The Public Ethernet—The next-generation broadband access network

Michael Begley

Thanks to its compatibility with IP, simplicity of deployment, and above all else, price-to-performance ratio, Ethernet is emerging as the leading standard for access technology. However, as currently deployed in the enterprise environment, Ethernet does not meet the requirements of a public telecommunications network. A broadband access network that carries real-time services must scale to handle large numbers of simultaneous users, offer excellent in-service performance, differentiate between services, and provide secure connections. Likewise, operators must have a management solution that helps them to roll out and maintain their networks and provide customer support.

This article describes Ericsson's Public Ethernet solution, which has been developed using standard Ethernet technology and incorporates the reliability, quality of service, security and manageability necessary for a public telecommunications network.

Broadband market trends

Telecommunications service providers are facing major changes in the services they provide and the infrastructure they use. In the broadband access market, some of the driving forces for these changes include

- demand for lower costs—the combination of a single service offering (high-speed Internet access) and relatively high access costs (especially in the backhaul connection) means that most operators are not yet generating sufficient revenue from their current broadband offering. Notwithstanding, to increase broadband penetration, they must lower their prices;
- increasingly complex home networks—the current broadband service offering (high-speed Internet access) is primarily intended for one PC per connection. But the emerging trend is for multiple devices per connection, such as two or more PCs, game consoles, and TV set-top boxes;
- new broadband applications—as broadband penetration increases and the market grows, the number of applications will also grow to meet the demand of increasingly technology-savvy users. Today's youth and professional early adopters, for instance, are driving the demand for the delivery of advanced applications, such as remote access, in-home health care, games on demand, and gambling;
- demand for greater bandwidth—advanced applications require greater bandwidth capacity, especially where multiple users share the connection. At present, applications such as online gaming account for nearly 9% of all backbone traffic in the USA, and as new services become available, end-users will indirectly demand greater bandwidth from the access network; and
- the need for service differentiation—the one-size-fits-all approach is not sustainable. Indeed, operators are already starting to differentiate service offerings with tiered bandwidth (512 kbps, 2 Mbps, 6 Mbps, and more). Ideally, the next step will be an access network that can deliver multiple services, such as a temporary bandwidth upgrade for, say, gaming or video conferencing.

For operators, the key to success in an increasingly competitive environment is to drive down the cost of delivering existing services and to open the door to new revenue-generating services. There are two aspects to delivering a cost-effective solution: scalability and cost (Figure 1). Ericsson's Ethernet DSL access node reduces costs substantially. Being uniquely scalable and making use of Ethernet technology, the IP DSLAM is a complete system on a board that facilitates expansion in steps of 10 or 12 subscribers at a time. The Ethernet DSL access node thus scales linearly to fit sites with as few as 10 subscribers to those with thousands of subscribers. As a backhaul link

BOX A, TERMS AND ABBREVIATIONS

ADSL	Asymmetrical DSL
ARP	Address resolution protocol
ATM	Asynchronous transfer mode
BRAS	Broadband remote access server
CAPEX	Capital expenditure
CESoIP	Circuit emulation services over IP
CPE	Customer premises equipment
C-WDM	Coarse wavelength division multiplexing
DHCP	Dynamic host configuration protocol
DSL	Digital subscriber line
DSLAM	DSL access multiplexer
E1, E3	Digital circuit based on PDH (Europe)
EAPS	Ethernet automatic protection switching
EDA	Ethernet DSL access
FE	Fiber Ethernet
FEA	Fiber Ethernet access
FRoIP	Frame Rely over IP
Gbit	Gigabit
GE	Gbit Ethernet
ICMP	Internet control message protocol
IEEE	Institute of Electrical and Electronic Engineers
IGMP	Internet group message protocol
IP	Internet protocol
ISP	Internet service provider
ITU	International Telecommunication Union
Kbps	Kilobits per second
LAN	Local area network
MAC	Media access control
Mbps	Megabits per second
OC3	Optical carrier level 3
OSI	Open systems interconnection
OSS	Operations support system
PBX	Private branch exchange
PDH	Plesiochronous digital hierarchy
PIM	Protocol-independent multicast
PPP	Point-to-point protocol
PPPoA	PPP over ATM
PPPoE	PPP over Ethernet
PVC	Permanent virtual circuit
QoS	Quality of service
RADIUS	Remote authentication dial-in user service
RFC	Request for comments
RSTP	Rapid spanning tree protocol
SDH	Synchronous digital hierarchy
SLA	Service level agreement
SONET	Synchronous optical network
T1, T3	Digital circuit based on PDH (North America)
TDM	Time-division multiplexing
VLAN	Virtual LAN
VMAC	Virtual MAC
VoIP	Voice over IP
WAN	Wide area network
WDM	Wavelength division multiplexing

("second-mile" transport), Ethernet is considerably cost-effective. The cost of an Ethernet uplink, for example, can be one-tenth that of an asynchronous transfer mode (ATM) uplink.

