

Control servers in the core network

Andreas Witzel

The horizontally layered architecture introduced by Ericsson is playing a key role in the migration from earlier architectures into third-generation multiservice networks.

A central element in the horizontally layered architecture is the network control layer. Here, control servers perform call and session control.

In this article, which describes Ericsson's move away from the traditional, vertically integrated networks, the author gives a broad overview of the new architecture supplied by third-generation networks. He also emphasizes the central network control layer and the control server related to it, by giving detailed examples of typical traffic cases for the different domains in the network control layer.

Introduction

However stable it was throughout earlier decades, the architecture of telecommunication networks has become subject to pressure in recent years. Not only are new demands being put on the network, but also a changing environment and greater awareness on the part of users are forcing telecommunication network architects to change some fundamental concepts.

The first step was taken in the early 1990s, when the GSM standard emerged in Europe. Although mobile access had already been available via various analog systems, GSM was the first digital system that allowed roaming between different countries in Europe. Today, with the GSM standard spread around the globe, worldwide GSM roaming is becoming commonplace.

The GSM network was introduced in parallel to the existing networks for fixed line access and cable TV. At the same time, networks dedicated to packet data traffic also emerged. Starting with X.25 in early 1976,

development subsequently progressed toward frame relay, which was faster.

These networks existed in parallel—each of them tailored to its own specific purpose—and were run by specific operators. Although in the early days various national authorities controlled these operators, the trend toward deregulation in Europe opened up the market for many new network operators. It also unleashed fierce competition between them.

The deregulated market dropped prices for telecommunication services and boosted the number of subscribers and the volume of traffic. However, the growth of the network has been outrun by the growth of the Internet, which started and grew as a network of its own, mainly built on computer connections between different local area networks (LAN). If we put all these factors together, it becomes clear that the existing networks are vertically oriented, with almost no synergies or interworking between them.

Now that the third generation of mobile systems has emerged, it has at last become possible to combine computing, communication and broadcasting technologies. Like the fixed telephone networks, digital mobile systems started off as a conveyor of the human voice. However, along with unlimited mobility, the third-generation systems will offer a broad range of information and multimedia services.

If this is to be achieved in a way that is economically feasible, the vertical structure of the network architecture cannot be maintained; it must be replaced by horizontal layering, which represents a more generic approach. By being one of the initiators of this approach, Ericsson has succeeded in anchor-

BOX A, ABBREVIATIONS

3GPP	Third-generation Partnership Project	GTP-C	Gateway tunneling protocol, control	PSTN	Public switched telephone network
AAA	Authentication, authorization and accounting	GTP-U	GTP user plane	RADIUS	Remote authentication dial-in user service
API	Application program interface	HSS	Home subscriber server	RANAP	Radio access network application part
APN	Access point name	IB	Interface board	RTP	Real-time transport protocol
BICC	Bearer-independent call control	IMSI	International mobile subscriber identity	SCS	Service capability server
CAMEL	Customized applications for mobile network enhanced logic	IP	Internet protocol	SDL	Specification and description language
CSCF	Call state control function	ISDN	Integrated services digital network	SGSN	Serving GSN
DNS	Domain name server	ISP	Internet service provider	SIP	Session initiation protocol
DPE	Distributed processing environment	MAP	Mobile application part	SMS	Short message service
EIR	Equipment identity register	MGCF	Media gateway control function	SMS-C	SMS center
GCP	Gateway control protocol (for example, H.248)	MGW	Media gateway	SSF	Service switching function
GGSN	Gateway GSN	MIP	Mobile IP	TSC	Transit switching center
GMSC	Gateway MSC	MPB	Multiprocessing board	UDP	User datagram protocol
GPB	General processing board	MSC	Mobile switching center	UMTS	Universal mobile telecommunications system
GPRS	General packet radio service	OSA	Open service architecture	VHE	Virtual home environment
GSN	GPRS support node	PDP	Packet data protocol	VLR	Visitor location register
		PEB	Power and Ethernet board		
		POP	Point of presence		

ing the concept of horizontal layering in the ongoing standardization of 3GPP, due for release 4 in March 2001.

