July 2003

WHITE PAPER

Public Ethernet

Enabling True Multi-Service Networks

ERICSSON

Contents

1 Executive summary

Delivery of broadband services sets new demands on a provider's network. To obtain maximum return on investment, the solution a network provider chooses must be cost efficient at delivering new and existing services, as well as scalable and flexible enough to support new services when demand rises. In addition, the chosen solution must integrate smoothly with existing management and business support systems, minimizing time to service without disturbing established revenue streams.

Ericsson has such a solution. Using Ethernet as the primary transport technology, Ericsson's Public Ethernet provides a cost efficient and scalable solution for the access provider's broadband infrastructure.

This whitepaper describes the key characteristics of Public Ethernet, including

- How to create a true multi-service broadband access network
- How Quality-of-Service (QoS) and security are implemented
- Management of the Public Ethernet

Using the broadband access infrastructure established by Ericsson's Public Ethernet, providers are offered the tools to create new and exciting revenue generating services, while at the same time minimizing the time and investment needed to develop and deploy them.

2 Background – Drivers toward a Broadband Multi-Service Architecture

Current broadband Internet access services are characterized by “best effort” quality of service, usually with a flat rate price structure and the only real service differentiator being raw access bandwidth. Stiff competition between providers has driven down the price for end-users, thus minimizing the providers’ revenue of selling access only.

Ericsson is firmly convinced that the future of broadband access networks is much more than just fast Internet connectivity. New applications are emerging, attracting both residential and business users. However, to take full advantage of this growing interest, the access network solution must support a multitude of applications, some of which for example streaming video, are much more bandwidth intensive than traditional Internet access.

The growth in applications is constantly increasing the requirements on network bandwidth and quality of service. The ability to increase bandwidth has drawn a lot of attention among providers and equipment vendors, and resulted in a number of standards, both for xDSL and fiber access. However, to enable a true multi-service network, the achievements in obtainable bandwidth on the first mile must be complemented by improvements in the network’s ability to ensure end-to-end QoS and capacity, from the end-user all the way to the Application Service Provider.

Access providers around the world have realized that offering only a basic Internet connectivity service will not generate significant revenues in the long run. Consequently, they are now investigating how new applications can create new revenue options. Figure 1 illustrates how the various drivers stimulate broadband access deployment causing increased revenue options for providers.

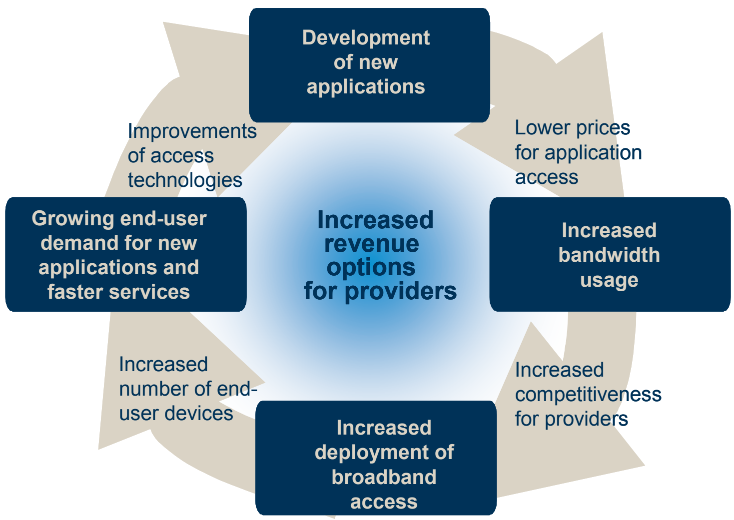

Figure 1. Drivers toward a broadband multi-service architecture

Increased broadband penetration and faster services have boosted the development of both new applications and broadband transmission technology. Applications include high-speed Internet access, audio and video streaming, on-line gaming, telephony services and virtual private networks for telecommuters and enterprises.

As a result of transmission technology improvements, especially within the xDSL area, new standards have emerged that enable much higher bandwidth in the first mile between end-users and the central office. Since many of the new applications are bandwidth hungry, these new standards provide a necessary foundation for deploying a multi-service network able to support the applications.

The development of new applications is also driven by the growth of networked devices in the home, which in turn increases traffic volumes and form the basis for still more applications. A wide range of connected devices, such as IP telephones, set-top boxes, and game consoles, accompanies the PC. An increasing number of household appliances are also built with network connections, enabling remote monitoring and control of their function as well as direct access to content providers.

3 Public Ethernet Key Characteristics

The growing number of applications is increasing the requirements on network bandwidth and quality of service. Real-time applications such as IP telephony and online gaming require low latency in the network, and bandwidth-hungry applications (such as video) increase the bandwidth demands significantly. The access providers' broadband network must be capable of handling all these requirements in a cost-efficient way, both now and in the future.

Ericsson's Public Ethernet is developed with a focus on these requirements.

In a nutshell, Public Ethernet is about bringing applications to end-users, by taking advantage of Ethernet technology. Ethernet serves as the transport mechanism between the end-users and the access provider, tying together the end-users' home network and the access network.

Ericsson's way of using Ethernet technology in Public Ethernet is unique, when building public access networks. Traditionally, broadband access networks have been based on ATM, but Ericsson has created a solution superior to these legacy and ATM-centric systems. Ethernet technology has already proven very successful in LAN environments where it today is ubiquitous and dominant. Some of the key features behind its success are low equipment cost, easy provisioning, simple maintenance, and good scalability. These benefits are also maintained when deploying Ethernet technology for broadband access.

Ethernet has continuously been re-invented over the last 30 years and is now moving into public networks with the 10GbE and Ethernet First Mile (EFM) standards. Ericsson has carefully chosen the best features of this technology, and matured it to become a solution that offers the reliability, scalability, QoS, security and manageability of a traditional telecom network, while maintaining the cost efficiency of Ethernet. The result – Public Ethernet – can be described by a number of key characteristics.

