

Modularity is key when designing packet backbone networks for mobile services

Araceli Calle, Alberto Fernandez Bravo, Mounir Merhi, Jens Poscher and Helena Stjerna

Ericsson has developed a verified reference network design that optimally integrates general packet radio service (GPRS) and wideband code-division multiple access (WCDMA) technology with site and backbone IP infrastructure. The solution, called Mobile Packet Backbone Network, or Mobile-PBN, unites Ericsson's core network products, and includes, among other things, optimum designs for the service network, the network management center, and network synchronization. The solution is offered as part of Ericsson's network design services.

The authors introduce the modular Mobile-PBN concept and describe some of its key modules—the multiservice backbone, circuit-switched layered architecture, and service network. They also describe the challenges associated with verifying the Mobile-PBN and how these have been solved in the context of an end-to-end network. The Mobile-PBN frequency synchronization solution has been described in a separate article.¹

Background

The introduction of new access and core technologies as well as the convergence of packet- and circuit-switched traffic in mobile networks is a challenge for mobile operators, because they are responsible for network infrastructure and must deal with net-

work complexity, diverging requirements and node limitations.

All too often, new nodes are introduced into operator networks using *ad hoc* network design with the aim of making the system work quickly. But the inevitable result of improper network design is an inflexible and unscalable network that cannot cope well with even minor changes. To avoid a complete redesign (which could jeopardize past investments) after the network has been taken into operation and the subscriber base is growing, operators are often forced to accept costly add-ons and workarounds.

Ericsson has collaborated with third-party networking vendors, such as Juniper Networks, Extreme Networks, NetScreen Technologies (now part of Juniper Networks), and F5 Networks, to tackle these issues and provide a strong solution—the Mobile-PBN reference design.

Operator benefits

The Mobile-PBN solution has been optimized and verified to work for networks of any size, including multi-site networks. Of-

BOX A, TERMS AND ABBREVIATIONS

3GPP	Third-generation Partnership Project	LSP	Label-switched path	RNSAP	RNS application part
A2EA	AAL2 service end-point address	MAP	Mobile application part	RSVP	Resource reservation protocol
AAL2	ATM adaptation layer-2	MGW	Media gateway	RSVP-TE	RSVP with traffic engineering extensions
AMR	Adaptive multirate	M-MGW	Mobile MGW	SAPI	Service access and protection infrastructure
ATM	Asynchronous transfer mode	MP-BGP	Multiprotocol BGP	SDH	Synchronous digital hierarchy
BGP	Border gateway protocol	MPLS	Multiprotocol label switching	SGSN	Serving GSN
BICC	Bearer-independent call control	MSC	Mobile switching center	SGW	Signaling gateway
CAMEL	Customized applications for mobile network-enhanced logic	MTP3	Message transfer protocol 3	SIGTRAN	Signaling transport
CAP	CAMEL application part	NAT	Network address translation	SLB	Server load balancer
CAPEX	Capital expenditure	NTP	Network time protocol	SNF	Service network framework
CoS	Class of service	O&M	Operation and maintenance	SN-IP	Service network IP infrastructure
Diffserv	Differentiated services	OPEX	Operating expenditure	SONET	Synchronous optical network
DNS	Domain name server	OSPF	Open shortest path first	SPF	Shortest path first
DSCP	Diffserv code point	PAT	Port address translation	SS7	Signaling system no. 7
FRR	Fast reroute	PBN	Packet backbone network	STM-1	Synchronous transfer mode data rate 1 (155.52 Mbps)
FWLB	Firewall load balancer	PCM	Pulse-code modulation	STP	Signal transfer point
GCP	Gateway control protocol	PDB	Per-domain behavior	TCP	Transmission control protocol
GGSN	Gateway GSN	PDU	Protocol data units	TDM	Time-division multiplexing
GMSC	Gateway MSC	PFE	Packet-forwarding engine	TrFO	Transcoder-free operation
GPRS	General packet radio system	PHB	Per-hop behavior	TSC	Transit switching center
GSM	Global system for mobile communication	PLMN	Public land mobile network	UMTS	Universal mobile telecommunications system
GSN	GPRS service node	PoI	Points of interconnection	UTRAN	UMTS terrestrial radio access network
IGP	Interior gateway protocol	PoP	Point of presence	VPN	Virtual private network
iLB	Internal load balancer	PSTN	Public switched telephone network	WCDMA	Wideband code-division multiple access
IMS	IP multimedia subsystem	PVC	Permanent virtual circuit	WRED	Weighted RED
IP	Internet protocol	QoS	Quality of service	WRR	Weighted round-robin
ISDN	Integrated services digital network	RAN	Radio access network		
ISUP	ISDN user part	RANAP	RAN application part		
IT	Information technology	RE	Routing engine		
L2/L3	Layer-2/layer-3	RED	Random early discard		
		RNC	Radio network controller		
		RNS	Radio network subsystem		