The horizontally layered network architecture

In the horizontally layered network architecture (Figure 1), functionality and nodes are arranged in layers according to their specific areas of use. This separation into independent layers, which is a key principle in modern networking, has been integrated into several standardization initiatives such as Megaco (in the IETF), Tiphon (in ETSI) and the Multiservices Switching Forum (MSF) led by several large operators and manufacturers. It is also an integral part of 3GPP. The layered concept of the network architecture introduced with release 4 of the 3GPP specifications comprises three distinct layers:

- an application layer;
- a network control layer; and
- a connectivity layer.

The application layer

The application layer is where the end-user applications reside. In modern networks, applications are implemented in mobile terminals and in dedicated application servers in the network. The application servers are often complemented with content servers, which host service-related databases or libraries (such as video-clip libraries or news history databases).

Concepts such as the virtual home environment (VHE) and open service architecture (OSA) were developed in the 3GPP to allow operators to provide unique services. Operators benefit from being able to differentiate themselves from one another by providing unique services, thus securing for themselves a higher position in the value chain. They also have the option of developing these services themselves or of obtaining them from third-party software houses—they can even get external service providers to run them. This flexibility allows the operator to choose from a huge portfolio of services that it can offer its subscribers.

The application layer (Figure 2) interfaces with the network control layer via a defined set of open application program interfaces (API). By using open APIs, application developers can make use of the features of standardized service capabilities, to design new services and applications.




Figure 1
Horizontally layered network architecture.

Figure 2
The structure of the application layer.





Figure 3
Domains in the network control layer.



The network control layer

The network control layer incorporates all the functionality needed to provide seamless, high-quality services across different types of network. The different networks can be seen as a set of domains, each of which houses control servers that are specific to a given network (Figure 3).

Among the domains listed is the traditional circuit-switched domain based on the GSM standard, which uses ISDN-related signaling. Also shown are the soon-to-emerge packet-switched domain based on GPRS technology, and the Internet multimedia domain; both will be introduced in release 5 of the 3GPP specifications in December 2001. Other domains control access to the fixed network or the delivery of cable TV.

Generally speaking, the network control layer houses several different kinds of network server. The servers are responsible for controlling mobility management, the set-up and release of calls and sessions requested by end-users, circuit-mode supplementary services, security, and similar functions. These domains can be owned by various individual operators or wholly owned by a single operator.

The home location register (HLR, known from the GSM standard) is being integrated into the home subscriber server (HSS). In contrast to the other control servers in the domains, the HSS plays a central role in the network control layer and so becomes a

multidomain entity. Because it contains all subscriber data in the network control layer, the HSS can handle authorization, authentication and location management queries from all domains.

The connectivity layer

The connectivity layer is a pure transport mechanism that is capable of transporting any type of information via voice, data and multimedia streams. Its backbone architecture incorporates core and edge equipment.

The core equipment transports aggregated traffic streams between the different nodes at the edges of the backbone. As a rule, core equipment is a backbone router or backbone switch that handles traffic streams either according to very simple classification principles, or according to routes that the network operator has predefined by means of traffic engineering.

Edge equipment collects customer-specific data and statistics for accounting and billing purposes, and provides the single bit-pipes that guarantee an appropriate quality of service. The edge equipment is usually a media gateway, which operates under the full control of the nodes in the network control layer. In addition, a media gateway allows the bit streams to be processed, thus providing coding/decoding of speech streams, canceling echo, bridging multiple party calls, and converting transport protocols. The nodes in the network control layer also control these manipulations.

This exertion of control down to the bit-stream level allows the variety of services and applications implemented by the different network control domains to be achieved via a common connectivity layer. At the same time, the services and application are independent of the transport technology applied, which may be mixed or vary over time as the network evolves.

Connectivity-layer solutions can be based on either asynchronous transfer mode (ATM) transmission or Internet protocol (IP) transmission (Figure 4). However, in the long run, the role of ATM as a transport technology is expected to diminish, leaving IP transmission as the predominant technology.