Multi-service support
Public Ethernet is a true multi-service enabler for broadband access. This includes the ability for each end-user to access multiple applications simultaneously, via different service providers, and with the QoS to ensure that the most sensitive and profitable services are prioritized.

Security ensured for end-users and providers
Public Ethernet uses several mechanisms for ensuring security and privacy for end-user data carried over the access network, as well as for protecting the network equipment from end-user attacks. Some of these mechanisms are based on established standards (for example virtual LANs) whereas others are enhancements developed by Ericsson to provide carrier-grade levels for services and security.

Standards based
Public Ethernet has complete interoperability with other network solutions thanks to its use of standard protocols and well-proven technologies. The basis is the use of switched Ethernet as standardized by IEEE, but also in all other areas the use of ITU-T and IETF standards constitutes the basis for the Public Ethernet architecture.

Cost-efficient with excellent scalability
The use of switched Ethernet technology has enabled Ericsson to provide a very cost-efficient solution for broadband access, compared to traditional ATM-centric solutions. At the same time, Public Ethernet proves to have superior scalability, making access networks of all sizes profitable, and allowing a seamless transition from smaller to larger networks as the broadband penetration and bandwidth requirements increase.
Smooth migration from legacy networks

Many incumbent providers have invested heavily in transport networks based on traditional technology. Public Ethernet enables these providers to reuse the backbone investment as an integrated part of their broadband solution, and use it either as the sole transport network or as a complementary transport network in parallel with other transport and access solutions.

Network resilience
Maximizing network availability is vital to a provider, because no operational network translates into no revenue. Public Ethernet features a number of mechanisms and protocols to provide robustness and resilience in the access network.

Complete suite of network management tools
To tie all parts of the access solution together, Public Ethernet comprises a suite of tools for managing the entire access network, and standard interfaces to easily integrate the Public Ethernet network with existing Network Management and Business Support Systems.

4 The Architecture of Public Ethernet

This document details the mechanisms used to deliver the key characteristics of Public Ethernet. As a foundation, the Public Ethernet network reference architecture and fundamental terms will be defined in this chapter.

4.1 Reference Architecture of the Public Ethernet Network

The reference architecture for an Ethernet centric access solution is depicted in Figure 2. The Public Ethernet portfolio primarily targets the access and regional network areas in this architecture.

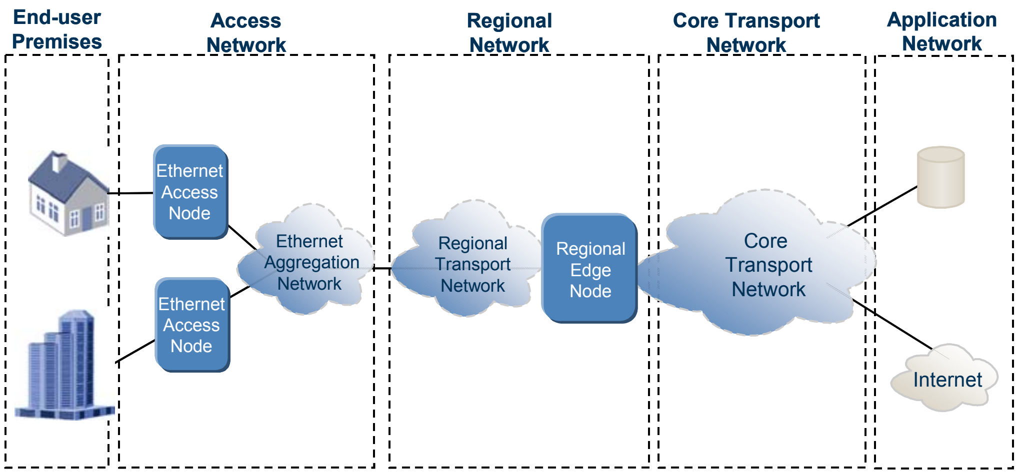

Figure 2. Network reference architecture

The access network comprises the Ethernet Access Nodes and optionally an aggregation network based on switched Ethernet technology. The (aggregated) Ethernet traffic may be carried over a regional transport network, possibly using other transport means than native Ethernet, such as PDH, SDH or ATM. The regional transport network may also be a metro Ethernet ring. The regional network is terminated in a Regional Edge Node, for example a BRAS (Broadband Remote Access Server), an IP router, or an Ethernet switch. Network Service Providers and Application Service Providers interface to the Regional Edge Node, possibly via a core transport network.

As in any other access system, the basic objective of Public Ethernet is to provide end-users with connectivity towards the application network. Thus, connectivity is a service provided to end-users, creating a logical connection – a *service binding* – between an end-user device and a network service or application.

4.2 Functional Architecture

The role of Ericsson's Public Ethernet may also be described using a layered model as shown in Figure 3. As indicated, the solution comprises three main functional areas: transmission technology, access/network services, and management. The applications are carried over the Public Ethernet infrastructure using various services and technologies.

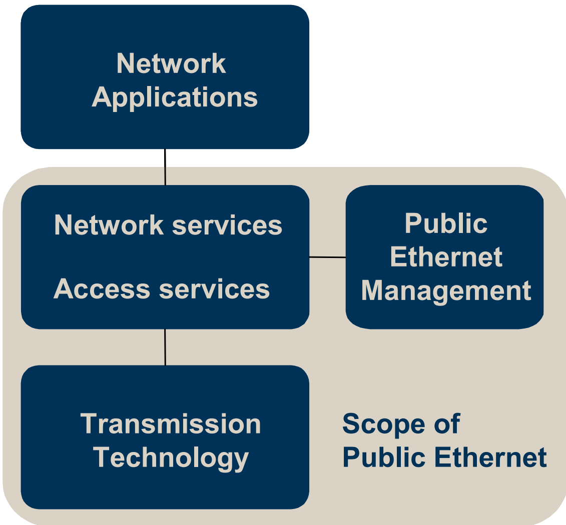

Figure 3. The scope of Public Ethernet

The top layer of the Public Ethernet model in Figure 3, *applications*, reflects real values that end-users are requesting. Using the definition made by the DSL Forum, applications are service offerings as experienced by the end-user via a device with an user interface (for example audio, visual display screen, joystick, remote control), typically located on or near the customer premises. Examples of applications include web access, telephony over IP, video streaming, and network gaming.