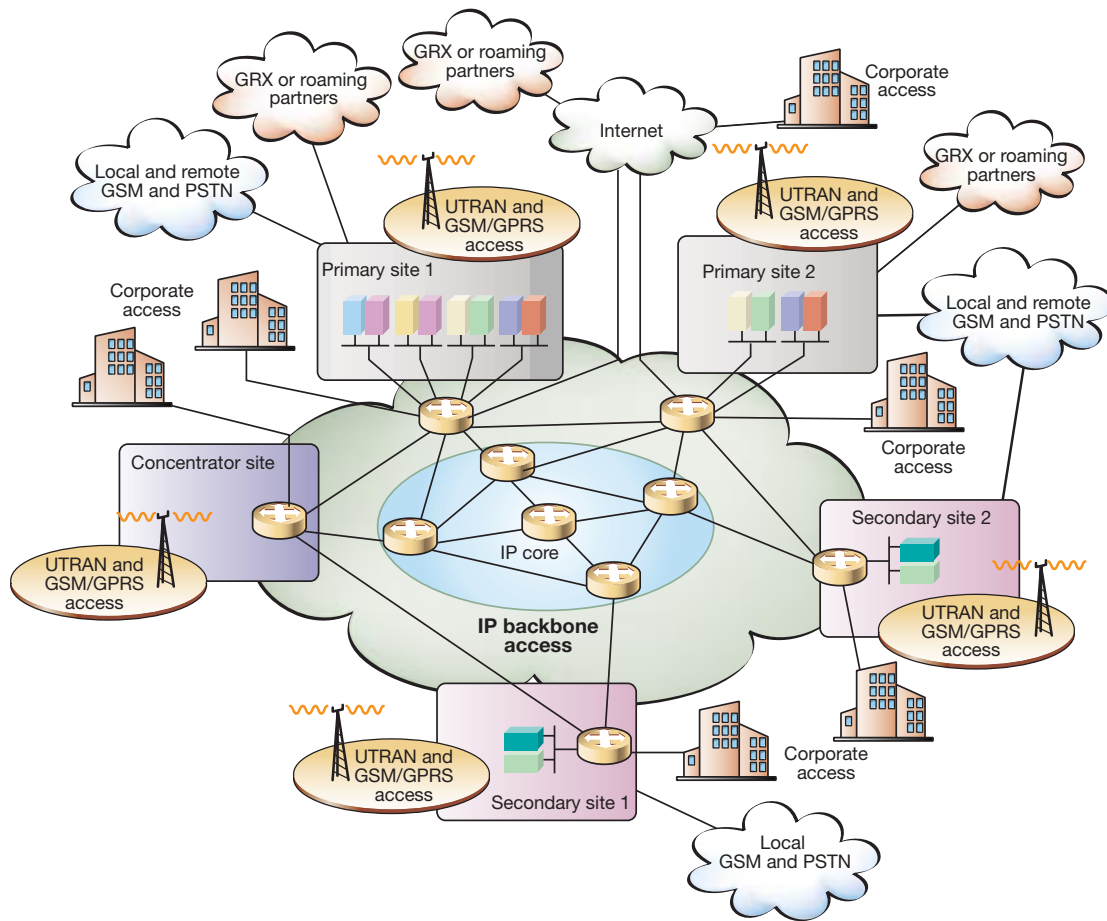


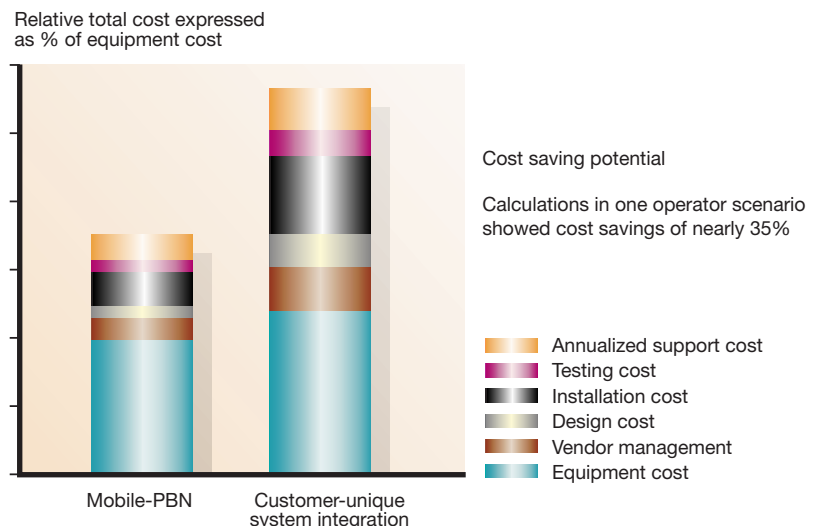
Figure 1
The Mobile-PBN reference network.

ferred as part of Ericsson's professional network design services, the solution

- reduces costs of network design, verification and type acceptance, which translates into shorter time to revenue (Figure 2);
- reduces risks—it has been verified, is scalable, and future-proof;
- relieves operators of the cumbersome task of network design—they can rely on Ericsson's professional services staff to adapt the solution to existing networks and instead concentrate more fully on crucial revenue-generating end-user services; and
- reduces supplier relationships—Ericsson owns the overall solution and assumes responsibility for designing, migrating and supporting the entire network, and guarantees integration into Ericsson core network and GPRS and WCDMA solutions.

In addition, operators receive a verified, highly flexible, and future-proof network

Figure 2
Benefits of Mobile-PBN: reduced time and costs of system integration.



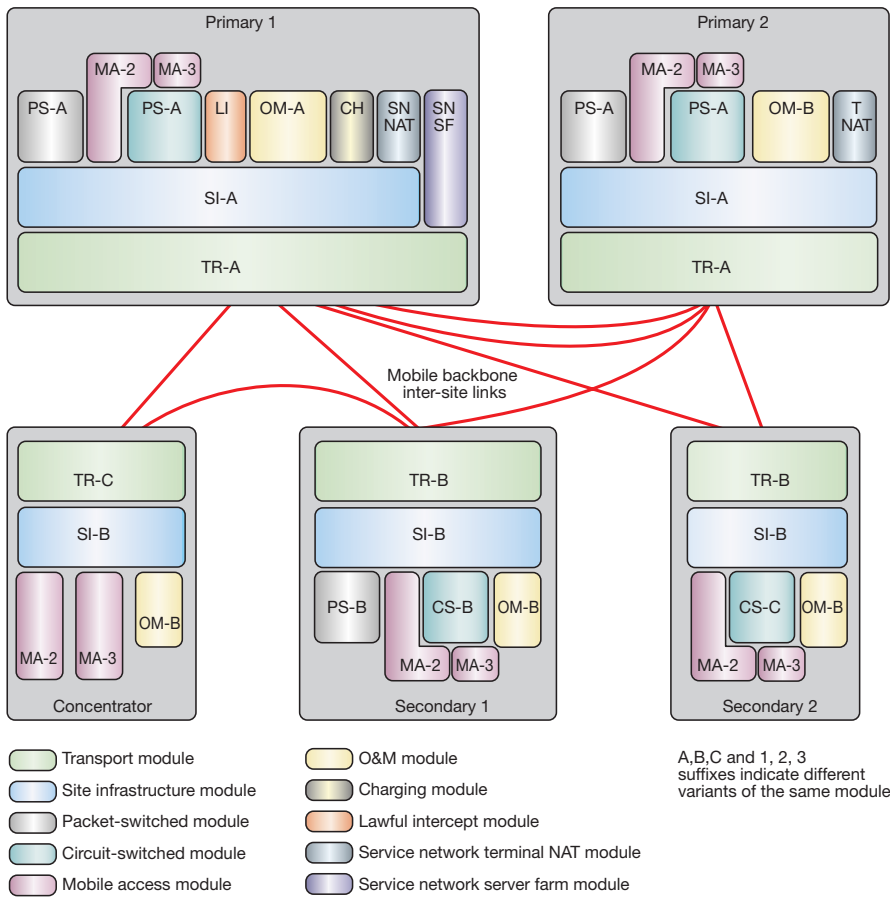


Figure 3
Mobile-PBN modules by site and function in a multi-site environment.

design to which any number of modules can be added—for example, corporate access, network management and data optimization.