The combination of scalability and low cost makes Ethernet an attractive technology for operators who want to deploy a broadband access network. However, before the technology can be deployed in a public telecommunications network it must be adapted to fulfill several key requirements.

The Public Ethernet network

A public broadband access network must support traffic separation, service differentiation (quality of service), be secure, support multicast, be robust (in-service performance), and have a telecommunications management solution to support operation and maintenance of the network.

Traffic separation prevents end-users from eavesdropping the traffic of other end-users. It also separates services and other service provider traffic, giving the network operator full control of who talks to whom, thereby guaranteeing that only authenticated users may use network resources.

The definition of different classes of quality of service (QoS) makes it possible to differentiate between services—for example, those that are sensitive to delay and packet loss and those that are not. This ensures that the most sensitive applications and best-paying services receive priority when there is congestion in the network.

Although most end-users are genuine, a small percentage of them can be malicious. Therefore, to avoid fraud and service outage, operators must put security mechanisms in place to protect the network and other end-users.

Over the years, numerous experiments have been made to deliver broadcast services such as TV and radio over broadband networks. At present, the only cost-effective method of delivering broadcast services is multicast, where a single stream of data is transported through the network as close to end-users as possible and then duplicated and delivered to multiple end-users.

In-service performance is important to operators because network downtime results in lost revenue. Likewise, a resilient network is needed to fulfill service level agreements (SLA), avoid poor reputation and prevent customer churn.

Operators must have a comprehensive telecommunications management solution in place before they can build out infrastructure, provide services and effectively manage their networks. Present-day Ethernet-based local area networks (LAN) deployed in the enterprise environment cannot fully meet these requirements. Ericsson has thus designed and developed the Public Ethernet network solution (Figure 2).

Traffic separation

Originally, Ethernet LANs were based on coaxial cables, with each network device tapping into the cable. This resulted in poor performance and made it difficult to isolate faults. As Ethernet LANs continued to grow, the network devices were linked to a repeater (Ethernet hub) using twisted-pair wiring. In an Ethernet network built using hubs, only one device can transmit or receive data at a time and every device on the LAN must thus compete for available bandwidth. At best, this solution makes use of only 40% of available capacity.

To apply Ethernet on a larger scale, for example, in a wide area network (WAN) or Public Ethernet network, Ethernet switches replace the hub in the network. Ethernet switches examine packets, and process and

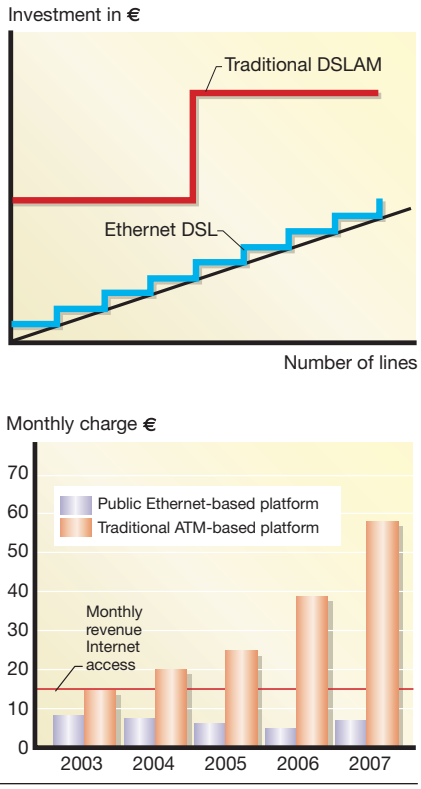

Figure 1
Overview of the scalability and cost competitiveness of the Public Ethernet network.

Figure 2
Overview of the Public Ethernet network.