The lowest layer, *transmission technology*, deals with the physical infrastructure of the network, and how information is transported from one end of the network to the other. The key transmission technology for Ericsson's Public Ethernet is Ethernet, or rather *switched* Ethernet as standardized by IEEE 802.3.

Public Ethernet utilizes various physical medias for conveying Ethernet over the *first mile*: cat5/6 cable, optical fiber and local loop wires (twisted pair, CAT3), see Figure 4. For cat5/6 cable and optical fiber, the Ethernet protocol data units are carried directly over the physical media. This is denoted *native Ethernet access*. For local loop wires, ADSL technology is used. This includes the use of ATM as an intermediate layer, between Ethernet and ADSL, to ensure interoperability with existing ADSL modems. A Public Ethernet access network may deploy all of these drop technologies simultaneously in one integrated system.

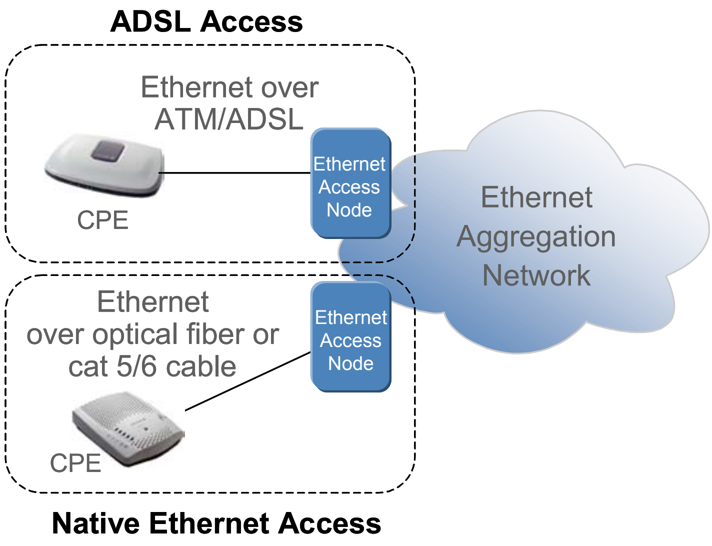

Figure 4. First mile technologies in Public Ethernet

Between the application layer and the transmission layer in Figure 3 are *services*. In the Public Ethernet context, services deal primarily with establishment of service bindings. A service binding has a number of attributes, including binding method and QoS characteristics.

In a traditional Ethernet (LAN) environment, for example within an office or home, the service layer is virtually non-existing. Ethernet technology is simply used to provide a cost-effective solution for interconnecting networked devices such as PCs and printers. Applications are executed over the LAN, which offers a single QoS level and where all devices compete for the available bandwidth.

However, the service layer of Public Ethernet is assigned a range of new responsibilities, corresponding to the previously described key characteristics of the solution. In other words, in Public Ethernet the service layer must also ensure multi-service support, QoS, security, resilience, etc in order to create a carrier-grade broadband access system based on Ethernet.

The service layer may be split into "*network services*" and "*access services*". *Network services* provide end-users with layer-3 connectivity, whereas *access services* deal primarily with layer-2 transport with certain traffic management functionality. Compared to applications, network and access services are not directly visible to the end-user, but they enable the end-user to access applications. Thus, applications sit on top of services. In general, services are offerings that support a value for the end-user, and from which the provider of the service can generate revenue. In general, a service is independent of the transmission technology, meaning that it can be offered via different technologies (for example fiber and ADSL).

5 Multi-service Support in Public Ethernet

Supporting the delivery of multiple services is a cornerstone of Public Ethernet. It enables the Access Service Provider to design and deploy services with a QoS profile that matches the requirements of the application(s) running over it. However, this ability to offer different service qualities must be complemented with a number of other system properties in order to create a true multi-service network architecture.

The fundamental aspects of Ericsson's Public Ethernet supporting multi-service delivery are as follows:

- Public Ethernet supports multiple Access Service Providers offering their services to all end-users on equal terms, a basic prerequisite for offering many-to-many access.

- Public Ethernet supports many-to-many access, enabling multiple users to utilize a single physical access connection (for example the local loop) to access more than one Access Service Provider simultaneously. Thus, from a single end-user's premises, multiple service bindings can be established towards different providers.

- Public Ethernet supports multiple applications over the same physical end-user connection, with different QoS and bandwidth profiles for each application. Different applications can have quite different requirements. Consequently, the access network must treat applications according to their service needs, giving preferential treatment to, for example real-time applications like telephony. Some applications, like Internet access, may also be offered with different service characteristics, for example different bandwidths or service guarantees.

- Public Ethernet supports any kind of application and combination of different applications, targeting both residential and business users. The service requirements from a business user view are not the same posed by a residential user. Business users generally have much higher requirements, and guaranteed service and network resilience are vital parameters to these users. In addition these services traditionally represent good revenue possibilities for providers.

- Public Ethernet supports AAA functionality (Authentication, Authorization, Accounting), enabling differentiated access to applications, and charging according to the specific service and application provided.

Traditional ATM-centric broadband access systems cannot compete in a cost efficient way with these properties of Public Ethernet. The DSL Forum has identified the major evolutionary requirements on access architectures in the document *Multi-Service Architecture and Framework Requirements*[1]. The list of improvement requirements draws a picture of traditional ATM-centric access architectures lacking most of the capability and flexibility to meet even basic multi-service requirements. In contrast, Public Ethernet addresses these requirements in creating an Ethernet-centric broadband access architecture that fulfills the multi-service requirements posed by the DSL Forum.