Mobile-PBN, a total network solution

The Mobile-PBN solution consists of network modules that combine to provide network functions built on backbone and site infrastructure. The modules contain functionality for GPRS and WCDMA, circuit-switched voice, packet-switched data (including inter-PLMN roaming and corporate access), service network IP infrastructure (including Internet access), charging, data optimization, network management and lawful interception.

The site infrastructure provides dense Ethernet connectivity, switching functionality, frequency synchronization, network time protocol (NTP), and domain name server (DNS) services. The transport modules use IP routers or asynchronous transfer mode (ATM) switches to implement the core, distribution and access layers needed for inter-site connectivity.

Mobile-PBN reference design

The modular concept simplifies the definition of different kinds of site. The Mobile-PBN uses these definitions to design and dimension a detailed reference network for a “virtual operator” (Figure 3). The network consists of

- two Primary sites—these include the majority of modules, in particular, the operation and maintenance (O&M) module, which is needed for managing the complete network;
- two Secondary sites—these have fewer modules but maintain core network functionality; and
- a Concentrator site—this is mainly for mobile access and corporate connectivity.

The reference network is dimensioned to support a very large number of GSM and WCDMA subscribers. Besides the large reference design, the Mobile-PBN contains two very cost-effective designs for small, single-site networks. The first is a pure packet-switched solution that makes use of Ericsson’s combined GPRS service node (CGSN); the second is a split serving GSN and gateway GSN (SGSN-GGSN) design with added circuit-switched functionality. Each network includes advanced security protection using firewalls and can be migrated to the multi-site Mobile-PBN design.

The reference design meets operator requirements for availability, redundancy, security and scalability. Security considerations include a policy and multi-layered security architecture that defines logical security zones and areas of physical access. Traffic is not allowed to move from one zone to another without fulfilling the conditions of pre-configured firewall policies.

Multiservice IP backbone

Most operator networks already have, or will include, geographically distributed sites connected by a transport infrastructure. Although a variety of switching equipment can be used to transport different services,

there is a more profitable approach, namely common transport. Simplifying the network in this manner reduces capital expenditure (CAPEX) and eliminates the overhead of operating and maintaining multiple transport networks. It also significantly lowers operating expenditures (OPEX). The Ericsson Mobile-PBN facilitates convergence by providing backbone transport modules built around best-in-class IP routers that can support the most stringent services required by present-day and future GPRS and WCDMA networks.

Challenges of designing a packet backbone network

Interfaces

The first requirement put on a multiservice backbone is that it must be able to replace legacy network equipment. In this context, it is clear that the transport modules in the backbone must provide the same kinds of interfaces as offered by existing networks (Figure 4). The site routers of the Mobile-PBN have been selected to offer a rich variety of interfaces. They allow connections to ATM and Frame Relay and are frequently used to connect core nodes and peering networks. For Ethernet-based interfaces, the preferred means of aggregation, and of optimizing costs, is to use the Ethernet switches in the site infrastructure.

Traffic separation in VPNs

Unfortunately, multiservice backbone integration is not as simple as “plugging” client nodes into a new box. Different client networks need distinct transport services in separate network layers (Figure 5). For instance, in the Core Network 3.0 timeframe, Ericsson’s mobile media gateway (M-MGW) can deliver circuit-switched traffic in the form of IP packets (limited availability) or ATM adaptation layer-2 (AAL2) ATM cells. The latter solution requires functionality not found in ordinary IP backbones. The Mobile-PBN backbone provides this functionality using virtual private networks (VPN) based on multiprotocol label switching (MPLS).

MPLS is a key technology enabler, which in conjunction with other protocols, provides a common framework that supports Kompella layer-2 and RFC 2547bis layer-3 VPNs in an effective and scalable way. MPLS separates the control plane (which remains IP) from the forwarding plane, there-

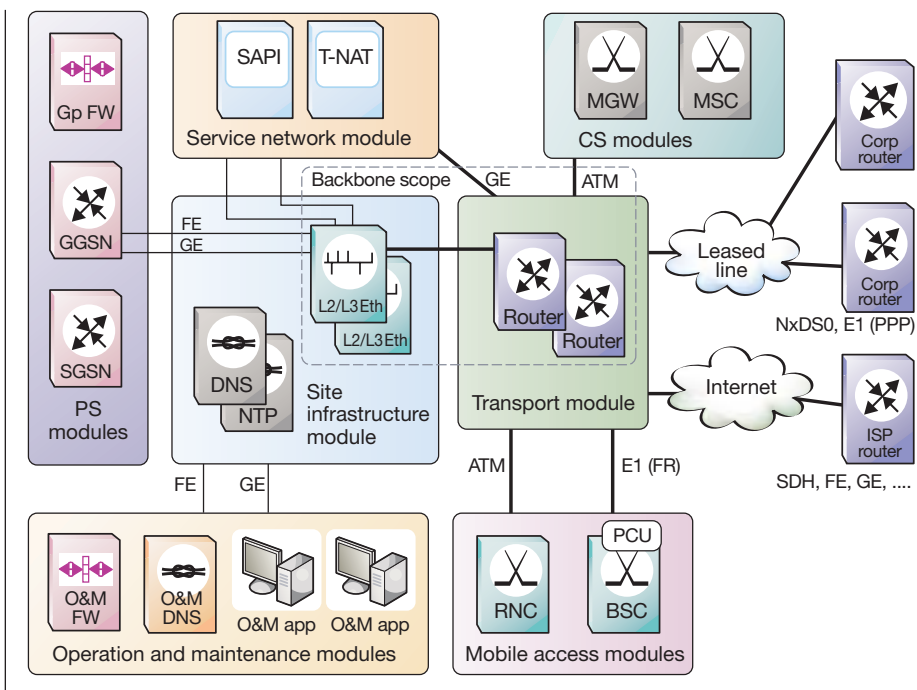
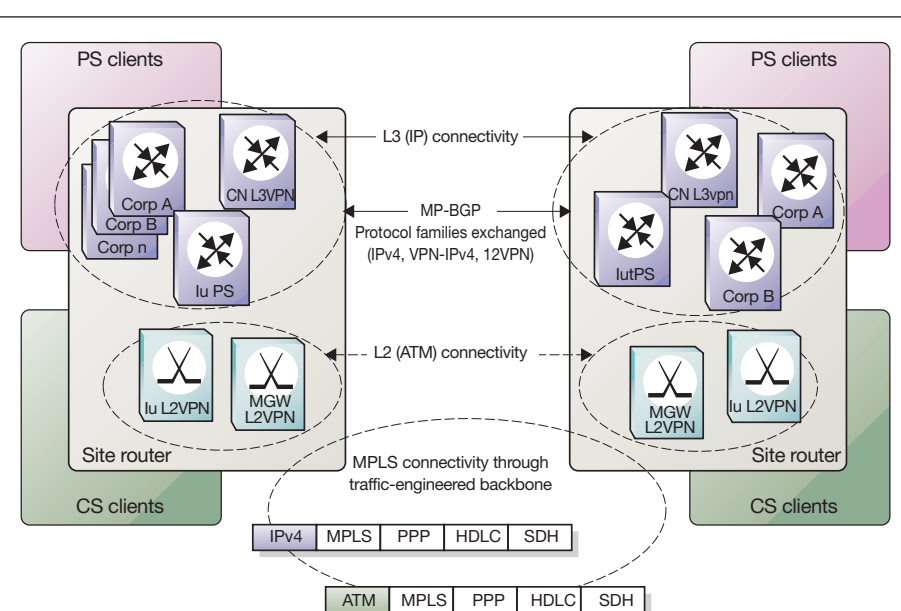


