# THE NEW ETHERNET AGGREGATION NETWORK

February 2009

White Paper

The focus of this paper is on addressing the challenges of deployment of Ethernet as a WAN aggregation technology. Both technical and economical considerations in designing such a network are discussed. The paper recommends that the best technology to complement Ethernet with is MPLS with Traffic Engineering.

ERICSSON

TAKING YOU FORWARD

ERICSSON
TAKING YOU FORWARD

**Contents**

# 1 Introduction

The technological transition of Ethernet from a pure LAN technology to a WAN-based service aggregator has been astounding. The benefits of Ethernet are well understood: simplicity, compatibility with the enterprise networks and lower cost. Challenges remain for service providers in deploying Ethernet in the metro region to aggregate its existing and new access networks. These challenges are: choice of network architecture, multi-vendor device interoperability, migration from existing technology, minimization of network downtime, security, and on the business side, achieving the necessary economic benchmarks, e.g., ROI to justify the investment. These factors require a careful and methodical approach to building an Ethernet aggregation network.

The service provider understands that the ability to provide emerging sophisticated service, such as triple play, requires a physical change in the network infrastructure. The choice of equipment for architecting the change becomes the main driver for the economic and technical longevity of the network. There are several critical parameters in selection of the right equipment - scalability, reliability, and a clear, well-defined growth path for instance. For Ethernet to be considered as a carrier-class technology, other important factors must also be considered: OAM, fault-tolerance, and QoS.

The focus of this paper is on addressing the challenges of deployment of Ethernet as a WAN aggregation technology. Both technical and economical considerations in designing such a network are discussed. The paper recommends that the best technology to complement Ethernet with is MPLS with Traffic Engineering. This recommendation is justified since, in recent years, numerous service providers have opted for and deployed MPLS.

The paper also discusses the merits of a converged architecture at the metro edge where, a single platform can operate as an Ethernet aggregator, BRAS, and a Provider’s Edge router. The right platform will provide a technically strong and highly economical solution for enabling emerging applications, such as VoIP, Video over IP, and mission-critical business applications such as VPLS and IP-VPN.

# 2 Technical Considerations

The combination of carrier Ethernet for service aggregation and MPLS has gained tremendous market traction worldwide. According to Heavy Reading, worldwide sales of carrier Ethernet switch/routers rose an additional of 14% in the first quarter of 2006 and are poised to break the $1.2 billion mark for 2006. With sales level of $284 million worldwide in 1Q06, this sector is on pace to surpass $1.2 billion for the year – an increase of more than 70% over the $697 million posted for 2005. This trend has caused investments to shift from legacy ATM/SONET/SDH-based networks to the more scalable, cost-effective Ethernet plus MPLS. This approach is capable of supporting both enterprise and consumer applications (point to point or multipoint) such as P2P, gaming, VPNs, VoIP, VoD. However, the existing infrastructure can not summarily be discarded as it continues to be a cash-cow for the service operator, and some end-users may opt to keep their current connections. A sound strategy of supporting coexistence of legacy and new becomes essential. Therefore, the aggregation network must be able to support different access technologies: ATM, Frame Relay, FTTx, cable, wireless, TDM, as well as Ethernet.

- The key functional elements in construction of an Ethernet-based metro network capable of aggregation of multiple services are shown in Figure 1. The Ethernet aggregation switch terminates access networks (including legacy), and provides L2 user differentiation via assignment of VLANs, and bandwidth management. Figure 1 depicts systems that require tight coupling with the Ethernet switch aggregators in a metro PoP

- Broadband Remote Access Server (BRAS) that provides subscriber management (user authentication, information for billing, service authorization/policy management).

- Provider's Edge Router (PE) that provides the entry to the MPLS core network and establish end-to-end Label Switched Paths (LSPs) to carry user traffic. These platforms provide logical tunnels for user services, 2547bis, VPLS, Multicast PIM, PIM-SSM/PIM-SM

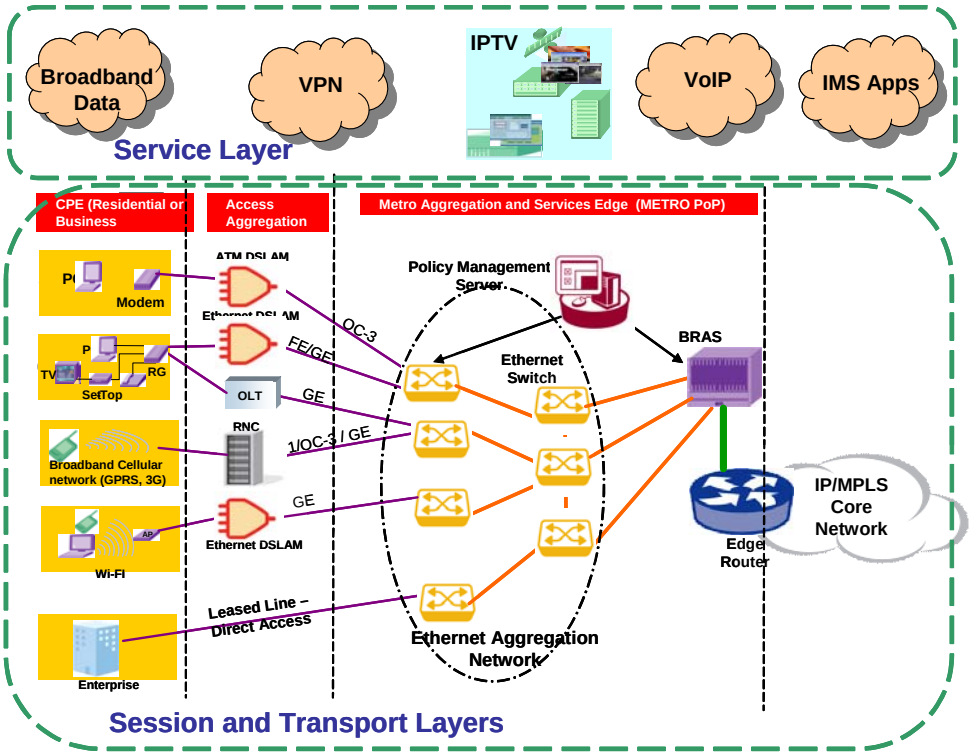


Figure 1) Functional Elements of a Broadband Services Network

Below, key factors to consider in construction of such a network are discussed. Choices in architectural approaches are also discussed. Strategies to use the network efficiently to enable high-capacity video service such as IPTV are noted. Economic advantages can be gained by consolidating these functionalities in one platform. This issue will be discussed in later section.

## 2.1 Applications and Requirements

The business choice in application and service delivery drives the underlying requirements for network construction. Emerging applications and a subset of their requirements is given in Table 1. These requirements enforce a variety of technical challenges: high-availability, scalability, speed of trouble shooting, speed of service provisioning, support for legacy access protocols and multi-vendor device interoperability. These requirements and challenges are discussed in the subsequent sections.

| Target Category | Application | Bandwidth & Throughput Requirements | Expected QoS Level | Expected Security Level |
|---|---|---|---|---|
| Business | Storage Networking | Very High | High | Very High |
| | L2/L3 VPN | Moderate | Moderate | High |
| | Video Conferencing | High | High – no jitter or delays | Moderate |
| | Live Web-casting | Moderate | Moderate | Internal – very high<br>External – Low |
| | Live Broadcast | High | Moderate – some packet loss accepted | Low |
| Consumer | Video On Demand | High for High Definition | Moderate – buffered video stream | Low |
| | IP-TV | High to very High | Moderate* | Low |
| | Interactive gaming | Low | Moderate – requires low delay | Low |
| Common | Web Browsing | Moderate | Moderate | Low |
| | VoIP | Low | High – no packet loss | Low to Moderate |
| | Email | Low | Low | Moderate to High |

* Requires the ability to change channel quickly.