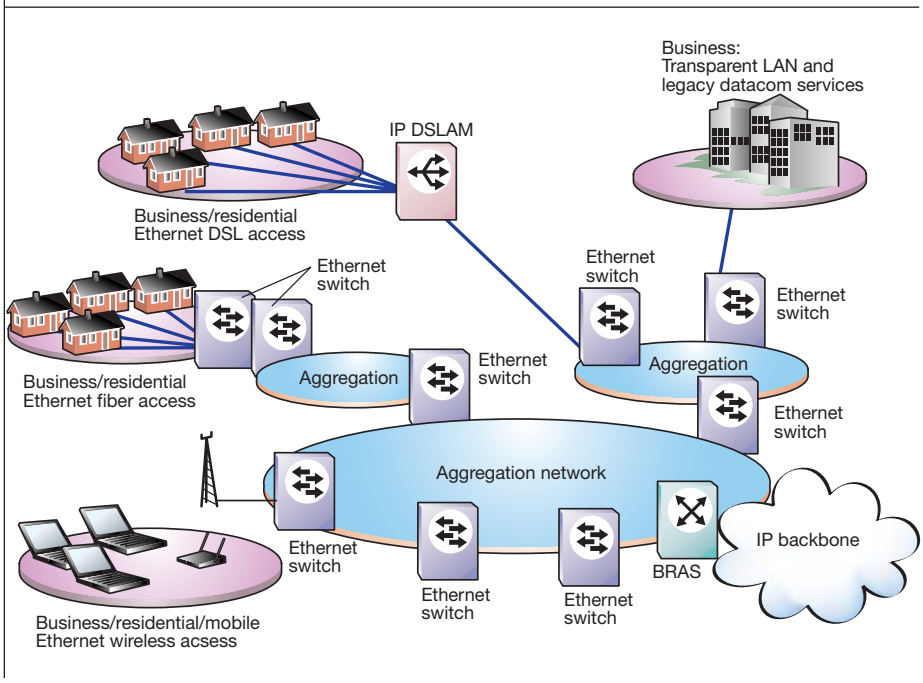

prioritize them instead of simply repeating the signal to every port in the network. Switches map the Ethernet addresses of the nodes that reside on each network segment and only allow pertinent traffic to pass through them.

Ethernet switching thus facilitates a significant increase in bandwidth by enabling simultaneous switching of data packets between its ports, and segmenting the network into different broadcast domains.

A LAN is defined as a single broadcast domain. Therefore, when a user broadcasts information on the LAN, every user on the LAN receives the broadcast. The virtual LAN (VLAN) standard described in IEEE 802.1Q is an extension of the Ethernet standard that describes how a LAN can be logically segmented into different layer-2 broadcast domains. Because the segmentation is logical and not physical, end-users do not have to be at the same physical location.

For residential subscribers, VLANs are used to separate services into service VLANs. To prevent unauthorized access to the services, end-users are solely connected to the service VLANs to which they have subscribed (Figure 3).

End-users who access the same service VLAN are in fact connected to the same broadcast domain. In traditional LAN environments, the end-users would have layer-2 visibility—that is, their network devices would be able to communicate with one another on Ethernet layer 2. This gives rise to two main problems:

- End-users connected to the same service VLAN might establish direct peer-to-peer connections on layer 2, bypassing the service provider's broadband remote access server (BRAS), which controls access, provides security filters, manages charging, and allocates capacity.
- A malicious end-user might execute specific layer-2 attacks on other end-users.

Ericsson's Public Ethernet solution resolves these problems. Layer-2 visibility between end-user devices requires knowledge of a peer's layer-2 address, or media access control (MAC) address. Devices obtain the MAC address of other devices by broadcasting an address resolution protocol (ARP) request. To prevent this from occurring, and to direct all upstream traffic to the BRAS, the Ethernet access node implements the MAC forced-forwarding mechanism. An ARP proxy in the Ethernet access node prevents end-users from sending broadcast requests into the network by always responding with the MAC address of the BRAS. This ensures that end-user devices obtain the MAC address of the BRAS but not of other users. The upstream traffic of users who try to violate this rule is filtered out. All end-user traffic is thus forced to pass via the BRAS (Figure 4).

For point-to-point protocol (PPP) access, the PPP over Ethernet (PPPoE) protocol creates a tunnel between the end-user device and the BRAS. Each PPP tunnel is separated from other PPP tunnels, and because the Ethernet access node will only accept PPPoE frames, the upstream traffic can only flow to the BRAS. Therefore, no two end-users can obtain layer-2 visibility.

For business subscribers, a VLAN can be allocated to a specific end-user (the network of a branch, regional or corporate office) to create an isolated layer-2 network that can cross the IP core network in a layer-2 or layer-3 virtual private network (VPN). The traffic—which could be from a corporate intranet, circuit emulation services over IP (CESoIP) for PBX interconnection, or Frame Relay—is carried transparently through the layer-2 network (Figure 5).

Quality of service

IEEE 802.1p, which is part of the IEEE 802.1Q standard, defines the Ethernet pri-

Figure 3
Overview of VLAN per service.