Control servers used in UMTS

What is a control server?

A control server is a network element that belongs to the network control layer. As mentioned above, the network control layer hosts a set of domains, and a control server can only occupy one domain. Control servers communicate with other control servers in the same domain by means of standard layer-3 protocols, such as ISUP and mobile application part (MAP). There is also some inter-domain communication, but this can be regarded as an exception to the rule.

A control server solely handles control communication—no user-plane streams pass through it. The user plane is located in the connectivity layer and control servers operate on it via the gateway control proto-

col (GCP). The control servers located in the circuit-switched domain are the mobile switching center (MSC) server, the gateway MSC server, and the transit switching center (TSC) server.

In the packet-switched domain, we find the serving GPRS support node (SGSN) server. Note: the gateway GPRS support node (GGSN) is regarded as a network element that handles the user-plane stream; consequently, it is positioned as a specialized media gateway in the connectivity layer.

Ericsson has recently defined the IP multimedia domain in 3GPP standardization. The call state control function (CSCF) server, which handles multimedia calls based on the session initiation protocol (SIP), is located in that domain.

MSC server

The MSC server, which is located in the circuit-switched domain, handles control layer functions related to circuit-mode communication services at the border between the access network and the core network. It hosts the following main functions:

- Call-control and related supplementary services—the MSC server supports call-control functionality as well as a comprehensive set of supplementary services for the teleservices and bearer services defined in GSM and UMTS.
- CAMEL and IN services—the MSC server supports the integrated service switching function (SSF) entity as defined by the customized applications for mobile network enhanced logic (CAMEL) specifications of GSM and UMTS.




Figure 4
The connectivity layer, based on ATM or IP transmission.

- Roaming—together with the HLR, the MSC server supports roaming by subscribers within the GSM network. Inter-system roaming allows the subscriber to roam from UMTS networks to GSM networks and vice versa. The MSC server supports mobility management, thus enabling mobile stations to attach/detach and to roam within the UMTS network, between UMTS networks, and between UMTS and GSM networks.
- Handover—the MSC server supports intra-MSC relocation for UMTS—that is, the target RNC is connected to the same MSC server. The MSC server supports inter-MSC and intra-MSC handover from UMTS to GSM so that the teleservices and bearer services that are common to UMTS and GSM can continue as the subscriber moves about.
- Media gateway control function—the MSC server supports the gateway control protocol (GCP) for controlling the circuit handling functionality of the media gateway.
- Integrated visitor location register (VLR) and subscriber data management—the MSC server, which has built-in VLR functionality, supports the standardized in-


terface to the HLR in order to manage subscriber data, such as international mobile subscriber identity (IMSI), subscribed services, and so on.

- Control signaling—the MSC server supports the radio access network application part (RANAP) for control signaling over the RNC-MSC interface, so that radio access bearers can be established and released. The MSC server also supports MAP signaling to the HLR and equipment identity register (EIR). Furthermore, the ISDN and bearer-independent call control (BICC) protocols are used for control signaling between the call-control servers in the circuit-switched domain and external integrated services digital networks (ISDN).

In addition to these functions, the MSC server houses the functionality of a GMSC and TSC server.

To illustrate the mechanisms in the network control layer, it is best to analyze a typical traffic case. Figure 5 shows the network level view of a call originated from a mobile terminal. Let us assume that the mobile terminal is roaming in a UMTS access network using a WCDMA air interface—the mobile terminal is already registered in the MSC

Figure 5
Example of a call originated from a mobile terminal.



server, and the integrated VLR has fetched and stored the subscriber data from the HSS.

When the mobile user initiates a call, the MSC server receives a RANAP message. This message specifies the destination of the call and the type of service—for example a voice call. Taking into account the location from which the call originates, the MSC server selects a media gateway (located at the edge) to the access network.