Table 2) Key Requirements for New Ethernet-Based Applications

[Note: At a higher level, it is desirable for the service providers Ethernet aggregator to be able to distinguish between business and consumer traffic.]

Per Table 1, Bandwidth throughput requirements and QoS are two important factors in delivery of emerging services. Concepts in traffic management influence both of these factors, as discussed in the next section.

**Traffic Management**

Traffic management is the vehicle that is used to ensure that the network services are utilized efficiently when providing the services that a user has signed up and paid for. Services such as web browsing require very little traffic management as best-effort can be used for its delivery. In this case, the end-user can tolerate bounded delays in reception of web pages. This does not hold true for new services (see table 1 above). End-to-end QoS guarantees are required for most applications per table 1. The rule of thumb is to provide the level of service that is equal to or superior to what a customer has experienced in the past. For example, grainy, choppy TV viewing is not acceptable; therefore, no noticeable jitter or delays in the network are tolerated. Techniques such as over-subscription will only work with a limited set of users. As more users sign up for these services, more and more bandwidth must be allocated and guaranteed. Traffic management becomes an essential tool to ensure network efficiency and at the same time deliver the service that was promised to the end-user.

Traffic Management (TM) functions such as policing, shaping, scheduling, queuing, back pressure flow control are all necessary for expected delivery of services and efficient use of resources in the network. Currently, since Ethernet does not include any form of traffic management methodology, it is not considered as an optimal technology to be considered as carrier-class. Therefore, traffic management must be designed at the product level in Ethernet platforms. In general, the ATM-based QoS model is a good benchmark to define and design these functions for Ethernet networks.

To achieve this, policing and shaping can be done on the ingress and egress ports of the system. As the packet enters the Ethernet aggregator, it can be policed, via a variation of Generic Packet Rate Algorithm with the Dual Leaky Bucket approach, and marked based on the user Service Level. Using three-color marking for user traffic enables the service provider to limit the inter-arrival rate of packets and assign relative priority to the traffic for delivery of tiered, flexible services. Three color rate shaping uses Committed Information Rate (CIR) for guaranteed traffic, Peak Information Rate (PIR) for Burst Services and marks non-conforming excess traffic to be passed given there is sufficient available bandwidth or otherwise dropped (figure 2).

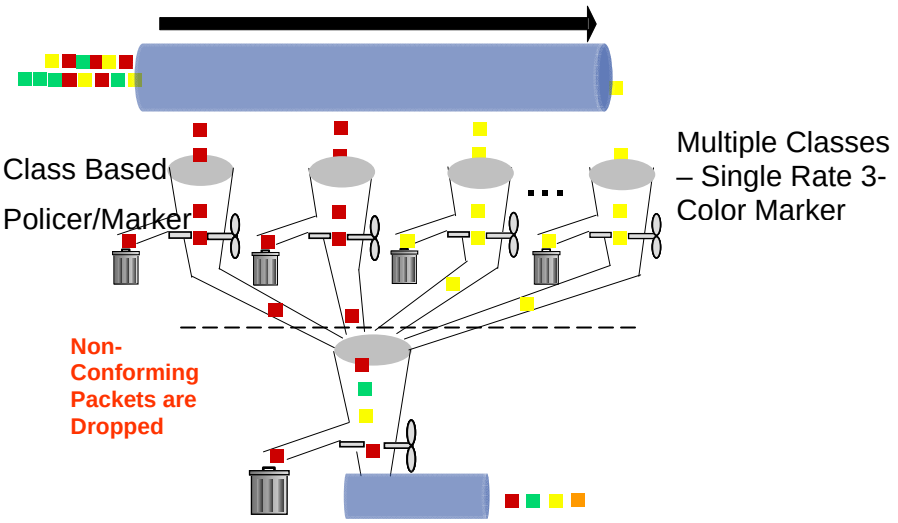


Figure 2) QoS with 3-Color Marking & Dual Leaky Bucket (Policing and Rate Limiting)

Traffic shaping ensures that network traffic is shared fairly among all users by setting per-user rule that limits traffic to or from each user to a certain level. Traffic limits can be either "hard" or "burstable." Burstable limits, or Maximum Burst Rate (MBR), allow traffic to exceed the base threshold value (at least up to a specified "burst limit") as long as bandwidth remains available and no higher priority application preemptively claims the bandwidth. In addition, shaping ensures that there is no buffer overflow in the system. And using back pressure flow, the system is capable to manage traffic

overload (exceeding the available bandwidth) that can occasionally occur at the egress (trunk side).

A technically advanced Ethernet aggregator switch is capable of utilizing a hierarchical QoS scheme to apply to different applications that a single user has subscribed to. For example, if a single VLAN is assigned for each user to deliver Video/Voice/Data, then it is technically advantageous to have an application-aware platform that uses the aforementioned scheme to ensure that each application is getting the priority that it requires. Many of currently available aggregators have implemented their QoS schemes based on layer 2 circuits. However, if the aggregator is layer 3-aware, e.g., IP-aware, utilizing L2 TM schemes are grossly inefficient. The ability to apply TM to application layer (L3 aware) can make the network operation much more efficient as all of the TM will now occur in one place where it can model the congestion points of the entire network. Indeed, higher granularity in traffic management implementation has included schemes to model congestion points, e.g., last mile link, DSLAMs, all the way to the point of aggregation (apply traffic management to each point 1, 2, 3 – see figure 3.)

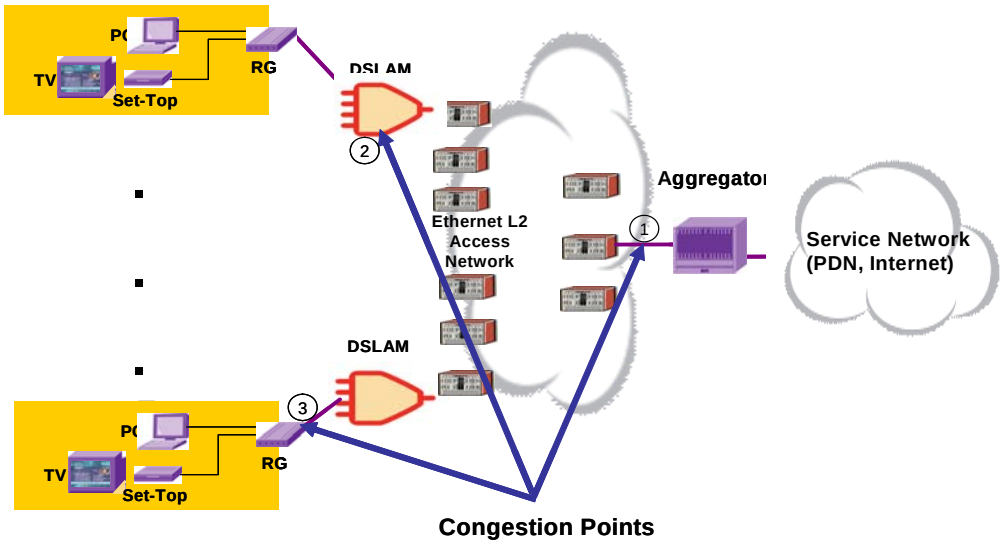


Figure 3) Application of Traffic Management Per Network Element

The aggregator in figure 3 is capable of L3 QoS application and is aware of data streams from each point of congestion. It can apply TM via “Hierarchical Scheduling” of packets per point of interest, where packets stay in queues until transmitted via some pre-defined policy. Consequently, network operation becomes highly efficient. This scheme is not the same as hierarchical queuing, as it does not increase packet latency, given that the implementation uses methods such as Priority Based Weighted Fair Queuing to transmit packets. In addition, Weight Random Early Drop scheme can be used for further congestion avoidance.