Figure 4 Mobile-PBN backbone interfaces in the Mobile-PBN.

Figure 5 Traffic separation in the backbone.



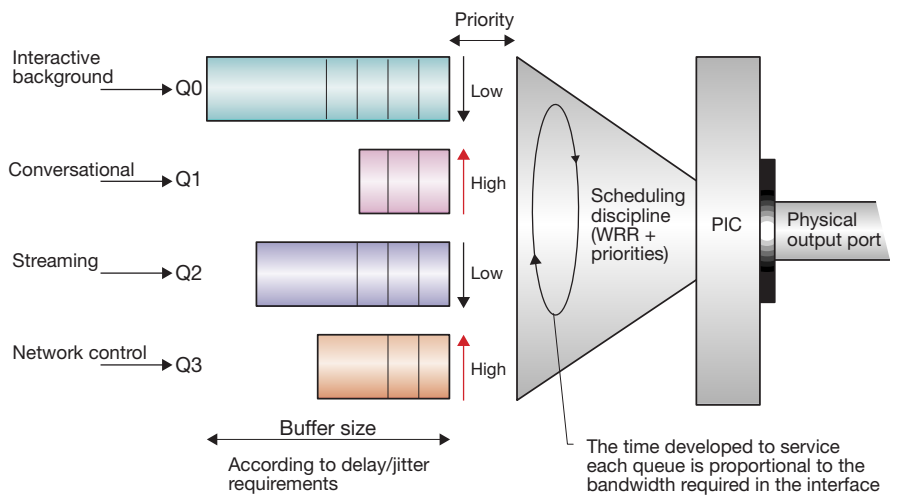


Figure 6 PHB implementation at site routers—queue, weight and priority configurations.

by solving the problem of forwarding non-IP protocol data units (PDU). Switching in the backbone is thus no longer dependent on looking up IP headers, because the forwarding of layer-2 and layer-3 (L2/L3) PDUs is based on MPLS labels, which are prepended to the PDUs when they enter the backbone. The core of the aforementioned control plane is based on the border gateway protocol (BGP) with multiprotocol extensions (MP-BGP) and traffic-engineering-enhanced link-state internal gateway protocol (IGP)—for example, open shortest path first (OSPF)—and finally, on the resource reservation protocol (RSVP).

The flexibility of a large-scale protocol like MP-BGP makes it possible to exchange VPN information (associated interfaces, encapsulation, topology, and so on) on the different VPN types. OSPF and RSVP-TE permit label-switched paths (LSP) to be implemented between network sites without the constraints imposed by shortest path first (SPF) algorithms inherent in link-state protocols. Thus, we see that transmission resources are used efficiently. Likewise, system availability is maintained because the SPF algorithms can quickly converge when the topology changes.

The ability to transport traffic on different layers is not the only benefit of VPNs. Addressing and numbering issues are also greatly simplified, because client networks perceive the backbone as if it were dedicated—a benefit that eliminates complex coordination between the client networks.

Traffic differentiation

The backbone is also greatly simplified thanks to a reduction in infrastructure and because traffic flows that were previously carried independently can now share network elements and transmission resources. If contention occurs, quality-of-service (QoS) mechanisms ensure that each type of traffic is forwarded according to its requirements.

The QoS solution for the backbone employs differentiated services (Diffserv) over MPLS. From a design point of view, the following points must be considered:

- ability of the backbone to implement forwarding treatments or per-hop behaviors (PHB) to which the traffic flows are mapped according to traffic class;
- ability of client nodes to communicate the correct QoS requirement; and
- mechanisms for dealing with local and backbone-wide congestion.

Implementing PHBs

The success of this type of solution relies on the selection of a scheduling algorithm that serves the queues used to implement the PHBs at each site router (Figure 6). The routers from Juniper Networks employ a combination of weighted round-robin (WRR) and priority disciplines, which define two priority levels. The approach upholds the strict delay and jitter requirements associated with voice traffic while giving fair treatment to lower-priority traffic classes.

Each PHB definition must also be complemented with buffer sizes that are appropriate for the delay or jitter characteristics of the traffic class to be served. In addition, for each PHB, a percentage of the available bandwidth must be allocated between each pair of sites according to relative traffic volumes. The coherent application of PHBs through the backbone routers leads to per-domain behaviors (PDB), which is to say that bearer services remain between required levels, regardless of which path is followed through the backbone.

Communicating QoS needs

Diffserv code points (DSCP) are used to map traffic flows to the appropriate PHBs. The DSCPs are inserted by the client node or can be allocated to the L2/L3 PDUs as they enter the backbone. Diffserv-capable nodes, such as GSNs, insert a DSCP that adequately reflects the PDP context. The common procedure, when the same organization owns the core and backbone networks, is to sup-

port this DSCP from an end-to-end perspective. If the IP nodes cannot handle Diffserv (for example, P2.1 RNCs), then the ingress site router inserts the DSCP based on a multi-field discriminator filter configured in the site router interface that connects to the client node.

Non-IP client nodes, such as media gateways that deliver AAL2 ATM cells, can use the site router resources by mapping the connecting interface (logical or physical) to the proper forwarding class in the ingress router. This mapping is defined when the interface is configured.