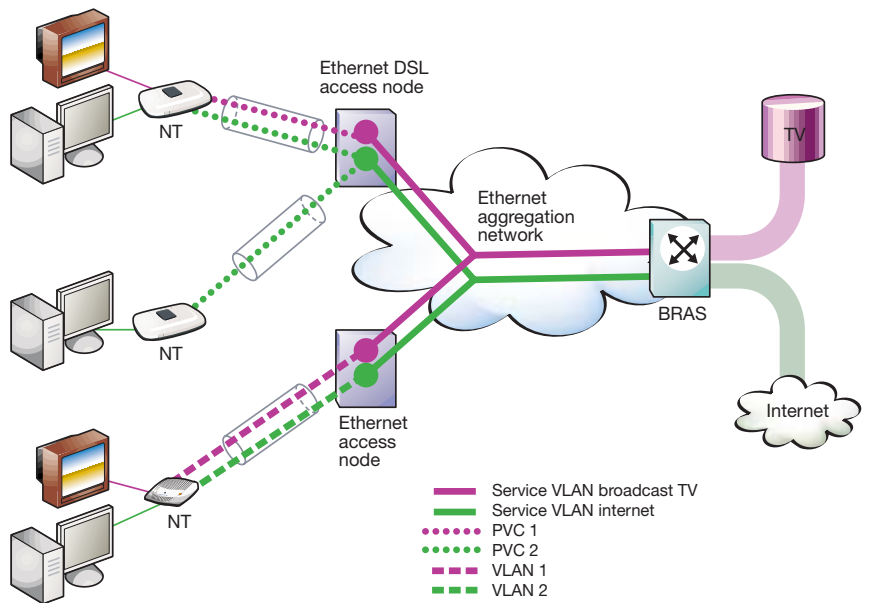

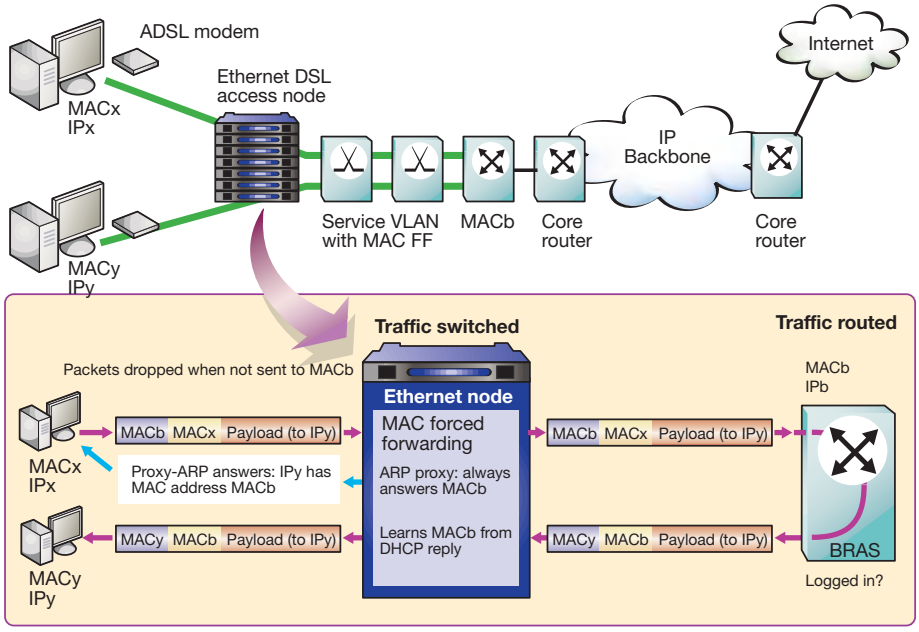

Figure 4
Overview of the service VLAN and MAC forced-forwarding mechanism.

ority bits used for QoS over Ethernet networks, stipulating eight levels of priority that can be assigned to a VLAN. If congestion develops in the network, the Ethernet switches first drop Ethernet frames marked with lower priority. Likewise, if bandwidth shaping or policing are applied, then it is possible to define a maximum and minimum bandwidth for the VLAN or end-user. Bandwidth shaping buffers traffic to compensate for traffic bursts, keeping traffic levels within the limits defined for the VLAN.