It is the media gateway's task to adapt the access network to the core network. Adaptation is required for the

- conversion from ATM to IP (as the transport technology);
- conversion from ATM-based quality of service (constant bit rate), to IP-based quality of service based on the differentiated services (DiffServ) concept; and
- conversion of the framing mechanism used at the *Iu* interface to standard IP framing based on the real-time transport protocol (RTP).

To allow traffic engineering in the core network, the media gateway must select an appropriate multiprotocol label-switching (MPLS) path to the destination. To increase transmission efficiency, besides compressing the lengthy RTP/UDP/IP header of the voice sample, the media gateway can multiplex several connections into one MPLS path.

The MSC server selects a second media gateway according to the destination of the call—that is, at the edge that leads to the external ISDN or public switched telephone network (PSTN). In the media gateway, the MSC server selects the voice coder hardware, which allows coded speech to be transported from the terminal up to the edge of the core network. Coded speech consumes only one seventh as much bandwidth as pulse-code-modulated (PCM) speech, thus saving bandwidth in the transmission backbone network.

When it begins communicating with the TSC server, the MSC server controls two media gateways. The BICC communication protocol is used. This protocol is an evolved version of the ISDN protocol. In contrast to the ISDN protocol, which is defined for time-division multiplexing (TDM) trunks, the BICC protocol allows any transmission media to be used.

The TSC server owns the TDM logical trunks that lead to the external ISDN/PSTN. The physical trunks are also located in a media gateway. As it receives the set-up request from the MSC server, the

TSC server also receives the identity of the selected media gateway at the edge leading to the ISDN/PSTN. The TSC server must then select a media gateway with TDM trunks to the requested destination. To minimize the use of network resources, the TSC selects a media gateway at the edge that leads to the ISDN/PSTN. This means that the media gateway provides both the voice coder and the TDM trunks in the same physical node, and is controlled by two control servers for the same call.


Finally the call set-up is extended to the external ISDN/PSTN where the called party is located.

SGSN server

The SGSN server, which is located in the packet-switched domain, handles control layer functions related to packet-mode communication services at the border between the access network and the core network. The SGSN server hosts the following main functions:

- session management—by session management is meant the establishment, maintenance, and release of end-user packet data protocol (PDP) contexts. This includes interworking with the GGSN for IP addresses. Session management also includes functionality for establishing WCDMA radio access bearers for end-user IP data transportation, as well as functionality for end-user QoS negotiation.
- mobility management—by mobility management is meant the functionality that supports roaming within and between GSM and UMTS mobile networks.
- integrated VLR and subscriber data management—the SGSN server supports the standardized interface to the HLR in order to manage end-user subscriber data, such as the international mobile subscriber identity (IMSI), quality-of-service profile, access point names, and so on.
- integrated domain name server—the SGSN server supports an integrated domain name server (DNS) to find GGSN IP addresses related to a given access point name (APN).
- GGSN interface—the gateway tunneling protocol-control (GTP-C) supports control signaling between SGSN servers and the GGSN. GTP-C is transported by UDP/IP and contains functionality for SGSN-to-GGSN tunnel management and control.

Figure 6
Example of PDP context activation initiated by a mobile terminal.



- MAP and RANAP control signaling—the SGSN server supports the RANAP protocol for control signaling over the RNC-SGSN interface for establishing and releasing radio access bearers. The SGSN server also supports MAP signaling to the HLR, equipment identity register (EIR), MSC server and SMS center (SMS-C).
- media gateway control function—the SGSN server supports the GCP protocol for controlling the packet-handling functionality of the media gateway.

To illustrate the mechanisms in the network control layer, let us analyze a typical traffic case. Figure 6 shows the network level view of packet data protocol (PDP) context activation initiated by a mobile terminal. Let us assume that the mobile terminal is roaming in a UMTS access network using the WCDMA air interface. The mobile terminal has already been attached to GPRS in the SGSN server, and the integrated VLR has fetched and stored the subscriber profile from the HSS.

When the mobile user requests that a PDP context be activated, the SGSN server receives a RANAP message. The message contains the requested QoS for the session and the related access point name. The SGSN server then identifies the target GGSN by checking the received access point name (APN), and using the integrated domain name server (DNS), resolves the IP address of the selected GGSN.