In every case, Diffserv treatment in the MPLS backbone network is encoded in the MPLS header using E-LSPs. This means that the applicable quality of service is derived from the experimental bits in the MPLS header (Figure 7).

Dealing with congestion

To work effectively, the QoS approach requires mechanisms for dealing with local and backbone-wide congestion. The mechanism for managing local congestion, called weighted random early discard (WRED), enables different drop profiles to be assigned to different kinds of traffic. A less aggressive random early discard (RED) drop profile is assigned to critical packets; a more aggressive profile is assigned to other packet types.

With the exception of transmission control protocol (TCP) traffic, this mechanism by itself is not sufficient to avoid long-term congestion. A backbone-wide solution which is based on traffic conditioning at the edges of the backbone and which is valid for all kinds of traffic must be employed to keep traffic injected by different sources within acceptable limits. Ericsson recommends that policers should be employed in the site router interfaces connected to client nodes. The policing mechanism in routers from Juniper Networks employs the token bucket algorithm, which enforces a limit on average bandwidth while allowing bursts of up to a specified maximum value.

Availability

The infrastructure must also offer the same level of availability found in existing public telecommunications networks. The Mobile-PBN view is that reliable services are built on top of reliable networks, which in turn, are based on reliable platforms.

The Mobile-PBN recommends the usage of IP routing devices with separate routing

and forwarding planes. This gives efficient deployment of graceful restart implementations of the protocols used in the control plane. Separation guarantees the performance of the packet-forwarding engine (PFE) even with high levels of route instability. Moreover, even extremely large volumes of traffic cannot limit the ability of the routing engine (RE) to maintain peer relationships and calculate routing tables. In short, a clean separation of these two functions yields superior forwarding performance and a highly reliable system.

Equipping the routers with redundant routing engines and packet-forwarding engines further increases availability. The routing engines support hitless failover, which means that in the event of failure or planned downtime (for software upgrades) the packet-forwarding engine can maintain forwarding while the redundant routing engine takes over (for a period of, say, 2-3 minutes).

The introduction of robust routers does not, in itself, guarantee the required levels of reliability—the backbone must still be protected against link failures and the risk of node failures. Because the Mobile-PBN is a converged MPLS infrastructure, MPLS traffic protection mechanisms are also necessary (Figure 8). To meet the stringent requirements imposed by voice traffic, it is

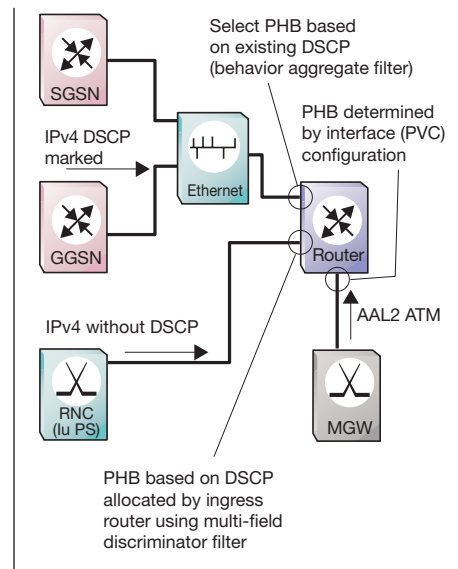
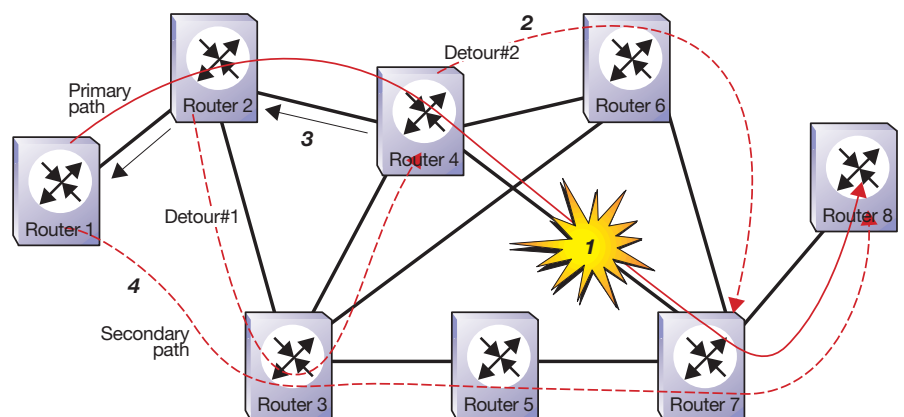


Figure 7 Determining QoS for different types of client nodes.

Figure 8 MPLS traffic-protection mechanisms applied to the backbone.



- 1 Link between routers 4 and 7 falls affecting LSP primary path
- 2 Router 4 switches traffic to detour#2
- 3 Router 4 notifies ingress router (router 1) about the failure
- 4 Router 1 switches traffic to LSP secondary standby path




Figure 9
Schematic view of the layered architecture.

ternative paths that the router upstream of an outage can use to quickly reroute traffic around a failed link or node. The same router notifies the ingress router about the failure, causing it to switch traffic to the secondary path. The use of pre-computed alternate paths significantly reduces recovery times. Indeed, it puts them on a par with SDH/SONET protection mechanisms. To be effective, the secondary paths and other alternative paths must not share fate with the primary label-switched path.

In large networks, the combination of these two mechanisms (FFR and secondary standby paths) provides optimum traffic protection. In networks with simpler topology, such as the Mobile-PBN reference network, secondary standby paths suffice.

To obtain the best results when applying these mechanisms, Ericsson recommends compliance, at physical and logical levels, with a set of connectivity guidelines for the backbone and different client networks.

recommended that two levels of pre-computed label-switched paths be used. At local levels, the fast reroute (FRR) mechanism is used. At backbone-wide levels, secondary standby paths are recommended.