As mentioned, the Ethernet aggregator can only do this if it can recognize and delineate each application. It can be argued that application recognition should occur in the BRAS or at the PE, but the closer this function is to the edge, the more efficient the operation will be as inter-system communication is reduced and traffic management need only be done in one platform not two or three. This fact makes a strong case for a consolidate platform.

**High Availability**

Service providers require that any device that is deployed in their network (edge or core) to be carrier-class. This remains true for Ethernet aggregation switches. NEBS compliance is the minimum requirement for a carrier-class platform. High availability for both hardware and software is critical. While each service provider may have a different view of what a "carrier-class" equipment should support, the following requirements have been consistent according to a study by Infonetics Research[1]:

- Separation of Control and Data traffic
    - Most L2/L3 switch/routers currently use this approach where the line cards and the switch controller cards are separate. Line cards receive, manipulate and forward user traffic but send control traffic, i.e., protocol messages to the switch controller card, where the protocol stacks reside. Most designs use redundancy to protect the switch controller card. Upon detection of fault in the primary card, a switchover to a hot standby controller card is initiated. The choice of internal communication among the cards affects the resiliency of the platform. Most of the current platforms utilize a common backplane for inter-card communications. Strict point to point communication between each card can result in better overall resiliency as it eliminates a single point of failure in the system. This architecture is also cost-effective as it eliminates the need for redundant switch fabric cards.
- Nonstop service
    - The flow of user data traffic must be minimally disrupted when a fault occurs in the system. Specifically, problems that may be encountered in the control plane must not affect user traffic. This is critical as the service provider must adhere to its SLA with the end-users.
- In service software upgrade
    - Frequently, a vendor may introduce new fixes for problems/issues that have been found in the field, or provide new features that may have

[1] Infonetics Research, "Service Provider Plans for IP Triple Play: North America, Europe, and Asia Pacific 2006

been added to the software. In some cases a completely new software image has to be downloaded and the system restarted. In any case, where, the platform has to go through any software change (minor or major), the user traffic must not be affected. Unless the underlying operating system is capable of processing “marginal” modifications, it is very difficult to provide this feature.

- Redundant hardware (power supplies, fan trays, port redundancy)
    - The MTBF for these units must be clearly provided by the vendor. To increase availability, reduction in MTTR is necessary; therefore, spares must be available to replace the faulty units on premises.
- Remote loop back testing and the ability to quickly pinpoint faults in the network
    - As faults occur within the network, they must be quickly pinpointed and resolved. A variety of standard tools can be used to help the service provider achieve this, .e.g., ping and trace, diverting traffic to a designated port in the network.
- Software reliability
    - This has proven to be even more important than hardware reliability. While both are necessary, without the right operating system it becomes difficult to provide the level of system reliability that the service provider demands. The operating system’s modularity is essential to separate different tasks, e.g., BGP, OSPF, RSVP, LDP, in such a way that faults in one does not become systemic and permeate to other tasks. Each task must have its own independent process space and protected memory area. Memory leakage must be avoided to ensure that adequate resources are available at all times [with protected memory, the impact of such a problem is strictly limited]. Additionally, a modular operating system should enable “In Service Software Upgrade” so when a bug-fix is applied to the BGP process and the process is restarted; it does not affect the other processes or force a system re-booting. Unfortunately, many of the current generations of routing platforms which are based on large monolithic operating systems and their architecture can not be considered as carrier-class. In these systems, instabilities in any of the processes can cause memory leakage or even worse, crash the system. Furthermore, these monolithic operating systems often lack proper task and process scheduling capabilities, which impact the overall system operation when under high load, caused, for instance, by a network event.

- Network Reliability
    - The network itself must be able to survive faults. Standard based schemes such as VRRP for dual platform redundancy (active/standby), Link Aggregation (802.3ad) for link protection or Label Switched Paths (LSP) protection via backup, nailed-down LSPs or Fast ReRoute (FRR) are required to maintain non-disruption of user services. By utilizing FRR, service providers can achieve SONET-like protection in order of milliseconds in path switchover and restoration. Another essential requirement is protection of routing and signaling protocols via stateful redundancies within the platform and standard-based graceful restart methods, e.g., for OSPF, BGP, LDP.

Several other factors can substantially influence the operation, growth, and competitiveness of the service provider's network. Identifying the target market, expected market penetration rate (number of expected subscribers per application) for the next "**n**" years influence the network design, and require solutions with the ability to scale and in-step feature growth schedule. The speed with which the service provider can establish services for its customers, pinpoint faults and defend against external threats to its network directly influence its level of competitiveness. Below, each of these is discussed.

**Network Topology**

The choice for architecting the network topology depends on the required size of the network to accommodate the type of services to be offered to the expected user population, the fiber plant and layout. One or two Ethernet switches may be adequate to support 10,000 users, whereas multiple number of these would be necessary in a metro region with a population of >1,000,000. The choices can be broadly be categorized into one of: 1. Star (Hub and Spoke), 2. Ring, and 3. A "Hybrid" of the first two. Figures 4 & 5 depict these topologies. Below, each topology is discussed and its deployment scenarios given.