The SGSN server then selects a media

gateway at the edge facing the access network and orders it to create a GTP tunnel transfer point. The media gateway sends a request, via the GTP-C protocol, to the GGSN, asking it to establish a GTP tunnel and to allocate an IP address for the terminal that will be used during the session.

The GGSN assigns a dynamic IP address, either from the range of IP addresses allocated to the public land mobile network (PLMN), or from an external authentication, authorization and accounting (AAA) server or remote authentication dial-in user service (RADIUS) server. This IP address is returned to the SGSN server and, finally, the GTP tunnel is established. As soon as the PDP context activation is confirmed to the mobile terminal and the user's own IP address is received, IP-based communication can start between the mobile user and the external packet data network via the GTP user plane (GTP-U) tunnel.

If global mobility is to be implemented for roaming between various networks, the mobile IP concept can be employed. To support global mobility, the GGSN must support "foreign-agent" functions. The corresponding "home agents" can be located anywhere in the global Internet.

CSCF server

The call/session control function (CSCF) server, which is located in the IP multimedia domain, handles control-layer functions related to SIP-based multimedia sessions. The CSCF server, which is currently

being defined by standards bodies, can host one or more of the following main functions:

- interrogating CSCF (I-CSCF)—the I-CSCF function interrogates an external location-service function to determine which CSCF is serving at a given time and acts as a SIP proxy firewall to fulfill network security and privacy functions;
- serving CSCF (S-CSCF)—the S-CSCF is the serving network element with which subscribers register in order to be reached when roaming. It temporarily stores user-profile-related data, which is downloaded from the HSS as registration takes place. The S-CSCF also triggers the call- and session-related services to which the user has subscribed; and
- proxy CSCF (P-CSCF)—the P-CSCF, which contains a very limited CSCF function, is the only Internet multimedia function placed in the network being visited; that is, when the subscriber's unit is roaming outside of its home network. The P-CSCF contains address translation functions in order to proxy the session request directed at the I-CSCF in the home network.

To illustrate the mechanisms in the network control layer, it is best to analyze a typical traffic case. Figure 7 shows a SIP call from a mobile terminal to a fixed ISDN phone. Let us assume that the mobile terminal roams to a network that provides UMTS packet access using the WCDMA air inter-

face. The mobile terminal is already attached to GPRS in the SGSN server. We can also assume that the terminal has performed a SIP registration to the home network—that is, a serving CSCF server has been allocated to the terminal, and data relating to the user profile has been fetched from the HSS and stored in the serving CSCF server.

When the mobile user wants to initiate a SIP call while roaming in another network, the GPRS bearer must first be established. The network does this by activating a PDP context, as described above.

When the GPRS tunnel has been established, a SIP message is sent to the proxy CSCF server in the visited network. The proxy CSCF server is able to resolve the address of the subscriber's home network and forward the SIP message.

The SIP message enters the home network in an interrogating CSCF server, whose task is to find the serving CSCF server in which the subscriber is currently registered. The interrogating CSCF server does this by sending a query to the location service in the HSS.

When it receives a reply from the HSS, the interrogating CSCF server forwards the SIP message to the serving CSCF server, which acts as host to the call-control logic. Because the destination address is located in the fixed ISDN network, the SIP message is forwarded to the ISDN gateway network, which may be physically located in either




Figure 7
Example of a SIP call from a mobile terminal to the ISDN.

the home network or the visited network. Within the ISDN gateway network, a media gateway control function (MGCF) converts the SIP message into appropriate ISDN messages and forwards them to the external ISDN network.

While the control signaling messages travel from the visited network via the home network to the ISDN network, the user-plane stream is sent directly from the GGSN in the visited network to the media gateway in the ISDN gateway network, taking the shortest optimal path to the destination. The media gateway at the edge facing the external ISDN network is controlled by the MGCF.