---

[1] Currently working text WT-080

5.1 Implementing Multi-service Access in Public Ethernet

As a cornerstone in the multi-service architecture, Public Ethernet supports multiple logical *service channels* on the first mile. The implementation of a service channel in Public Ethernet depends on the transmission technology used on the first mile (ADSL or native Ethernet access):

- For ADSL access, Public Ethernet supports multiple ATM permanent virtual circuits (PVCs) per end-user. Each PVC represents a service channel.
- In Public Ethernet based on native Ethernet access, the service channels are represented with different user *VLANs* (virtual LANs).

A service channel serves two purposes:

- It provides a certain transmission service for the end-user traffic carried on it. The characteristics of each service channel are configured individually, enabling different bandwidth and QoS settings for each channel.
- It conveys traffic pertaining to a certain access network service. The Ethernet Access Node performs a service binding between the service channel and a service VLAN within the access network.

For ADSL access the ATM layer is terminated in the Ethernet Access Node (the DSLAM), as opposed to traditional ATM-centric DSL access systems. This enables an access network where all drop technologies interface with a common switched Ethernet aggregation network, as illustrated in Figure 4.

In the aggregation network each end-user service is associated with a service VLAN. The Ethernet Access Node, interfacing the first mile connection (the U-interface), performs a service binding between the service channel and the service VLAN. Figure 5 illustrates the binding between service channels (PVC and user VLANs) and service VLANs. Note that both ADSL technology and native Ethernet access technology can interface with the same service VLAN simultaneously, thus creating an integrated access scenario where both transmission technologies can seamlessly coexist.

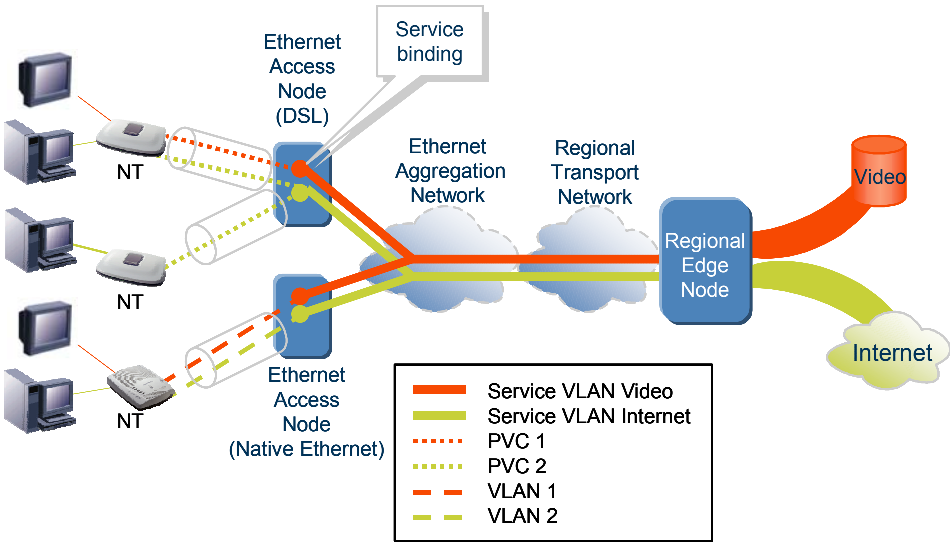

Figure 5. Service binding between service channels and service VLANs

The Ethernet Access Node is, in the ADSL access case, an IP DSLAM, and in the native Ethernet access scenario, an Ethernet switch. In both cases, the Ethernet Access Node performs *advanced Ethernet bridging*, meaning that it is capable of filtering and manipulating the information being bridged over it, both on layer 2 (Ethernet) and layer 3 (IP) -functions crucial to ensure traffic separation and end user security.

The logical mapping between user devices and service channels in the user Network Terminal (NT) can be based on various criteria. For example, each device may be attached to a specific physical NT port, for example one port for voice, one for TV and one for Internet access, and with each port corresponding to a specific service channel. Alternatively, the mapping may be based on more dynamic settings, or done automatically as a learning bridge.

In the first mile, each service channel has a well-defined service characteristic. For ADSL access, each PVC is associated with a number of QoS parameters, depending on the type of PVC (e.g. UBR, VBR-rt). For native Ethernet access, the maximum bandwidth of each user VLAN is policed in the Ethernet Access Node.

The use of ATM over DSL in the first mile provides compatibility with existing end-user modems, and QoS on the relatively slow ADSL links that must support real-time applications, for example Telephony over IP. Such applications may experience problems regarding delay and jitter if long Ethernet frames pertaining to other user applications are not segmented. Consequently, separate ATM PVCs make it possible to ensure that the real-time traffic obtains adequate treatment on the low speed link. On higher speed links and in the aggregation network where capacity is abundant, delay and jitter are much less of a concern since queuing situations are less likely to occur.

VLAN technology is a cornerstone in the multi-service architecture of Public Ethernet. By using this technology, the switched Ethernet access network is split in logically separated networks, each representing a specific service. VLAN technology is standardized in IEEE802.1Q.

Basing the access network on the IEEE802.1Q standard enables prioritization of Ethernet frames within the network. Before entering the aggregation network, each frame is tagged with a priority label by the Ethernet Access Node, and the Ethernet switches, constituting the aggregation network, give preferential treatment to higher priority traffic. The standard defines 8 levels of priority, but fewer levels (for example 4) are normally sufficient to ensure adequate multi-service functionality.

For upstream Ethernet frames, the Ethernet Access Node attaches the priority label according to the characteristic of the service channel, thus preserving the QoS for that channel. For downstream traffic the priority label is attached to the Ethernet frames before entering the switched Ethernet.

Within the Ethernet access network, each service VLAN may be allocated a specific bandwidth, managed by a connection manager tool, thus enabling service and tariff differentiation for the Access Network provider.