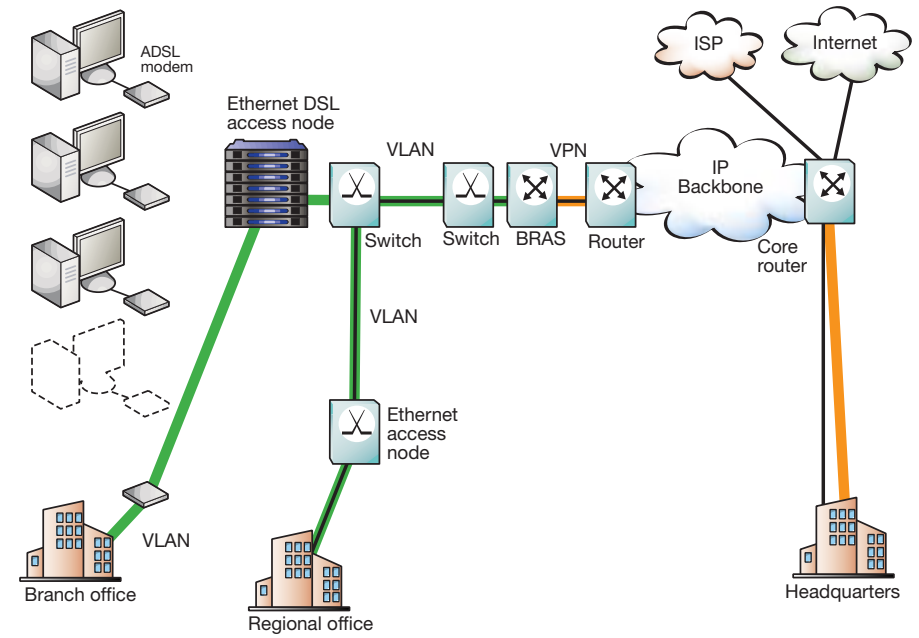

Figure 5
Overview of VLAN per customer and transparent connection.

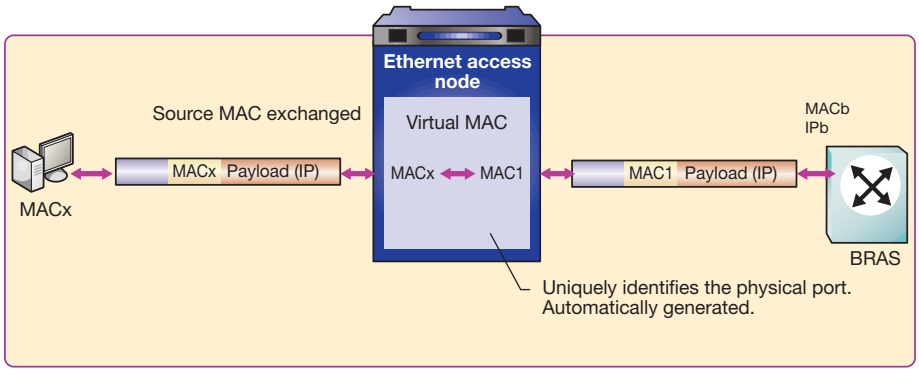

Figure 6
Overview of the virtual MAC address function.

The technique reduces traffic loss, but introduces delay and jitter into the traffic stream.

Bandwidth policing causes packets to be dropped as soon as the limit defined for the VLAN is reached. More traffic is lost, but traffic that does not exceed bandwidth limits experiences low delay and jitter. This is an attractive option for voice over IP (VoIP) and video traffic. In summary, a combination of IEEE 802.1p and bandwidth shaping or policing makes it possible to assign a QoS profile to a VLAN (the end-user's logical connection that carries the service).

When end-user traffic reaches the IP core network (the BRAS or edge router), the different IEEE 802.1p priorities are mapped into differential services (DiffServ) classes or multiprotocol label switching (MPLS) experimental bits to maintain the priority of the traffic as it is routed through the IP core network.

Figure 7
Overview of DHCP option 82.

