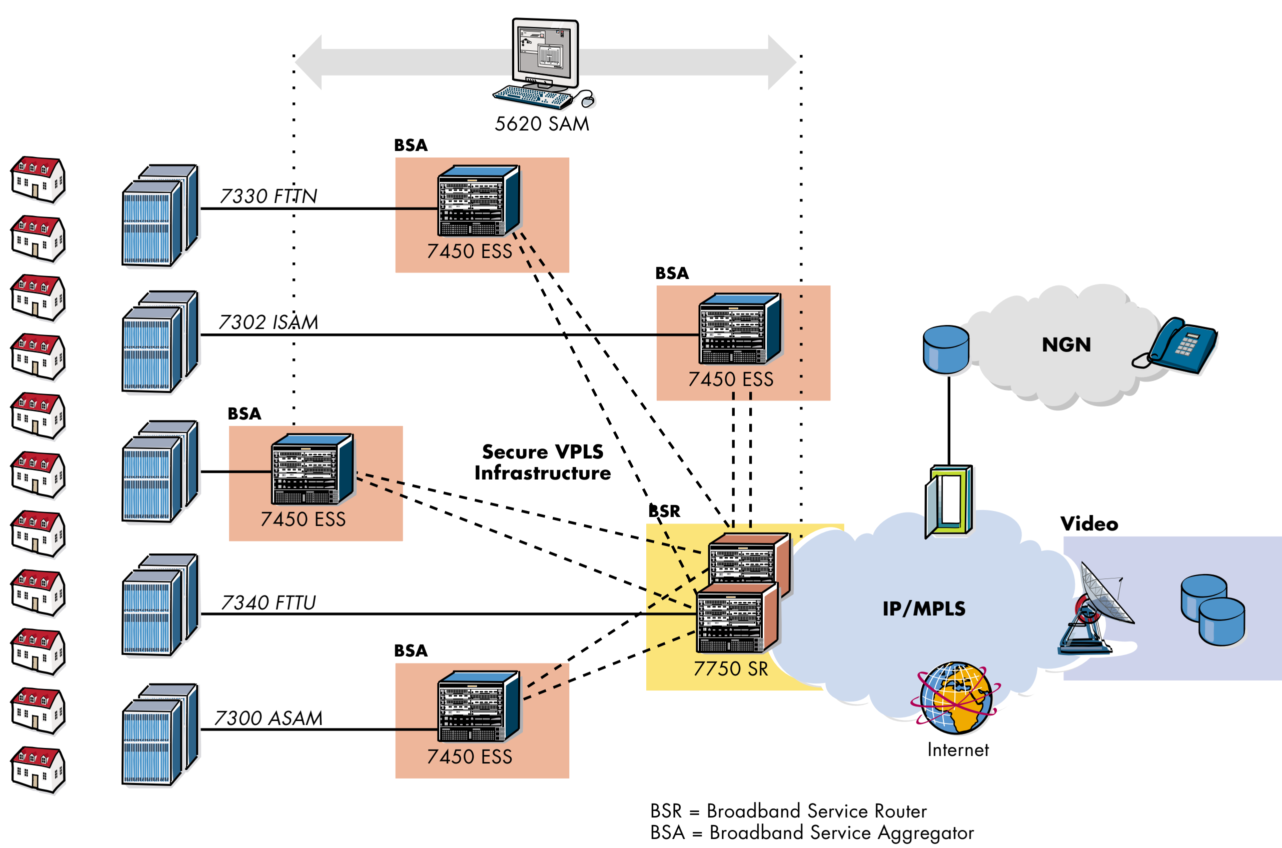

The 7450 ESS and 7750 SR, respectively, provide the BSA and BSR functionalities in Alcatel's service delivery architecture. Both are managed as a single virtual node using Alcatel's 5620 SAM, which provides a unified interface for streamlined service and policy activation across the distributed elements of the architecture, including VPLS, QoS, multicasting, security, filtering, and accounting.

Additionally, the service delivery architecture seamlessly integrates with Alcatel's leading first-mile access solutions for both fiber and copper-based rollouts:

> *Alcatel 7300 ASAM:* Industry-leading DSLAM which offers ADSL2plus, very high speed DSL (VDSL), and multiple GE network interfaces.
> *Alcatel 7302 ISAM:* The industry's first IP-based platform capable of delivering 100% triple play services.
> *Alcatel 7330 FTTN:* Small remote IP DSLAM that meets the growing need for a deep fiber access solution. It enables service providers to offer IP TV and other ultra-high bandwidth applications while leveraging the existing copper plant.
> *Alcatel 7340 FTTU:* Enables established and emerging service providers to cost-effectively deliver high-revenue, next-generation broadband services today over a single optical fiber.

Conclusion

In order to meet the combined demands of voice, video, and HSI services, providers should examine the benefits of a versatile Ethernet-centric aggregation network architecture that can adapt to service penetration growth, changing requirements, and technology choices over the next few years. Specific recommendations within the Triple Play Service Delivery Architecture are as follows:

> *Deploy an architecture that can support multiple service edges:* This enables network operators to achieve faster times-to-market for new services and retain the existing BRAS/PPPoE mode of operation for wholesale and retail HSI, while minimizing disruption to existing subscribers. It also enables service providers to leverage purpose-built platforms for voice, video, and HSI services to meet the specific requirements of each service.
> *Introduce video and VoIP services using DHCP:* This eases the introduction of services delivered through set-top boxes and VoIP phones, which were designed for use with DHCP, and to simplify and scale multicast video services, while reducing network bandwidth.
> *Continue to use PPPoE for retail and wholesale HSI:* This enables network operators to minimize disruption to existing subscribers and bring new services to market faster.
> *Optimally distribute multicasting functions across the access and aggregation nodes:* This is essential to subscriber and service scaling and reduces the bandwidth required in the aggregation network.

www.alcatel.com

3CL 00469 0778 TQZZA Ed.01 18926

ALCATEL