A service binding made in the Ethernet Access Node provides *basic connectivity* to a service node, for example an application server, an IP router, or BRAS. In the latter example, the BRAS complements the basic service binding with an option for dynamic service selection.

5.2 Service Access with and without a BRAS

Traditionally a multi-service access architecture is dependant on one or more Broadband Remote Access Servers (BRAS) for creating and administering service bindings. In a typical ATM access scenario, end-users log into the BRAS using the PPP protocol, and their login credentials are used both to validate the end-user and determine the requested service.

Public Ethernet fully supports this architecture. The BRAS may be used to perform a number of functions, including a wide range of access protocols, tunneling, VPN, full IP routing, MPLS, QoS shaping and policing, and AAA.

The BRAS may be complemented with a policy server which dynamically allocates traffic capabilities per IP session to end-users. The policy server enables end-users to dynamically change their service profile within the BRAS, for example for changing bandwidth or QoS settings.

A BRAS is normally located at a central point in the network, because it must aggregate traffic from a large number of end-users to minimize the cost per end-user. In this case end-user traffic must often traverse an intermediate transport network to obtain connectivity with the BRAS, where network and/or application access is granted and controlled.

The Public Ethernet network supports one or more BRASs, located in different service VLANs. The Ethernet Access Node is provisioned to perform a semi-static service binding between a service channel and the service VLAN containing the BRAS (c.f. Figure 5), whereas the BRAS itself is capable of performing a more dynamic service binding based on, for example, login credentials provided by the end-user.

However, in some scenarios the basic service binding created in the Ethernet Access Node may be considered sufficient to manage the end-users' access to applications. In that case the functionality of the BRAS is not required. Instead, a service node may interface directly towards its corresponding service VLAN in the access network. This approach has certain advantages:

- The BRAS is offloaded or simply not needed.
- Applications are easily added to the network. In principle, any type of service node may be connected in this way.
- The service nodes may be placed locally in conjunction with the access network, thus minimizing the traffic load on the transport network and reducing transmission delay for real-time applications.

A good example of a locally placed service node is for distribution of video. The video content is stored in a local server, which streams the signal into the access network, within the service VLAN assigned for this video service.

5.3 Multi-provider Support in Public Ethernet

A Public Ethernet scenario supports a number of different provider roles. The most important are:

- The Access Network Provider, owning and maintaining the access network equipment, and offering services to Access Service Providers.
- The Access Service Provider, providing access connectivity to end-users on the Ethernet level, based on the services created by the Access Network Provider.
- The Network Service Provider, providing network connectivity to end-users, typically towards an IP network (for example the Internet).
- The Application Service Provider, offering applications and content via the broadband network.

In some scenarios, two or more provider roles collapse into one entity, or a role may even not exist.

The Access Service Provider has two main options for providing connectivity to Network or Application Service Providers. The BRAS provides an advanced set of mechanisms for QoS, authentication, service selection and charging, primarily done at the IP-level. Alternatively, if the requirements for such mechanisms are less demanding, the Ethernet access network may interface directly to the Network Service provider or Application Service Provider, see Figure 6. As previously mentioned, this direct interface may enable a fast and cost-efficient deployment of new services, possibly distributed close to or in conjunction with the access network. As indicated in the figure, both options may seamlessly co-exist in a single Public Ethernet access network.

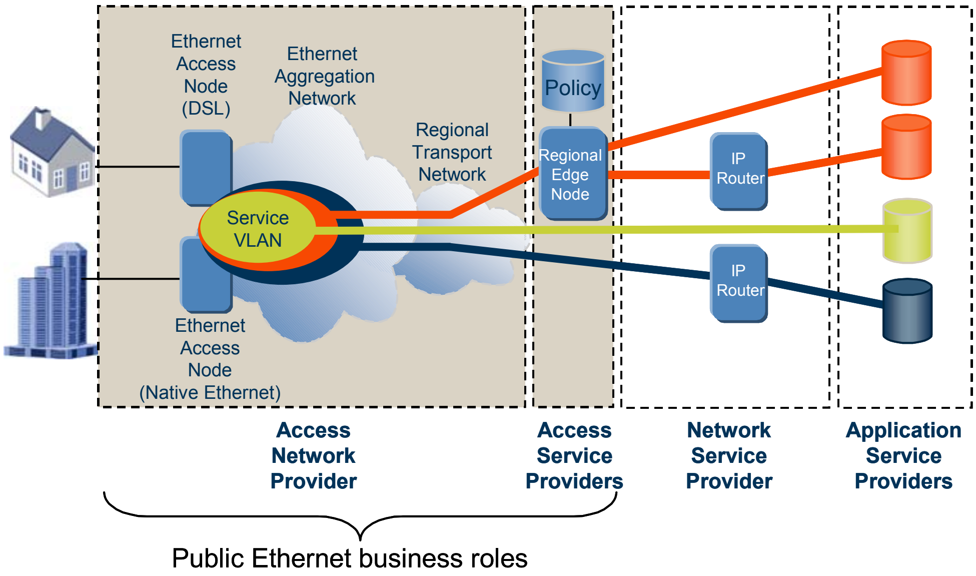

Figure 6. The different roles in the delivery of broadband services

6 Ensuring Security in Public Ethernet

All networks are susceptible to attacks from both external and internal parties. The attacks can have many forms, including information theft, denial-of-service attacks, and corruption of programs or information.

A traditional Ethernet network is based on broadcast media, as opposed to the circuit-oriented architecture of legacy access systems based on ATM. Consequently, Public Ethernet uses several mechanisms for separating different traffic types, ensuring the privacy and integrity of the transported data, and protecting the equipment owned by subscribers and providers.

6.1 Traffic Separation with Virtual LANs (VLAN)

By employing VLAN technology an operator can separate different traffic types. A VLAN is a logically separated broadcast domain within an Ethernet network. Communication between VLANs must always traverse a layer-3 device, typically an IP router. Thus, an immediate use of VLAN is to separate management traffic and subscriber traffic in different VLANs. In this way, malicious subscribers cannot address or abuse the Ethernet domain by tampering with the Access Network Providers' equipment. The IP routers at the network edge are configured to prevent traffic in subscriber VLANs from accessing the management VLAN.