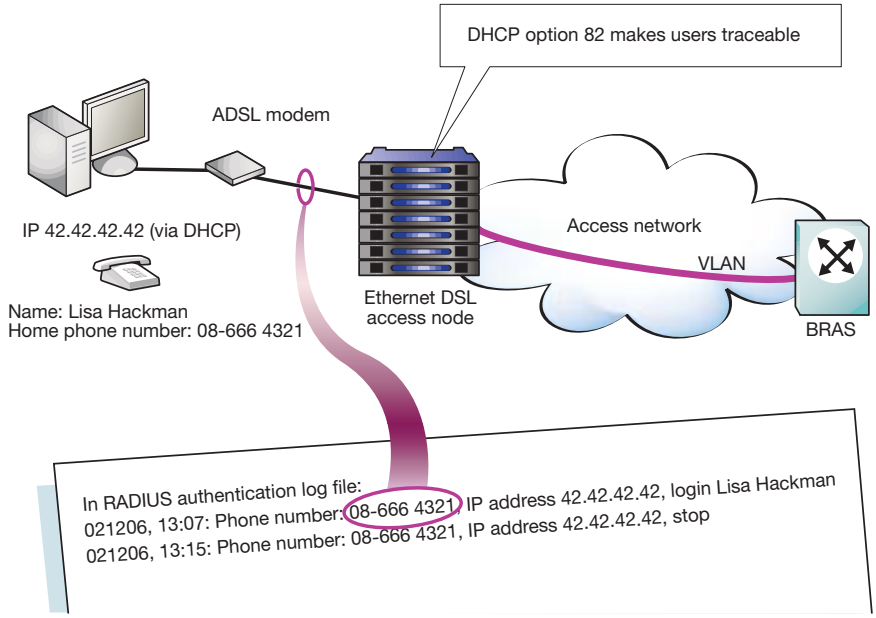

Security

One method of identifying end-users who have accessed the network is to record their MAC addresses. However, this method has some fundamental drawbacks—for instance, it is difficult to trace a MAC address to a specific end-user, end-users can easily modify their MAC addresses, and multiple devices can appear in the Ethernet access network with identical MAC addresses, either intentionally (spoofing) or by accident.

To counter these problems, the Public Ethernet network solution makes use of virtual MAC addressing. In the Ethernet access node, the end-user MAC address is mapped to a unique, locally generated MAC address (virtual MAC). This is used within the access network as an alias for the end-user's MAC address. The algorithm used to generate the virtual MAC address identifies the port and logical connection to the end-user. Consequently, the virtual MAC address, which is outside the control of the

end-user, can be used to identify the exact source of all upstream traffic. In addition, the function enables operators to protect the network from malicious end-users, by blocking their attempts to broadcast multiple MAC addresses into the network to overload the MAC address tables in the Ethernet switches (Figure 6).

Operators who use the dynamic host configuration protocol (DHCP) to assign IP addresses to end-users can enhance this function using DHCP option 82. When end-users request an IP address, the Ethernet access node detects the requests and modifies them by adding a string that identifies each end-user. The Ethernet access node can automatically generate this string to identify the end-user port from which the request originated. Alternatively, the operator can define the string to use via the management system. The BRAS can then send information such as time, date, IP address and identifier to the RADIUS log to keep a record of which end-users logged onto the network, and from which ports the end-user traffic originated. The operator can thus trace end-user activities if necessary (Figure 7).

Filtering rules implemented in the Ethernet access node provide additional security. For instance, it is possible to set up rules for processing traffic on layer 2 or layer 3. To prevent IP spoofing,the Ethernet access node can be configured to solely accept traffic from the IP address or IP subnet assigned to the end-user. The Ethernet access node learns the assigned IP address or IP subnet by snooping the DHCP reply sent from the DHCP server. For end-users with a static IP address, it obtains them from the management system.

Multicast

There are several broadcast services, such as TV and radio, that send the same data stream to multiple end-users simultaneously. Two methods are used for delivering the data stream: unicast and multicast. With unicast, an end-user requests access to the data stream and it is sent to him. Although simple to implement, the unicast solution is not scalable because the same data must be sent 100 times to 100 different end-users (Figure 8).

With multicast, a single data stream is transported through the network as close to end-users as possible before it is duplicated and delivered. This approach saves bandwidth in the transport network and reduces load on the content server, providing a much

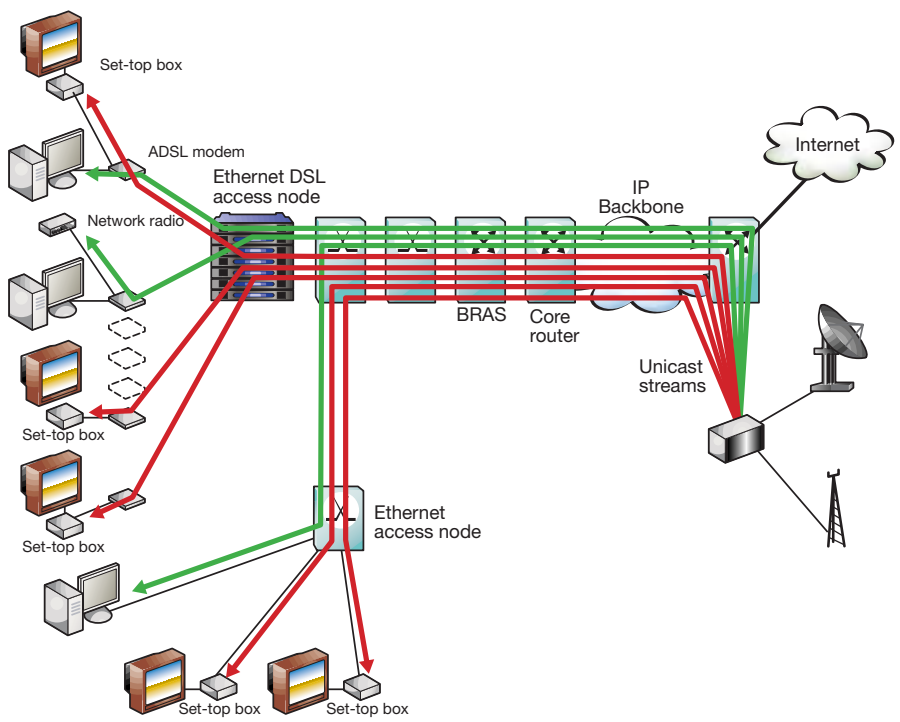

Figure 8
Overview of the unicast implementation.