The fast reroute mechanism provides al-

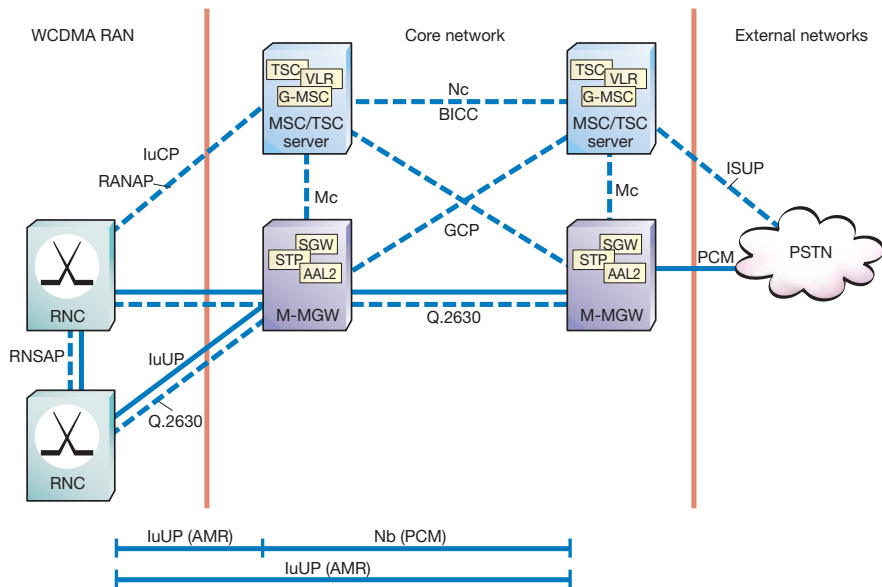
Circuit-switched layered architecture

The Third-generation Partnership Project (3GPP Release 4) has standardized the circuit-switched layered architecture in response to operator requirements

- to cut costs in the transmission network;
- to aggregate O&M-intensive nodes in central sites; and
- to spread the public switched telephone network (PSTN) points of interconnection (PoI) close to the destination, to obtain local switching.

Ericsson's solution to the circuit-switched layered architecture, which is based on mobile switching center (MSC) servers and mobile MGWs, enables operators to introduce the new technology using ATM in the transport network (for the user-plane) and signaling system no. 7 (SS7) signaling. A step-by-step evolution simplifies migration of the core network to IP. The first step entails the introduction of SS7 over IP signaling transport (SIGTRAN) for control signaling. Later, the user-plane in the core network will also be based on IP transport.

Figure 10
Interfaces used in the layered architecture.



Core network architecture

The control and connectivity layers in the circuit-switched layered architecture are logically separated. However, the servers (MSC, GMSC, TSC and SGW) retain the signaling protocols used in the non-layered architecture. Two additional protocols—

the bearer-independent call control (BICC) and gateway control protocol (GCP)—are also implemented (Figure 9). BICC introduces control of ATM and IP bearers in the connectivity layer. GCP enables the servers to control and manipulate the resources of a media gateway. Legacy narrowband protocols might also need a signaling gateway (SGW) to convert signaling bearers.

Benefits

The new architecture helps reduce the costs of operating the network. In traditional networks, the MSC and transit switching center (TSC) nodes are not centralized. Consequently, daily maintenance and hardware and software upgrades require O&M personnel to work from several locations.

Leased lines in the TDM-based core network are dimensioned for 64 kbps voice channels. Existing voice compression systems are not end-to-end solutions. However, setting up a network in the layered architecture to place the voice transcoders at the edge of the network permits true end-to-end transmission of coded voice to save bandwidth in the backbone. In addition, voice quality enhancements, such as transcoder-free operation (TrFO), which was standardized in 3GPP Release 4, will soon become available.

Another benefit of the new architecture is the decentralization of media gateways, which allows the opening of multiple PSTN points of interconnection and enables operators to carry calls as far as possible in their own networks.

Technology

The initial solution for the layered core network is based on an ATM network with AAL2 switching. The functionality and benefits of AAL2 switching with statistical multiplexing in the UMTS terrestrial radio access network (UTRAN) have been discussed previously.² Figure 10 shows all relevant payload and signaling interfaces.

The Ericsson mobile media gateway combines all transcoder, echo canceller, and multi-part device functionality. A signaling gateway is added on top of the mobile media gateway to convert the SS7 signaling bearer from ATM to TDM and vice versa. The M-MGW serves as a cross-connect for ATM permanent virtual circuits (PVC) and as an AAL2 switch. AAL2 switch functionality is especially important for the codec-at-the-edge feature, and for redundancy at higher layers.

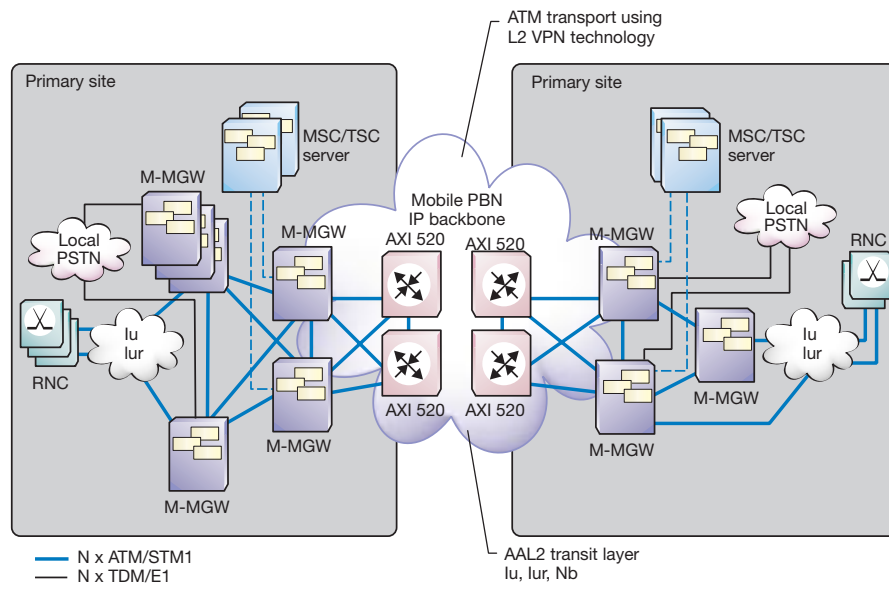


Figure 11
Mobile-PBN circuit-switched reference design.

Thanks to AAL2 switching and dedicated bearer control protocols (Q2630 or Q.AAL2) in the transport layer, the *Iu* interface can be extended through the core network domain to serve the nearest M-MGW of the region to which a call is to be delivered. The radio network controller (RNC) sets up the AAL2 connection hop-by-hop to the final M-MGW. It uses the AAL2 service end-point address (A2EA) to address the destination node. Only one media gateway is needed to transcode AMR-coded voice over *Iu* from the RNC to pulse code modulated (PCM) voice. Operators thus save bandwidth in the backbone—PCM requires the equivalent of 84.8 kbps in bandwidth, whereas AMR-coded voice over *Iu* requires only 12.2 kbps end-to-end. The same AAL2 network also facilitates *Iur* connections between RNCs. The next step is to build a core network that is as resilient and delivers the same perceived quality of service as TDM-based networks.