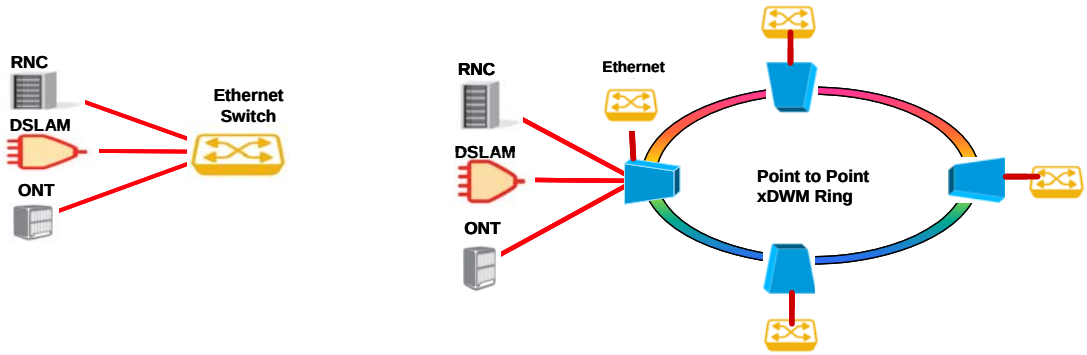


Figure 4) Star and Ring Topologies

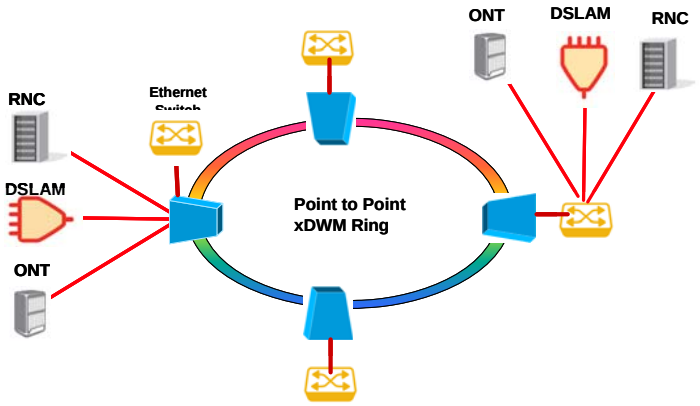


Figure 5) Hybrid Topology

### Hub and Spoke

A hub and spoke topology can be used when the service provider is providing services to a limited set of users within a geographical region. A single aggregation Ethernet switch can be placed in service provider's regional PoP to receive incoming flows from multiple DSLAMs (both ATM and Ethernet connectivity). Since dedicated links are used between the DSLAMs and the aggregator, bandwidth is not shared among the DSLAMs but scalability of the aggregator becomes an important factor in selecting it as the number of required physical ports is the same as the number of DSLAMs. The dedicated physical links may increase the cost of installation as each connection requires its own physical fiber. See figure 4. Protection of a link from an access device (e.g., DSLAM) can be provided via Link Aggregation 802.3ad.

### Ring

The ring topology is used in a larger metro area where the geographical service coverage demands topological efficiencies. It is more scalable than a hub and spoke but its potential drawback is the fact that it is a share media - bandwidth is shared among the aggregators. The shared characteristic of the ring topology translates into a lower required number of physical ports for the aggregation node. Expected traffic volume based on the user population and service bandwidth requirements must be considered in selection of the ring bandwidth. With MPLS, using traffic management schemes, e.g., RSVP-TE, efficient use of the ring's bandwidth can be ensured.
The ring topology is more prone to problems, e.g., fiber cuts, and proper protection is required to ensure minimal network disruption. Assuming that the Ethernet switches are MPLS capable, fast reroute can be used to bypass the cut in the network as shown in figure 5. An LSP can be established in the opposite direction to restore the failed connection.

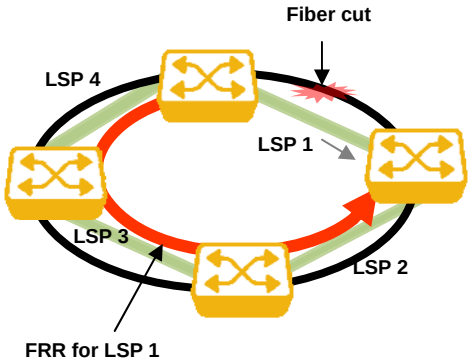


Figure 6) Ring Protection via FRR LSP

### Hybrid

The Hybrid topology is a mix of Hub and spoke and Ring topologies. One or more Ethernet Aggregators that are placed on the ring may have direct connections to access devices through additional physical ports or connectivity to other standalone aggregators that are not placed directly on the ring itself. Figure 6 depicts this architecture, where ONTs, DSLAMs or RNCs are directly feeding user data on to the Ethernet switch aggregator that is also part of a ring itself.

### Scalability

Expected market penetration and the corresponding services to be delivered are important parameters in determining the level of required scalability for the target Ethernet Aggregation platform. For example, delivery of bundled services to consumers in a large city (e.g., greater than a 1000,000 in population) will require platforms that are highly scaleable otherwise the size of the network can become too large to manage it effectively. Consider when a single (or even multiple) VLAN(s) is assigned per user to deliver triple play services to a city of this size, the service provider can ill-afford to use a system that is bound to a maximum of 4096 VLANs. Even with Q-in-Q support, most service provider use the outer Q (S-VLAN tag) for their own circuit differentiation and not for user assignment, therefore, they are still limited to 4096 VLAN assignments per system. Support for assignment of tens of thousands of VLANs is essential in the new generation of intelligent, scalable Ethernet aggregators.

For delivery of emerging applications, e.g., IP-TV, the scalability of MPLS-based services becomes important. For example, if an L2 based (e.g., VPLS), approach is used for delivery of this service, scaling of VPLS instances is necessary. Each VPLS instance can be used to support a different set of users within one geographical region. Also, since tens of thousands of VLANs may be mapped into a single VPLS instance, MAC address scaling becomes an important issue as thousands have to be learnt and stored.

Using a L3 approach for delivery of IP-TV, e.g., PIM, or PIM-SM, support for a large number (upward of 100s of thousands) of users and groups become critical. Furthermore, the video head-end systems; such as the MPEG-encoders, Video-on-Demand servers, Emergency-Alert-Systems, Middle-ware servers all have Ethernet (FE/GE) connectivity requirements increasing the need for scalability.

To be able to support the above scaling, several key factors in selection of an Ethernet aggregator must be considered: switching capacity, number of line cards/physical interfaces, internal memory architecture and software architecture.

**Operation and Maintenance – OAM**

OAM functions can be summarized as one of configuration, fault-detection and statistics collection. This topic deserves its own white paper. Important issues are covered in this section for reference.

Efficient OAM capabilities are gauged through fast provisioning time, quicker fault-detection and correction and availability of ample statistics information to make adjustments in the network and gauge its performance.

Quick provisioning provides efficiency in the use of service provider's assets and increase customer satisfaction. Revenues can be collected faster as a result. In general, the speed of provisioning pseudo-wires to support ELL or VLL must be taken into account. For example, according to a study by Ovum[2] provisioning strictly the VPLS circuits for a 5-site VPLS service took 10 to 350 minutes depending on the architecture and the platform used. Provisioning VPLS within the core network itself took an average of 25 minutes for establishing 5 VPLS circuits. These are not benchmark figures, but a good indication for the current speed of service delivery.

For the service provider, access to network statistic is critical to ensure that user SLAs are verified. Statistics that provide information on packet latency, loss, jitter, and throughput are necessary for gauging the performance of triple play services. The Ethernet aggregator needs to provide these statistics on the down to the packet level. After collection of the information, it can be analyzed against the benchmark levels that were set by the service provider. Adjustments to the network may be prescribed after this step.

In a network of Ethernet Aggregator, a strong capability to detect, pinpoint and correct faults can minimize network downtime, and ensure adherence to customer's SLA. Customer contracts usually have clauses that penalize the service provider when their circuits are lost. [The customer does not care why a fault happened!] A variety of tools are available to pinpoint faults in the network, e.g., ping and trace at the LSP and MAC levels, and Bi-directional Forwarding Detection (BFD) to monitor

[2] Ovum, "VPLS growth requires carrier-class network features", October 2005

protocol (e.g., OSPF) operation. In an MPLS based network, protection of LSP, e.g., through FRR, is essential in bypassing faulty segments.

**Security**

Network security is a major issue for all the players - service providers, vendors and end-users. Attackers have been getting increasing sophisticated in their schemes in disrupting the network and its service. The attackers usually do it by trying to exploit network weaknesses to extract sensitive user information. For an Ethernet aggregation platform, stringent security policies must be considered as it is the first point of entry to the service provider's network. This includes security for L2 and L3 protocols, as well as application specific security support. The Ethernet aggregation platform must ensure defense against a malicious attacker to its control plane, sensitive internal tables, e.g., MAC address tables. In addition, it has to ensure that it does not compromise the network operation by undue delivery of disruptive packets, e.g. broadcast storms.