As indicated in Figure 5, the VLANs are used to separate different services into different service VLANs, thus preventing unauthorized service access from end-users. A VLAN may also be allocated to a specific end-user, in order to create an isolated layer-2 network for this end-user alone, e.g. for a transparent LAN service for an enterprise or other organization.

6.2 PPPoE and MAC Forced Forwarding

End-users accessing the same service VLAN are in fact connected to the same broadcast domain. In a traditional LAN environment this causes end-users to have *layer-2 visibility*, meaning that their network devices can access each other on the Ethernet layer. These "local connections" causes problems for two reasons: End-users connected to the same service VLAN can establish peer-to-peer connections between each other directly on layer 2, or a malicious end-user can perform specific layer-2 attacks on other end-users.

The problem with these connections is that this traffic will not pass through the provider's service node, for example the BRAS, IP router or application server. Consequently this node cannot provide access control, security filtering, charging or capacity allocation.

Public Ethernet deals with these security issues in two ways, depending on whether the end-user access protocol is IP-over-Ethernet or PPP over Ethernet.

In general, layer-2 visibility between end-user devices requires knowledge of the other peer's layer-2 address – the MAC address. In an IP-over-Ethernet access scenario, a device will normally learn this MAC address by broadcasting an ARP request. To prevent this, and to direct all upstream traffic to the Access Service Provider's service node, the Ethernet Access Node implements a mechanism called *MAC Forced Forwarding*. An ARP proxy located in the Ethernet Access Node manages the core of this functionality. This ARP proxy will ensure that the end-user device sends all upstream IP traffic towards the MAC address of the service node (typically the default gateway). If the end-user tries to violate this rule, the upstream traffic is filtered out. In this way, all end-user traffic is forced to go via the service node (for example a BRAS) where access control, filtering and charging can be applied.

For PPP access, the PPPoE protocol creates a logical connection between the end-user device and the service node (the BRAS). Security is ensured with filtering rules in the Ethernet Access Node: Only PPPoE frames are accepted on this service channel, and broadcast traffic is never sent downstream. These rules ensure that upstream traffic can only flow to the BRAS, and thus no two end-users can obtain layer-2 visibility.

6.3 End-user Identification Using DHCP option 82

A service-VLAN offering IP-over-Ethernet access may deploy a DHCP server for dynamic handling of end-users' IP address configurations. Normally, the provider of the IP address must know the exact identity of the requesting end-user, in order to enable back tracing of malicious users.

For this purpose, the Ethernet Access Node can insert the *DHCP relay agent information option* (DHCP option 82) field into the DHCP request. This field contains the identity of the Access Node end-user port that originated the DHCP request, i.e. the copper line identity or switch port identity used by the subscriber.

6.4 Virtual MAC Addressing

One reason that a device's MAC address is not usable for end-user identification is that the end-users themselves easily can modify MAC addresses. Additionally multiple devices may appear in the Ethernet access network with identical MAC addresses, either intentionally (i.e. MAC spoofing) or by accident.

To counter this problem Public Ethernet offers a functionality called *virtual MAC addressing*, in which each end-user device MAC address is mapped in the Ethernet Access Node to a unique locally generated MAC address[2]. This local MAC address is used within the access network as an alias for the device MAC address.

The local MAC address is constructed in a way that indicates the service channel. Consequently, a virtual MAC address can be used to identify the exact source of all upstream traffic. In PPPoE access scenarios this feature enables end-user traceability, because it provides a method to identify the physical line from which a PPPoE session initiation was received.

6.5 Filtering in the Ethernet Access Node

Filtering rules implemented in the Ethernet Access Node provide additional security. It is possible to set up rules for traffic processing on layer 2 (Ethertype, MAC addresses), layer 3 (IP addresses, ICMP messages), and layer 4 (port numbers). Specifically, for IP-over-Ethernet access the source IP address in upstream traffic is verified, in order to avoid IP spoofing. In DHCP based access, the valid IP address is obtained by snooping the DHCP reply sent from the DHCP server to the end-user. Likewise, for end-users with static IP addresses, the Ethernet Access Node is configured from the management system to accept only traffic from the assigned IP address or IP subnet.

6.6 Preventing broadcast storms

Broadcast Ethernet frames from end-users are generally not allowed to pass the Ethernet Access Node, because they could severely load the access network. Exceptions are broadcast frames related to ARP, DHCP and PPPoE. However, ARP broadcast is, limited to a minimum by the use of MAC Forced Forwarding. Furthermore, the number of outstanding DHCP requests and PPPoE session initiations are limited by the Ethernet Access Node.

7 Basing Public Ethernet on Standards

Public Ethernet is based on widely adopted standards and technologies. This enables Public Ethernet to be seamlessly integrated with the Access Provider's existing access solution, and creates a solid foundation for interfacing with Application Service Providers and Network Service Providers.

The Ethernet technology itself is defined in a number of IEEE standards. For Public Ethernet, the IEEE802.1Q standard is of special significance, because it gives the

2 Local MAC addresses are specified in IEEE802.

foundation for *switched Ethernet* – a core technology in making Public Ethernet a carrier-grade solution, because it enables the use of VLAN:s and priority indications within the Ethernet-based access network. The IEEE802.1Q Ethernet connections are all point-to-point and full duplex.

At the end-users' premises, Ethernet is the natural choice for hooking up one or more hosts to the users' network termination (NT) in a fast and easy way. The ease of installation at the end-users' premises is a major requirement posed by Access Service Providers, resulting in Ethernet being the prevailing and most popular modem interface today.