Mobile-PBN circuit-switched design

The Primary sites aggregate the MSC and TSC servers as control layer nodes. The MSC servers control the RNCs. For optimal call routing, the MSC servers also control every M-MGW in the network. The Secondary sites only contain M-MGWs. The MSC




Figure 12
AAL2 connections.

servers in the Primary sites control the RNCs and MGWs in the Secondary sites via GCP. The Concentrator sites solely contain the WCDMA access module.

The sites are interconnected by the multiservice backbone. The different sites tunnel ATM traffic through using layer-2 VPNs. This modular approach enables Mobile-PBN networks to scale from one-site solutions to very large WCDMA networks.

Connectivity

Ericsson's layered architecture solution calls for static ATM PVCs between each node in the AAL2 network. Each ATM PVC carries one AAL2 path, which in turn, can carry 248 AAL2 connections or calls (Figure 12).

Obviously, given the growing number of M-MGWs, a full mesh of physical links would be uneconomical and difficult to manage. An alternative approach is to connect the M-MGWs by means of cross-connected ATM PVCs, but this is also quite complex. Therefore, to achieve full connectivity of M-MGWs and the required redundancy on the AAL2 switching layer, Ericsson has introduced access layer and transit layer hierarchy.

The concept is similar to legacy TDM or IP networks. Two M-MGWs at each site belong to the AAL2 transit layer. Every AAL2 transit M-MGW is fully meshed over the

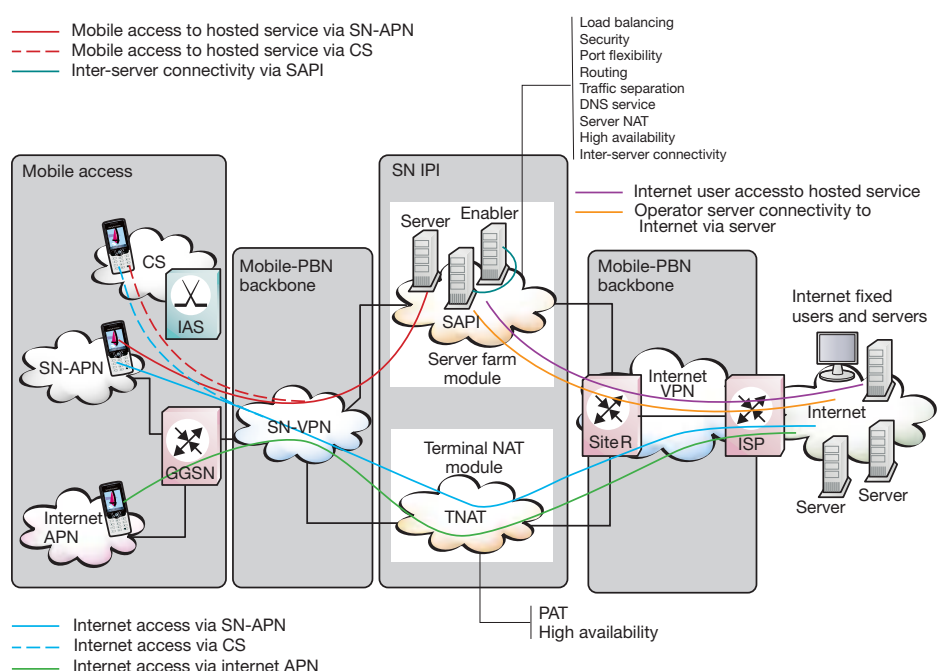
backbone using label-switched paths. Therefore, should the AAL2 path or node fail, the destination site can still be reached. The backbone is responsible for internal AAL2 path redundancy (Figure 8).

The remaining M-MGWs at the site connect on the AAL2 layer to the two transit M-MGWs. From the viewpoint of the AAL2 layer, the RNC must adhere to this same structure—that is, each RNC must connect to two different M-MGWs over redundant links. Thanks to this modular architecture, every AAL2-based interface (*Nb, Iu, Iur*) can use this transport layer. If a link, board or the next AAL2 switch fails, an alternate route is always available.

This strictly hierarchical network also allows certain nodes to be grouped in A2EA addressing—for example, the nodes at a site. This, in turn, makes it possible to use single routing entries for sites, because the transit AAL2 switches must only analyze the longest prefix in the address field. What is more, the addition of nodes only requires changes in a few routing tables.

To guarantee redundancy, the MSC and TSC servers must connect to two M-MGWs, which also serve as signal transfer points (STP). The internal STP functionality of the M-MGW forms the new broadband SS7 network. The broadband SS7 layer mirrors the hierarchical structure of the AAL2 layer. The transit M-MGWs, which make up the

Figure 13
Role of SN-IP infrastructure in mobile networks.



transit SS7 STP network, route all inter-site signaling (RANAP, RNSAP, GCP, MAP, CAP, BICC and ISUP) on the message transfer protocol 3 (MTP3) level.

IP infrastructure of the service network

The service network framework (SNF) is an architectural framework that consists of reusable designs for products and solutions in the service layer.³ This includes the IP network used for deploying the service network.

The SNF deployment view provides guidelines for ensuring that the IP network contains a set of common services and qualities on which every deployed system can rely. Examples of common services include naming, addressing, routing, load-balancing, firewall, and security gateway services. Common qualities include performance, scalability, flexibility, security and high availability. The Mobile-PBN thus employs a service network IP infrastructure (SN-IPI) that gives users of the Mobile Internet access to services hosted at the server farm, and mobile users access to the Internet.

The SN-IPI introduces two functional modules: the server farm module and the terminal NAT module. The server farm module provides services to mobile and In-

ternet users. It contains the service access and protection infrastructure (SAPI), service enablers, and application servers located in the mobile operator's service network. The terminal NAT modules provide the network address translation (NAT) functionality needed to map a private address to a public IP address when mobile subscribers want to access the Internet. As seen in Figure 13, the Mobile-PBN backbone makes use of SN-IPI modules and IP transport services to accommodate several traffic flows.

Server farm modules

In the context of the SN-IPI, the main system in a server farm module is the SAPI (Figure 14). The other systems of the server farm (enablers, application and content servers) are integrated into the SAPI, which consists of a set of firewalls placed between two pairs of load balancers—in this case, firewall load balancers (FWLB). This setup ensures incremental scalability. In other words, operators can add more firewall protection as capacity increases. The internal load balancers (iLB) also serve as load balancers for servers that scale horizontally. L2/L3 switches are used to provide flexible, high-port density to connect all enablers and services. Figure 14 shows how these components are interconnected.