The main goal is to anticipate, identify and defend against a set of well-known and unknown attacks. Several major categories to consider are:

- **(Distributed) Denial of Service**
    - Flooding [Flooding excessive packets to cause the network to behave abnormally and operate in a degraded mode.] Consider a scenario where an infected IPTV set-top box (STB) continually transmits IGMP join/leave messages to the network emulating a scenario where viewers are constantly switching channels. This can be categorized as DDOS attack if multiple STBs are infected and are sending the same messages. Placing limits on such messages can prevent such a problem.
    - Duplicate MAC and other sensitive tables [Ability to disable learning on a specific logical interface. Placing a limit on maximum number of MAC addresses that can be configured on a logical interface, to control the number of entries in the table and avoid overflowing it with bad addresses.]
- **Address spoofing** [An attacker tries to redirect traffic to a destination of his own choosing by forging a MAC address. This way the attacker can extract and disseminate sensitive user information from packets].
- **Frame Tagging** [Transmission of frames with invalid MAC addresses such as incorrect headers or bad contents. This type of attack can overflow the internal tables. Therefore, the platform must discard any packets with incorrect headers, e.g., a MAC address that contains all "0"s or a MAC address with the Multicast bit is set in the source MAC address of a packet.]

- **VLAN Hopping** [In this case, packets are intercepted or redirected from the originating VLAN to either a legitimate or bogus VLAN. A method to defend against this is to ensure that the VLAN is provisioned by the service provider who has a secure login for network configuration. Also, unknown VLAN IDs must be automatically dropped by the system]

- **Anti Replay** [The attacker repeatedly transmits a packet or packets by traversing and storing them. A sound security solution must be able to detect replayed packets and discard them].

At application levels, techniques such as “stateful packet inspection firewalls” can ensure that packets are intercepted and filtered at the network layer. Stateful packet inspection firewalls analyze packets in terms of sessions by examining all incoming data transmission and if a packet happens to be a correct reply to a previous request from within the network, the firewall allows it to go through. Otherwise, access is denied. With stateful packet inspection firewall, individual ports and connections are tracked as well so that no ports or connection need to be operational other than the required ones based on preset security policies.

Other techniques such as Lawful Intercept which allows for the inspection of traffic to the full extent of law, and Internet Key Exchange (IKE), which ensures security for VPNs negotiation and remote host or network access, can be used to further protect the network and its services.

## 2.2 An Efficient Network Architecture

In previous sections, discussion was focused on the major requirements and considerations for selection of an Ethernet aggregation platform that would be capable of delivery of new multimedia services. Topics such as traffic management, network provisioning and security were discussed.

Further economic and technical efficiencies can be gained if the Ethernet aggregator platform is integrated with the other network devices. Based on figure 1, the natural candidates for integration are the provider edge router (PE) and the Broadband Remote Access Server (BRAS). The industry refers to these platforms as Multi-Service Edge Routers (MSRE). A few vendors in the industry are offering platforms that integrate two or all of these functions: Ethernet aggregator with that of a providers’ edge MPLS router and/or BRAS. The majority of vendors offer a two-platform solution to provide all three functions. Figure 7 shows the architectural simplification that can result from a consolidated platform relative to the network in figure 1.

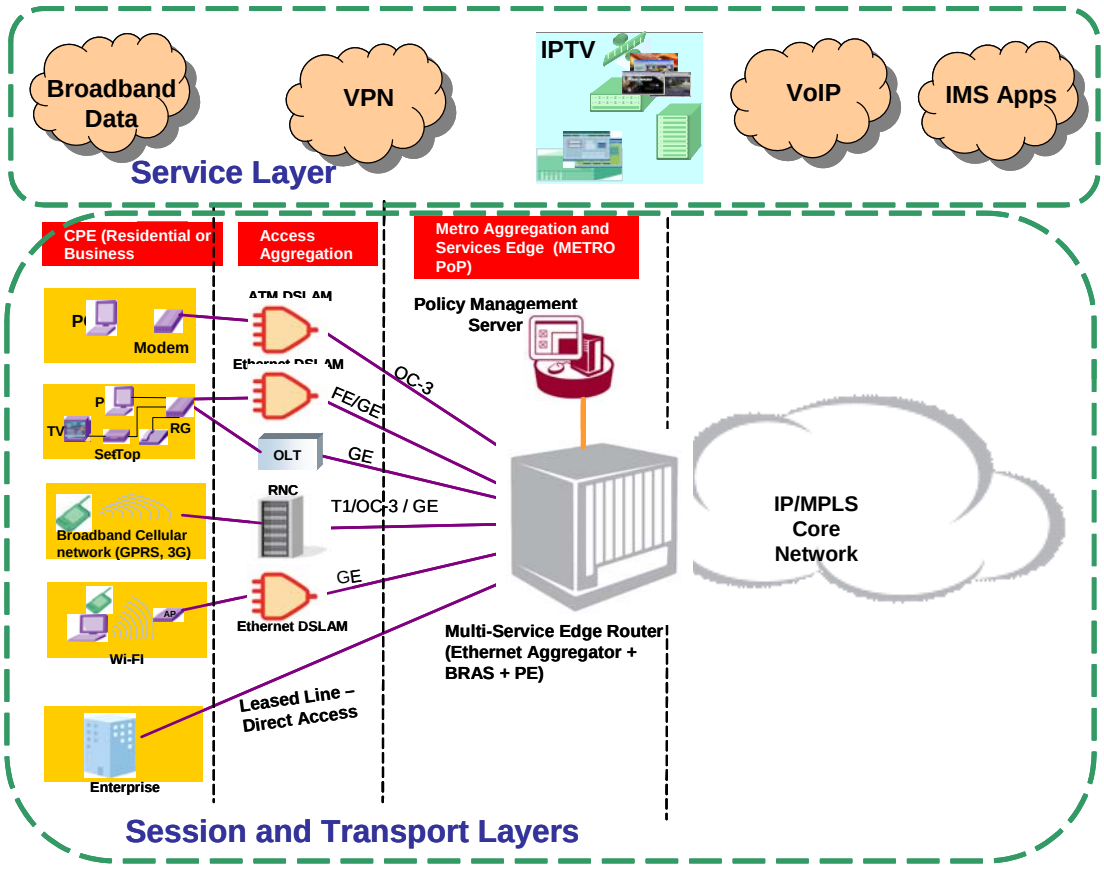


Figure 7) Multi-Service Edge Router – Consolidated Metro Aggregation and Service Edge

In addition to simplifying the network architecture, a consolidated platform can reduce latency and increase performance by eliminating the need for inter-system communication in a 3-platform architecture. Signaling and routing messages that must be exchanged between an aggregator and a PE would also be eliminated. Common software, tables, or hardware modules can be shared among different the three “functions” [Ethernet aggregation, BRAS, PE routing] for further system efficiency. Common security features in one consolidated platform can reduce network susceptibility to attacks and failures since the platform will be able to defend itself from a variety of attacks without having to “export” problems to other devices that may crash them. For example, since a central platform can now validate the source address of all incoming packets, the probability of attacks can be reduced as the source of the attack can be readily identified via PPPoE or DHCP traffic that comes into the consolidate platform. If the fault can be identified and taken care of close to the access, it will not be allowed to travel to edge of the core of the network where problems may magnify the initial attack.