Ethernet is the natural choice of protocol for transporting the Internet Protocol (IP). The majority of all IP traffic originates and terminates on Ethernet networks. It is believed that more than 90% of all IP traffic originates and terminates in Ethernet networks. The Public Ethernet network itself relies on IP technology: All management traffic uses IP, and Ethernet Access Nodes are able to perform security filtering and manipulation of IP traffic to and from end-users.

The Public Ethernet management tools use industry standard interfaces such as SNMPv2 and CORBA.

The Public Ethernet supports multiple access protocols from the end-users:

- PPPoE (RFC2516)
- PPPoA (RFC2364)
- IP-over-Ethernet
- Transparent LAN service (Ethernet-based VPN)

All these protocols can seamlessly co-exist in a single Public Ethernet access network. The type of access protocol is configured on a per-service-channel basis, meaning that a multi-service enabled user may use different protocols to different providers.

In an ADSL based network, the Ethernet protocol layer is transferred into ATM AAL5 according to RFC1483/2684. Use of the PPPoA access protocol is only relevant in an ADSL network, and the Ethernet Access Node converts it to PPPoE. Between the Ethernet Access Node and the BRAS, PPP-based access is always run as PPPoE.

For IP access over Ethernet, end-users may use static IP addresses or obtain IP addresses via DHCP (RFC2131).

Multicast sessions are supported via IGMPv2 (RFC2236). The Ethernet Access Nodes and the access network Ethernet switches are able to perform IGMP snooping allowing for an efficient distribution of this type of traffic.

Finally, the Public Ethernet architecture is completely neutral with respect to first mile drop technologies, allowing easy integration of any existing and new standard drop technology such as ADSL, ADSL2, ADSL2+, Ethernet First Mile (EFM) 10GbE, GbE, 100-Base-LX10 and 10Base-TS.

8 Public Ethernet: Cost-efficient with Excellent Scalability

Public Ethernet focuses on minimizing both capital and operational expenditures (CAPEX, OPEX), making it possible for providers to offer attractive services and applications at competitive prices.

The use of Ethernet technology is a major factor in the cost-efficiency of Public Ethernet. Ethernet technology has proven itself very successful in LAN environments where today it is ubiquitous and dominant. Some of the key features behind its success are low equipment cost, easy provisioning, simple maintenance, and good scalability. With Public Ethernet these benefits are also obtained when deploying Ethernet technology as an access technology.

As opposed to ATM, an Ethernet network is not connection-oriented, so individual traffic flows need no prior configuration of the access network. Instead, by inspecting the source addresses of incoming frames the access network switches quickly build up knowledge about the network topology in order to direct traffic in the right direction only. This self-configuring feature of switched Ethernet ensures a simple and stable network installation. Another advantage, is that the bandwidth of the Ethernet access network may effectively be shared between all or a group of end-users, as opposed to the more static bandwidth allocation known from ATM-centric systems.

Ethernet technology is particularly well suited as the foundation for a multi-service network architecture. Some of the most demanding applications, such as real-time broadcast TV, can be transferred over Ethernet as one multicast stream, instead of having a stream per end-user as in ATM-centric access networks. Thereby precious bandwidth of both the access network and transport network is saved.

High scalability is a major asset of Public Ethernet, because it can be adapted to both very small sites as well as large sites. Especially with DSL rollout, downwards scalability is often important to create an acceptable business case in rural areas or locations with low DSL penetration. Public Ethernet offers its unique IP DSLAM for this purpose. Each IP DSLAM is a system-on-a-board, handling 10 end-users. The IP DSLAM is a non-blocking Ethernet Access Node for ADSL, enabling every end-user to reach his theoretical maximum capacity of 8 Mbit/s.

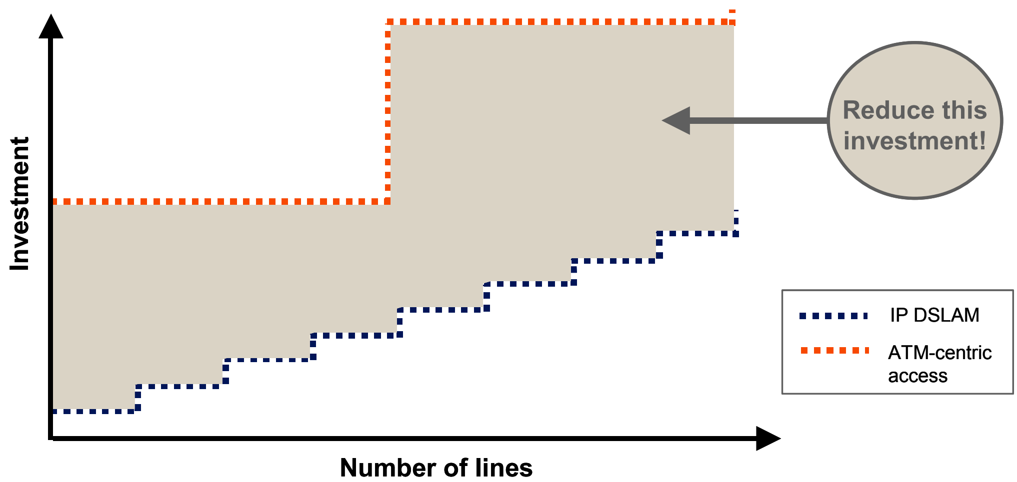

Figure 7. Public Ethernet for best scalability

In conclusion, a healthy business case is made possible with practically any access network size – ranging from a handful of end-users to tens or hundreds of thousands of end-users. As broadband penetration grows, the access network is gradually expanded to match the demand and bandwidth requirements, see Figure 7.

9 Enabling Smooth Migration towards Public Ethernet

For existing ATM-centric access providers Public Ethernet represents a paradigm shift. Although the many advantages of Ethernet-centric broadband access are compelling, some of these providers may find it difficult to migrate their existing network architecture towards this new concept. However, Public Ethernet also provides for easy migration from an ATM-centric environment.