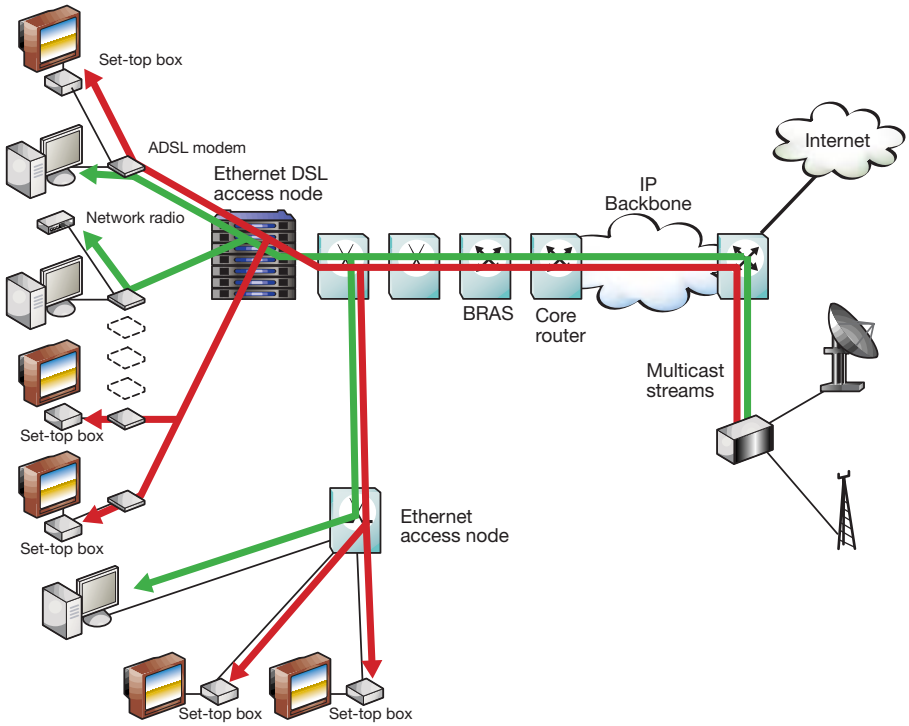

Figure 9
Overview of a multicast implementation.

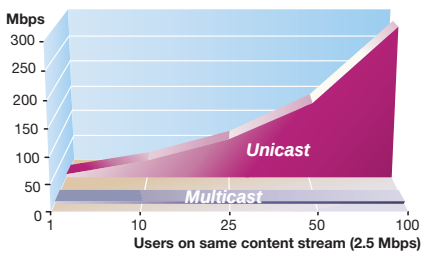

Figure 10
Unicast bandwidth increases uniformly with more users, whereas multicast bandwidth remains constant.

more scalable and cost-effective solution for delivering broadcast services (Figures 9 and 10).

For multicast to work, a protocol is used to determine the multicast group, so the router or switch can determine which traffic to forward to each end-user. The Public Ethernet network uses the Internet group management protocol (IGMPv2 RFC 2236). End-users send join and leave messages to tell the network which multicast session they want to participate in. The Ethernet DSL access nodes and Ethernet switches in the aggregation network snoop the IGMP requests to detect and maintain the multicast groups.

The Ethernet access node can also keep a "white list" of the IGMP groups to which a given end-user is allowed to access. Operators can thus set the end-user port on the Ethernet access node to restrict access of a specific end-user to, say, IGMP groups 1-20 (for example, TV channels 1 - 20 in a basic subscription package). The Ethernet access node blocks any attempts by end-users to select IGMP groups to which they do not have access.

Robustness and availability

To fulfill service level agreements, prevent customer churn, and avoid poor reputation, operators must have a resilient network. Equipment can feature internal redundant hardware—some equipment is fully redundant, some only partially redundant. Network redundancy, which protects against the failure of a link or network element, is achieved by having alternative paths for end-user traffic through the network.

The use of two fibers between Ethernet switches provides 1+1 physical link redundancy with two physical links to carry traffic between Ethernet switches and with failover time of less than 50 µs.