Further performance is gained in such a platform as traffic management can be applied in one place which can provide a higher throughput for user data. If both the

aggregator and the PE conduct traffic management for the same stream, this can increase latency and reduce speed of service delivery. Common H-QoS methods can be applied directly to the application packets once instead of a two tiered QoS processing of single stream – once in the Aggregator or the BRAS and once in the PE.

Integrating subscriber management functions such as authentication, application session establishment, and billing into this platform can provide further resource efficiencies. For example, a single VLAN assignment can uniquely identify an end user via DHCP (option 82). This way no interface identification is required for uniquely distinguishing an end-user, as needed for DHCP.

Finally, network management of a single consolidate platform will become simpler. Instead of having potentially three element management systems plus a policy manager, all can be consolidated in to one central platform that can provide policy management, configuration, provide statistics and pinpoint faults.

# 3 Economics Considerations

While delivery of services by deployment of Ethernet-based access has gained noticeable momentum, the economic fundamentals have not changed. End-users do not care what access technology is used as long as they get the same services at reasonable prices or more services that they consider useful/interesting at marginally higher prices. To loosen resistance to change, initial incentives are necessary to capture market share. As competition increases in the consumer service delivery segment, the business model of providing bundled services [VoIP, Broadband Internet and TV (VoD + SDTV + HDTV)] becomes more prevalent. ARPU for this bundle is expected to be in the range of $100 to $130 per month at equilibrium. For delivery of business services, the average revenue per customer is substantially more but customer services may offset some of the gains, e.g., 24x7 monitoring of service. In general, the cost of an Ethernet VLL, ELL, VPLS should be less than the existing leased lines, FR, or ATM and must provide tangible service improvements for an enterprise to shift its connectivity.

To offer emerging services via an Ethernet network, the service provider is faced with a couple of economic forces: 1. price elasticity of demand and 2. performance elasticity of demand. In other words, a noticeable change in price levels or performance levels can decrease (or increase) the number of interested end-users to sign up for the new services – businesses or consumers. For example, if a large difference exists between an Ethernet based access solution (performance or price) vs. an ATM-based access solution, the number of consumers can get lower (or higher) from one to the other service. Combine these factors with the reality that the

price expectations are low for Ethernet services, providers must look to two parameters to make a profitable business case: *economies of scale* and *product differentiation*.

## 3.1 Economies of Scale

When selecting a platform to offer emerging services, studies have shown that there is a clear economic advantage to service provider in selecting a platform that consolidates that functionality of an Ethernet Aggregator, Edge Routing and subscriber management (BRAS)[3]. This study shows a 22% advantage in TCO, with a 39.8% IRR and an NPV advantage of 51.8%. An integrated platform can provide the economies of scale as each incremental resource, e.g., a single port, needed to service an additional set of customers can provide functionality for all three sub-platforms. In other words, increasing geographical reach and therefore, the number of customers will not require an exponential increase in purchasing additional resources [$O(\mathbf{n})$ in new resources as opposed to $O(\mathbf{n}^2)$]. A consolidated platform, therefore, has a major impact in achieving economies of scale.

The target integrated platform must still be carrier-class, scalable, and allow for efficient use of resources, through application-aware H-QoS and traffic engineering. Scalability is of specific interest since its lack can greatly increase the cost of expansion. Using the scenarios provided in the Yankee Group study, service delivery for a set of customers starting with pool of 100,000 and increasing to 700,000 in a span of four years, let us make the following assumptions:

1. All customers have signed up for triple play services
2. Each customers gets on VLAN assignment for delivery of said services

Based on point 2 above, the service provider must be able to assign 100,000 VLANs (or PPP/DHCP sessions) in the first year of operation. The minimum number of VLANs that are supported in an entry level Ethernet aggregator is 4096 (even with Q-in-Q, assume S-VLAN has local significance to the network). Therefore, a total of 25 switches will be needed to architect the network. With a platform capable of providing 64,000 assignable VLANs, only two are needed to satisfy a pool of 100,000 customers. As the penetration rate increases and the subscriber pool grows to 700,000, only 11 platforms are needed. Compare that with an entry switch that only scales to 4096 VLANs and 700 of them will be needed to satisfy a pool of customers of this size! Even with a 3:1 oversubscription, the number decreases to 234 for entry switches vs. 4 for a highly scalable one.

Let us further assume that on the average each customer has the following services:

[3] Yankee Group, “Integration of Edge Routing, Ethernet Aggregation and Subscriber Management Yields TCO advantages”, 2006

- Two Standard TV sets
- One High Definition TV set
- Three VoIP lines
- High Speed Internet access at 5Mbps

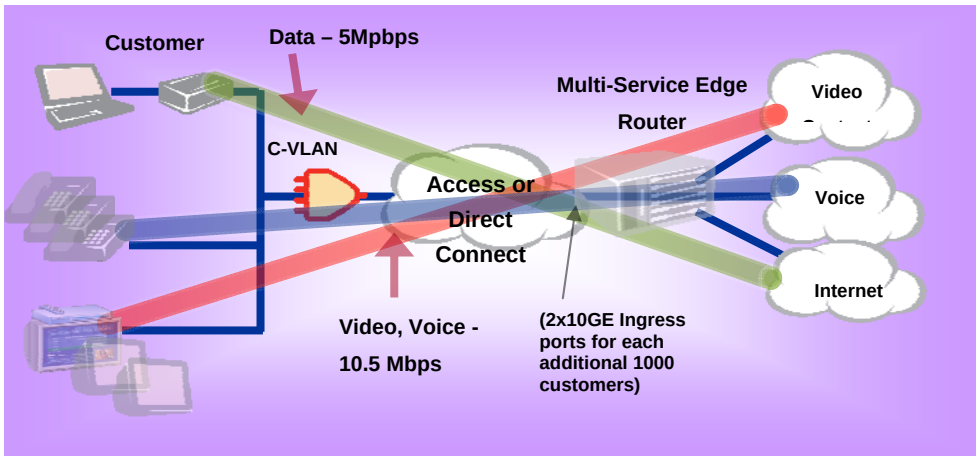


Figure 8) Bandwidth and Port Requirements in a Typical Multimedia Scenario

Therefore the total bandwidth requirement per customer will be about 15.5Mbps [7.5Mpbs for HD-TV+(2*1.5Mbps for SD-TV)+5Mbps for Data]. VoIP traffic compressed per G723.1 to 6Kbps requires negligible bandwidth. Therefore, for each addition 100 customers, two new GE ports will be required to accommodate the new pool. Or alternatively for each additional 1000 customers 2 x 10GE ports will be required (Figure 8). Having a non-consolidated based network, each additional 1000 customers will require 6 x 10GE ports (2 ingress to the aggregator and 4 to interconnect it to the PE). It is easy to see how the cost can grow geometrically in this case.

## 3.2 Product Differentiation

Service provider's product differentiation will come in the form of its ability to provide granular, verifiable SLAs to its customers. The integrated platform must be able to provide differentiation with QoS per application. For example, if the customer requires that its traffic burst beyond the Committed Information Rate (CIR) for an extended period of time and receive proactive congestion management, it should be able to do so. The service provider must be able to allocate excess network capacity on a fair basis among all users and provide optimum protection of network resources by metering access to the network.