The migration to Ethernet may be addressed in different ways, depending on the existing architecture and on the migration plans of the provider. In general, Ethernet technology can be deployed either as an access technology or as a transport technology. For example, a common use of Ethernet today is for upgrading ATM-centric access solutions to interface with an Ethernet-based transport network (for example. an Ethernet metropolitan network), because this technology provides advantages in terms of cost-efficiency and bandwidth capacity.

The Public Ethernet portfolio comprises elements to create true Ethernet-based solutions for both the access and the transport part of the broadband network. Furthermore, these solutions may be integrated, to tailor the migration path towards Ethernet to the needs of the individual network provider.

Although more and more providers are migrating to pure Ethernet transport solutions, legacy transport technologies will continue to exist in the foreseeable future.

Consequently, Ericsson offers a range of solutions that interfaces the Ethernet access network with such legacy transport technologies, as ATM, SDH, and PDH. This product portfolio covers:

- Ethernet over SDH or PDH, either as E1, E3 or STM-1 (2-155 Mbit/s)
- Ethernet over radio link, point-to-point or point-to-multipoint (2-155 Mbit/s)
- Ethernet over next generation SDH network
- Ethernet over ATM AAL5
- Ethernet over C-WDM

For example, the “Ethernet over ATM AAL5” solution uses a gateway to map between the VLANs of the Ethernet aggregation network and the PVCs of the regional ATM transport network. In this way, the traffic that was divided in different service-VLANs within the access network is still segregated over the transport network, and may be connected to different service providers via the ATM network.

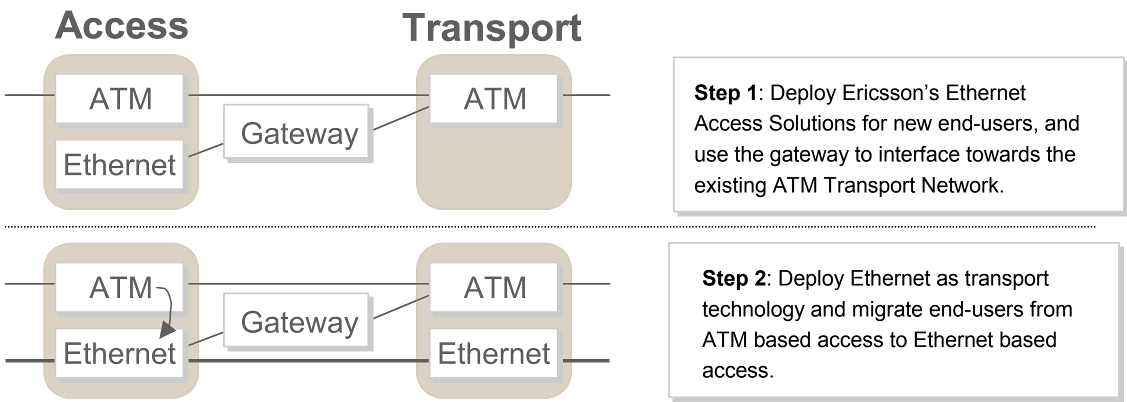

Figure 8. Ethernet migration scenario for ATM-centric providers.

Although Ethernet is the prevailing technology at the end-users’ premises, some modems are still using PPP over ATM (PPPoA) as the access protocol. To address this user segment, the Public Ethernet IP DSLAM supports PPPoA access, and converts this protocol into PPP over Ethernet to a provider’s BRAS.

10 Management of the Public Ethernet

Telecom management is one of the most important factors in creating a carrier-grade solution based on Ethernet technology. Traditionally, Ethernet networks have required very little provisioning and supervision when deployed in office environments or in homes. However, adequate tools for performing operation, administration, maintenance and provisioning of the network and its users must accompany the introduction of Ethernet in public networks.

Ericsson's Public Ethernet offers a complete suite of management tools, each addressing different aspects of the management tasks. Some of the tools are optional, because they are applicable only in specific network scenarios.

For the most basic management of the access network part of Public Ethernet, an element manager tailored to the characteristics of the individual IP-DSLAMs is required. For example, it must handle the definition of access network services, as well as the assignment of these services to end-users, even in situations with several Access Service Providers using the same Access Network Providers offering (unbundling).

Provisioning of services to end-users comprises handling of the settings in the Ethernet Access Nodes, for example regarding ADSL and ATM parameters for each end-user. The element manager must also handle installation and upgrade of access network nodes, as well as basic fault and performance management tasks for the access network. In a DSL scenario this includes, for example, functionality for local loop testing. To ensure smooth integration with existing management systems, the element manger must offer northbound management interfaces based on industry standards such as CORBA, and use profiles to hide the implementation details of the nodes.

The element manager may be complemented by additional management tools, depending on the nature of the provider's situation and Public Ethernet solution. Such management tools comprise:

- An Ethernet connection management system, for managing virtual layer-2 Ethernet connections end-to-end.
- Element mangers for the Ethernet switches of the Ethernet aggregation and transport network.
- A configuration server, for managing advanced end-user devices, such as residential gateways.

The Ericsson offering for management of Public Ethernet includes Element Managers for all Access and Regional Network elements and a Connection Manager application to support service provisioning. These applications provide well-defined open and published integration reference points, making them easy to integrate into a provider's existing management environment.

11 A True Multi-service Broadband Access Solution

The future of broadband access will create a number of challenges for service providers. The networks must provide existing services at a lower cost and simultaneously support many new types of revenue-generating applications, demanding high bandwidth and real-time traffic handling.

Ethernet is superior as the technology for broadband access networks, due to its scalability, efficiency, low cost, and easy deployment. The Ericsson Public Ethernet enables providers to perform a smooth migration from legacy ATM access networks to Ethernet-centric access, minimizing the investment and reusing their existing infrastructure in the future.

With Ericsson's Public Ethernet, providers are optimally equipped to meet the future. Ericsson's approach is to combine standard Ethernet functionality with advanced features, thus creating a true multi-service access network solution with carrier-grade performance.