Apart from providing physical connectivity and protecting servers, the SAPI also

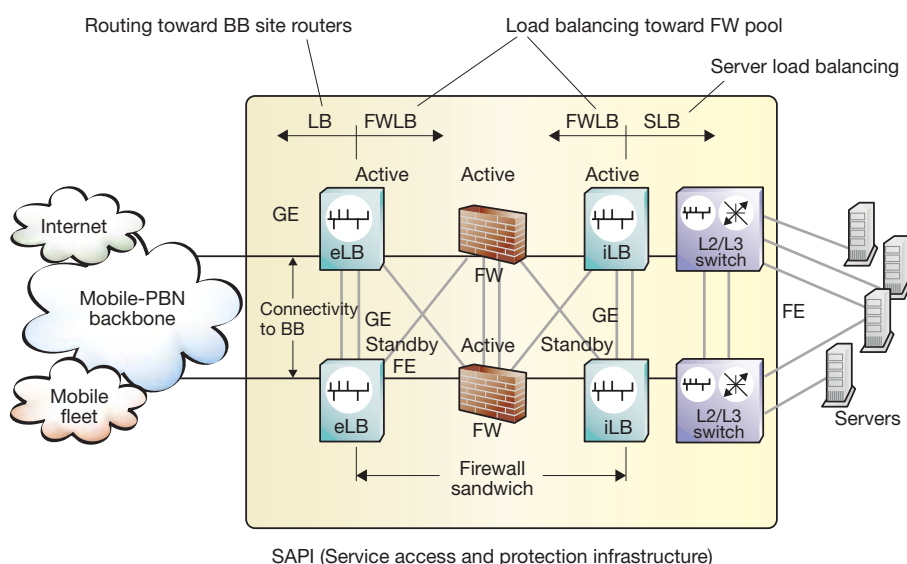


Figure 14 Service access and protection infrastructure.

provides address translation—it supports one-to-one and many-to-one address translation.

A small package solution is offered for smaller deployments of server farm modules. The module replaces external load balancers with standard L2/L3 switches. There are two variants for the internal server side. The first includes internal load-balancer units, used only as server load balancers and not as firewall load balancers. If necessary, standard L2/L3 switches can also be used to increase the number of ports or to connect to certain nodes, such as those that are based on TSP. The second variant includes internal load balancers. This solution is provided for operators who do not need load-balancing functionality.

The Mobile-PBN also provides support for integrating the following enablers and application servers into SAPI: USIS 1.0, DNS IPWorks 4.1, HTTP/FTP proxy (proxy + SSL), EMA 3.2, SNOS 1.0, MIEP 2.0 and MMS 3.0.

Terminal NAT module

The terminal NAT module consists of two switches that serve as NAT devices in an active-standby mode—if one of them fails, the other automatically takes over. In terms of deployment, the terminal NAT functionality is distributed in the same way as Internet point of presence (PoP). Ideally, a NAT

module should be deployed near the point of presence (PoP). This reduces the mobile operator's transmission costs by delivering traffic to the Internet as soon as possible.

Verification of the Mobile-PBN solution

Although the integration of new equipment and features is a critical issue in modern telecommunications networks, all planned and unplanned downtime disturbs active traffic and decreases revenue. To shorten integration times, the Mobile-PBN verifies the basic and new aspects of each design module before the design can be released.

The Mobile-PBN solution is set up and tested inside an end-to-end environment. The Ericsson Eurolab Deutschland GmbH verification center and the Ericsson AB Hot-Lab have several interconnected labs that provide the basic infrastructure for integrating every piece of equipment—site routers, switches, firewalls, and load balancers—with core network nodes (SGSN, GGSN, MSC, and M-MGW).

The modularity, flexibility and scalability of the Mobile-PBN are vital features for verification in the end-to-end test networks. These network characteristics ease the integration of a customized solution into an existing customer network.

TABLE 1, VERIFICATION ACTIVITIES COVERED BY THE MOBILE-PBN SOLUTION

- Connectivity and integration
- Security
- Resilience/redundancy
- Quality of service traffic management
- Operation and maintenance
- Load and characteristics
- Service network integration and the service network IP infrastructure with server farm and terminal NAT
- Data optimization
- Flexible bearer charging
- Small network (single-site solution)
- Network synchronization

Test cases are defined to expose potential problems and to verify design concepts before they reach the field. Verification reveals different kinds of faults or unexpected behavior in the equipment and network configuration. Because the network is operational, the test cases also show how it can be tuned for optimum performance. Traffic simulators and automated test patterns are used to simulate real radio network loading based on Ericsson's experience. These are used to stress the network to ensure correct behavior under load, which is essential for verifying operation of QoS delivery. Trouble Reports are written on all Ericsson and partner products. All problems and solutions are documented in test reports.

By discovering and resolving faults and unexpected effects, Ericsson makes the Mobile-PBN a more reliable network solution for the end-customer. Table 1 lists the test areas and extent of the Mobile-PBN verification activities.

Conclusion

This article demonstrates the strength of Ericsson's Mobile-PBN modular design concept:

- functional modules such as the circuit-switched (layered architecture) and service-network modules, can easily be

plugged into the transport modules (multiservice backbone and site infrastructure);

- other modules can be added easily when needed;
- the evolution of the backbone and site infrastructure is optimally aligned with the transport needs of the core network nodes (circuit- and packet-switched) and new packet-based services; and
- the network can readily be expanded using larger modules or adding new ones.

The design uses methods that cover a wide range of interconnected network areas and provide the following network-wide attributes:

- strong security at every level;
- class-of-service (CoS) differentiation;
- high availability (through physical and logical redundancy mechanisms);
- traffic separation (allowing overlapping IP address ranges for connected corporations); and
- intra- and inter-site connectivity.

A backbone and site infrastructure network design based on Ericsson's verified Mobile-PBN solution provides operators with many advantages—not only in terms of reduced costs and minimized risks, but also in terms of network interoperability, scalability and migration. With this solid foundation, operators can concentrate on providing attractive and cost-effective end-user services.

REFERENCES

1. The importance of network synchronization—Stand-alone products that support the design of synchronization networks, Ericsson Review no. 1, 2004
2. AAL2 switching in the WCDMA radio access network, Ericsson Review no. 3, 2002
3. The service network framework—An architectural blueprint for the service network, Ericsson Review no. 1, 2003
4. IP technology in WCDMA/GSM core networks, Ericsson Review no. 1, 2002