# 4 Case Study

A European service provider, a leader in delivery of Ethernet services to corporate and wholesale customers, had seen a 44% increase in its enterprise customers from 2004 to 2005. Its main products are VPLS, Ethernet point-to-point services within metro regions and beyond. Ethernet services are offered at speeds of 10Mbps to 1Gbps, with an average of 100Mbps. L3 based applications are also offered with VoIP and IP VPN. The service provider's core network is based on IP/MPLS which interconnects multiple regional networks.

## 4.1 The Issues

The service provider did not have its own xDSL network, and was forced to lease DSL service from its competitors. This was expensive and the lease was restrictive – no ADSL2+ or VDSL. This fact was threatening to stunt the service provider's business growth as it could not compete effectively in the residential market.

The service provider did not have L2 Ethernet transport for L2 services as it was employing FR, ATM and SDH to provide access to enterprise customers. These technologies were not scalable and were becoming expensive relative to Ethernet.

Because of the access technologies that the service provider was using, it was unable to offer services that require multicast capability, e.g., IP-TV, and this deficiency was beginning to erode the service provider's market share.

## 4.2 Initial Plan, Incumbent's Solution and Shortcomings

The service provider decided to aggressively expand in 2005-2006 upgrading the capacity of its backbone network with new 10GE routers (both in the core and at the

edges). The edge routers were to act as hubs for VPLS instances. Additional revamping was to replace the current, performance-deficient PE routers. It also decided to deploy its own next generation IP DSLAM network to provide unbundled DSL services. The upgrade plan promised a sustained growth rate as the service provider would expand its markets in large enterprises, small to medium sized businesses and residential segment. The service provider decided to use its existing fiber layout to provide Ethernet on the access side.

After assessing the feature/functionality requirements of their incumbent vendors, they decided that their solution had the following shortcomings:

- **Limitation of the incumbent Ethernet aggregation devices for L2 services**
    - Lack of 10GE support
    - Inadequate QoS capabilities – no support for Hierarchical QoS
    - Roadmap was not clear on support for VPLS
- **For VPN (2547), the incumbent lacked**
    - Bandwidth scalability
    - QoS deficiencies with increasing subscriber base
- **The Core Network was bounded to**
    - STM-1/4 and no growth path to STM-16 or STM-64
    - Overall bandwidth deficiency

In addition, the architecture that was proposed by the incumbent vendors became too complicated as there were too many devices to install and manage. The proposed architecture in each PoP is shown in figure 9. Deployment of new Ethernet aggregators, BRAS and PEs was proposed. This architecture required the use of multiple management systems and multiple ports to interconnect each distinct device. The service provider had to pay for these ports but could gain no revenue from them. Scalability in this architecture was bound by the least common denominator as two of the platforms were from one vendor and the third from a second vendor.

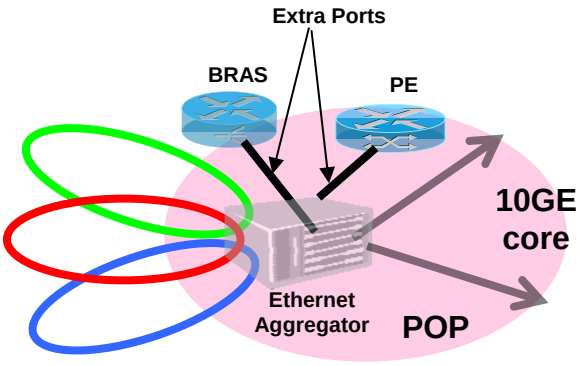


Figure 9) Incumbents' Recommended Architecture

The service provider decided to look for solutions that could revamp its core with higher bandwidth scaleable equipment, and reduce the number of platforms at the edge and at the same time, keeping some of the existing regional networks to minimize disruption in services for the existing customer.

## 4.3 Ericsson Winning Proposal

Ericsson proposed a simple, consolidated architecture (figure 10) that addressed all of the service provider's problems in one platform: the SmartEdge 800 Multi-Service Edge Router (MSER). The SE800 is a highly scalable consolidated platform that offers the functionality of a BRAS plus an Ethernet Aggregator and a PE router. It is a carrier-class platform that uses a highly available modular operating system with In Service System Upgrade to minimize network downtime and service disruption. Scalability is prevalent in this platform offering tens of thousands of VLANs, PPP sessions, and DHCP sessions. It solved all of the deficiencies that were present in the incumbent's product as noted before.

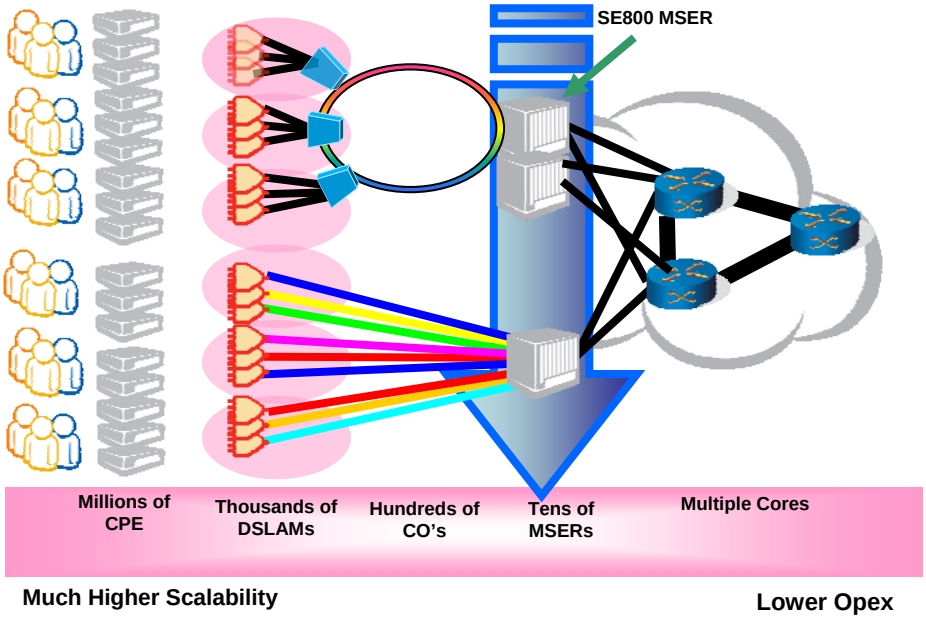


Figure 10) Simple, Elegant Architecture with Ericsson's SE800 MSER

Each PoP now consisted of two SE800s (for redundancy purposes). Note that the service provider opted to deploy one SE800 in certain PoPs. The SE800s were directly introduced in to the existing customer's fiber ring. The new IP-DSLAMs would be connected via this ring through DWDMs. One Lambda was set for Primary and another Lambda as Secondary or backup. In some areas with lower traffic, the IP-DSLAM had direct connections to the SE800. In essence the "hybrid" topology that was discussed earlier was proposed.

Certain regions that were using the incumbent's existing Ethernet aggregation platform were also integrated into to the SE800 network. Using H-VPLS to interconnect the two networks provide scaling in the existing access networks as shown in figure 11.