The solution also employs rapid spanning tree, a relatively lightweight protocol specified in IEEE 802.1w that enjoys wide support in the industry. The protocol makes it

possible to redirect end-user traffic (typically within 5 seconds) via an alternative path when the primary path fails. It is a cost-effective redundancy option at the edge of the network (for example, at the Ethernet DSL access node) where nodes are often remotely placed with a small number of Ethernet switches, links and VLANs, and relatively few subscribers are affected by a single failure.

The Ethernet automatic protection switching (EAPS) protocol, specified in RFC 3619, works in networks built with a ring topology, where a master node redirects traffic in response to failure. EAPS has a failover time of less than 50 ms, providing the in-service performance required in the Ethernet aggregation network, where a single failure can affect a large number of users.

These redundancy mechanisms enable operators to build a Public Ethernet network with 99.999% availability.

Telecommunications management

Operators need to have a comprehensive telecommunications management solution in place before they can build out infrastructure, provide services, and effectively manage their networks. An element manager, called PEM, offers fault-management, configuration-management, performance-management and security-management functionality for managing the Ethernet access nodes.

Likewise, the Ethernet connection manager (ECM) is a tool for configuring VLANs and network resources end-to-end in the Public Ethernet network. It supports a wide range of network topologies and can be adapted to support equipment from a variety of vendors. The ECM collects a vast amount of data on the network and network topology in its database. This data also has a useful secondary value outside the ECM and can be used to assist operators in engineering traffic for network expansion and usage.

Finally Ericsson has a service offering to assist operators in rolling out and installing the Public Ethernet network as well as in integrating it into systems for service provisioning, network management, work practices and business support.

Transport solutions for the Public Ethernet network

Although the many advantages of Ethernet broadband access are compelling, operators with extensive ATM and synchronous digital hierarchy (SDH) transport networks must find a way of fully utilizing and migrating existing network architectures. This migration can be addressed in different ways, depending on the architecture and operator plans.

Ericsson offers a range of solutions that interface the Ethernet access network to legacy transport technologies, such as ATM, SDH and plesiochronous digital hierarchy (PDH). The product portfolio covers

- Ethernet over dark fiber;
- Ethernet over course wavelength division multiplexing (C-WDM);
- Ethernet over radio link—point-to-point or point-to-multipoint;
- Ethernet over SDH or PDH, either as E1/T1, E3/T3, or STM-1/OC3;
- Ethernet over the next-generation SDH network; and
- Ethernet over ATM.

Migration of traditional TDM- and ATM-based services to Public Ethernet

Although the Public Ethernet network opens up new and exciting opportunities, legacy applications based on E1 leased lines and Frame Relay connections will continue to exist. To cater for these applications, transparent LAN connections (VLAN) are created in the Ethernet access network, and equipment is placed at the subscriber's premises (customer premises equipment, CPE)—the CPE converts subscriber traffic to circuit-emulation services over IP (CESoIP, structured and unstructured) or Frame Relay over IP (FRoIP). These services are then carried transparently through the Public Ethernet network with high priority and minimum delay and jitter.

Conclusion

The continuing evolution of broadband access will create numerous challenges for service providers, who in an increasingly competitive environment must drive down the cost of delivering existing services and simultaneously open the door to new revenue-generating services.

Ericsson's Public Ethernet network solution has been designed and developed to meet these challenges. By drawing on standard Ethernet technology, Ericsson can offer a broadband access network with an unbeatable price-to-performance ratio, while maintaining the carrier-class functionality, security, quality of service and scalability of traditional telecommunications networks.

REFERENCES

1 ITU-T Recommendation G.902, Framework recommendation on functional access networks (AN). Architecture and functions, access types, management and service node aspects.
2 IEEE 802.3, Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications (Ethernet)
3 IEEE 802.1Q, Virtual LAN
4 IEEE 802.1p, Traffic Class Expediting and Dynamic Multicast Filtering
5 IEEE 802.1w, Rapid Spanning Tree
6 IETF RFC 2364 PPP over AAL 5 (PPPoA)
7 IETF RFC 2516 (Feb 1999), A Method for Transmitting PPP Over Ethernet (PPPoE)
8 IETF RFC 2684 (September 1999), Multiprotocol Encapsulation over ATM Adaptation Layer 5
9 IETF RFC 2236, Internet Group Management Protocol, Version 2
10 IETF RFC 2662, Definitions of Managed Objects for ADSL Lines
11 IETF RFC 1757, Remote Network Monitoring Management Information Base (RMON MIB)
12 IETF RFC 3619, Ethernet Automatic Protection Switching
13 ITU-T Recommendation G.992.1, Transmission Systems And Media - Asymmetrical Digital Subscriber Line (ADSL) Transceivers