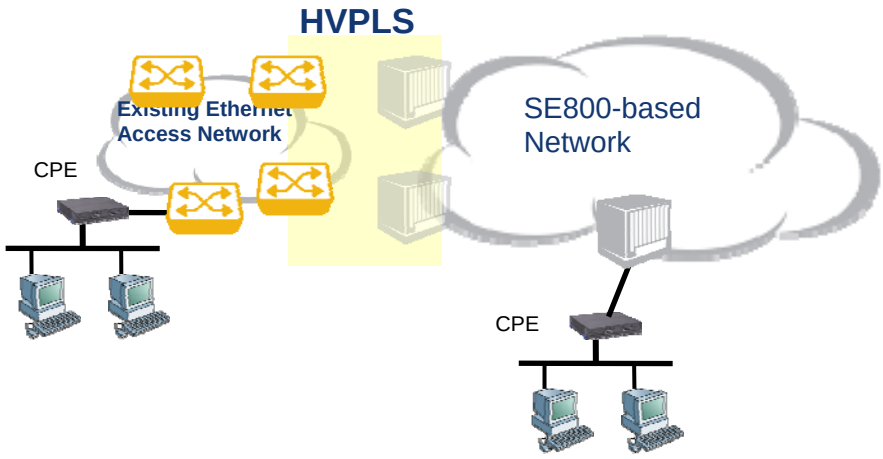


Figure 11) H-VPLS for Mixed Networks

## 4.4 Immediate Benefits of an Ericsson Solution

The economic benefits of the SE800 MSER followed closely to the discussion that was given in previous section. The following are some of the technical advantages that the service provider immediately benefits from:

- With a single platform, the management of the network becomes simple as services can be provisioned, established and monitored faster. The SE800 can provide per subscriber policies, per subscriber accounting, and consistent subscriber profiles across any access method. The service provider has the option to introduce new access methods as it sees fit in the future.

- High availability was a very important factor for the service provider. With In-Service System Upgrade, full hardware redundancy, modular operating system and extensive provisions for network resiliency, the service provider owns a carrier-class operation per its stringent requirements.

- Application-aware H-QoS scheme and the sophisticated traffic management in SE800 MSER provide a high performance network that uses bandwidth resources in a very efficient manner. For example, if VPLS services are used for data and VoIP services, the SE800 fairly assigns high priority to VoIP flows and lower priority to data services within the same incoming C-VLAN, utilizing Priority based WFQ.

- By using MLPPP and applying QoS on bundled PPP sessions, the service provider is able to increase and assign high bandwidth to its customers for delivery of additional services.

- The service provider can use the Q-n-Q capability of the SE800 to map and multiplex each S-VLAN to be used for VPLS or L3VPN (2547bis) simultaneously. This feature saves resource consumption of VLAN assignment by 50% for the service provider. See figure 12 below. For example VPLS, L3VPN or BRAS traffic can traverse the same physical port and provide management, VoIP, or Internet access for the same VLAN.

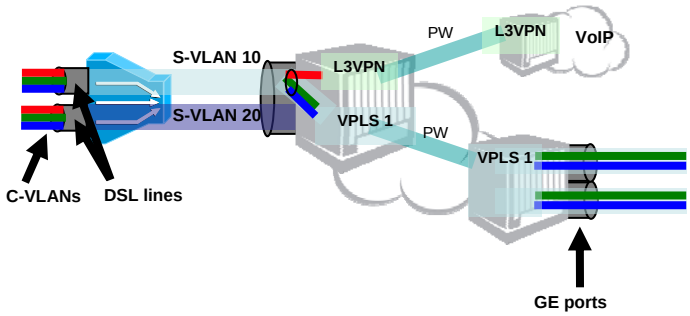


Figure 12) Multiplexing L2 and L3 VPNs on the same VLAN

Furthermore, the SE800 based network enables the service provider to enter into a new lucrative business of wholesaling bandwidth to major retail chains that provide residential services, such as IPTV. The service provider manages the entire network as its customers do not have the network infrastructure or the labor force. This is a win-win scenario for both businesses as each is focused on its own strengths. Indeed, the SE800 became an ideal platform for the blueprint to construct a highly advanced, service-rich network.

# 5 Conclusion

Service providers have accepted the benefits of deployment of Ethernet as an aggregation point for access networks. The service provider needs to carefully assess its present and near-future applications and services that it intends to deliver to its market. Depending on these services, different architectural and topological approaches are available, e.g., Star, Ring, Hybrid utilizing multiple platforms or one consolidated system. Carrier-class requirement of the platforms must be stringent to avoid service disruption and revenue streams. Efficient use of network resources must be carefully be taken in to account as multimedia applications are bandwidth thirsty. The solution must be scaleable both technically and economically.

Initial deployments of these networks focused on point-to-point and hub-and-spoke architectures, since services, such as corporate LANs and VPNs, were the primary

applications that the service providers required to maintain. But more recently support for multipoint Ethernet, and VPLS in particular, have been available. The expansion of service provider's business due to the emergence of triple play services has made the need for deployment of metro Ethernet aggregation networks more prevalent. These networks must include features such as multicast that make delivery of multimedia applications possible.

This paper showed that an architecture that uses a platform with consolidated functionalities is preferable to a multi-platform one. A typical case study that demonstrated technical complexities that currently hamper the deployment of these aggregation networks have been addressed and resolved by the SE800 MSER. This powerful platform provides carrier-class availability, service differentiation, and ease of management. Sophisticated features have been designed in to enable service providers to fine tune their traffic to guarantee their service level agreements, and provide reliability and quality of service to today's most demanding applications, VoIP, and multimedia on demand.

# 6 Glossary of Terms

| | |
|---|---|
| **ATM** | Asynchronous Transfer Mode |
| **BGP** | Border Gateway Protocol |
| **BRAS** | Broadband Remote Access Server |
| **CDR** | Committed Data Rate |
| **CIR** | Committed Information Rate |
| **DHCP** | Dynamic Host Configuration Protocol |
| **FRR** | Fast ReRoute |
| **GigE** | Gigabit Ethernet |
| **IGMP** | Internet Group Management Protocol |
| **IMS** | IP Multimedia System |
| **IRR** | Internal Rate of Return |
| **L2/3** | Layer 2 / Layer 3 |
| **LDP** | Label Distribution Protocol |
| **LSP** | Label Switched Path |
| **LSR** | Label Switching Router |
| **MPLS** | Multi-protocol Label Switching |
| **MSR** | Multi-Service Router |
| **MTBF** | Mean Time Between Failures |
| **MTRR** | Mean Time To Repair |
| **NPV** | Net Present Value |
| **P router** | Router at the Core of an MPLS network - LSR |

| | |
|---|---|
| **PE router** | Router at the edge of an MPLS network - LER |
| **P2P** | Peer to Peer |
| **PIM** | Protocol Independent Multicast |
| **PIR** | Peak Information Rate |
| **PoP** | Point of Presence |
| **PPP** | Point to Point Protocol |
| **PWE** | Pseudo-Wire Emulation |
| **QoS** | Quality of Service |
| **RNC** | Radio Network Controller |
| **ROI** | Return on Investment |
| **RSVP** | Resource Reservation Protocol |
| **SM** | Sparse Mode |
| **SSM** | Source Specific Multicast |
| **STB** | Set-Top Box |
| **TCO** | Total Cost of Ownership |
| **VBR-(n)rt** | Variable Bit Rate-(non)real time |
| **VLAN** | Virtual Local Area Network |
| **VoD** | Video on Demand |
| **VoIP** | Voice over IP |
| **VPLS** | Virtual Private LAN Service |
| **VPN** | Virtual Private Network |
| **WAN** | Wide Area Network |

# 7 Further Information

Product Specifications are subject to change without notice. Ericsson assumes no responsibility for any inaccuracies in this document and reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

ERICSSON
TAKING YOU FORWARD

**Rights and Restrictions**



www.ericsson.com