HSS

The HSS, which holds the master database for any given user, contains the subscription-related information necessary to support the network control layer entities' handling of calls or sessions. For example, the HSS provides support to the call-control servers, in order to complete the routing or roaming procedures, by handling authentication, authorization, naming and addressing, and location dependencies. The HSS hosts the following main functions:

- location service—this function keeps track of the domain in which a user is currently located. When the user is roaming in the Internet multimedia domain, the location service returns the identity of the S-CSCF;
- AAA service—this function is used to au-

thenticate and authorize users in the Internet multimedia domain;

- database—the database contains all relevant information pertaining to mobile subscribers, including dynamic data, such as the subscriber location and the status of supplementary services, and permanent data, such as subscriber-associated numbers and category information. Authentication and ciphering data for each mobile subscriber are also included;
- administration support—this function handles the commands given by the operator to connect mobile subscribers and to define their corresponding subscriber data;
- analysis functions—analysis functions are used to analyze mobile subscriber numbers, such as IMSI, MSISDN, and forwarded-to numbers; and
- mobile application part (MAP)—the MAP receives messages and makes sure that the appropriate action is taken. This consists of the call-handling part, which routes calls that terminate in a mobile terminal, handles mobility data and supplementary service subscriber procedures, and analyzes and handles bearer capabilities.

The HSS is involved in almost all traffic cases of the different domains.

Control server platform

MSC server

The platform used for the MSC server, the TSC server and the GMSC server is a continuation of Ericsson's GSM MSC/VLR, which is based on AXE technology. However, to cope with the demands that a pure control server is expected to meet, the AXE platform is being streamlined for optimized performance and characteristics:

- the media gateway control function has been introduced and group switch hardware has been removed;
- signaling connectivity has been improved to support high-speed ATM and IP signaling interfaces;
- a new CPU (APZ 212 40) has been introduced that is based on open commercial products;
- parallel code execution on multiple CPUs is targeted for the near future; and
- a new design environment has been introduced that is based on the specification and description language (SDL) standard of ITU-T.

SGSN server

BOX B, THE APZ 30 CONCEPT

APZ 30 is Ericsson's common architecture for the next generation of network servers and controller nodes. This architecture

- ensures carrier-class in-service performance on highly scalable hardware and software components;
- supports open standards through the use of common APIs and industry-standard development languages;
- makes extensive use of off-the-shelf hardware to give a high-value, future-proof investment; and
- secures a common user-interface and hardware structure for a wide range of network nodes.

The purpose of APZ 30 is to allow each application to be written in the most suitable programming language, and to be run in the most suitable environment on the most appropriate hardware. Common interfaces and hardware ensure optimal synergies between applications, besides facilitating business with satisfied customers.

APZ 302 configuration

This is the TelORB configuration. It consists of TelORB³, which is a distributed operating system and middleware for real-time applications. Thanks to TelORB mechanisms, processing capacity can be scaled linearly. The configuration has been certified for use with 2 to 28 Pentium processors. Furthermore, TelORB is highly available (zero downtime) and is an open system. Applications on TelORB can be written in Java and C++. TelORB runs on Pentium processors.

APZ 303 configuration

This is the NSP/Ronja/DPE configuration. It consists of Ronja, which is middleware that (with the DPE component) merges the advantages of high availability, distribution and scalability. The application, which is written in Java, runs on Solaris. Ronja/DPE and Solaris run on SPARC processors.

The SGSN server is built on the APZ 303 platform. The physical enclosure is a cabinet with subracks, fans and air-inlet units.

The processor and interface boards can host different functions, realized through the assembly of different computing modules. The boards are mounted in slots in the backplane of the subrack.

The backplane of the subrack distributes power and handles internal connections. Dual high-capacity connections facilitate communication between subracks. The backplane provides redundant power distribution and a redundant network for board interconnections.

The device board can carry different processor modules and interfaces. It is a generic and stable platform for present and future designs. Generally, the device board is very flexible and allows for many different configurations:

- general processing board, GPB—supports high capacity processors. The GPB has an on-board disk for secondary storage.
- multiprocessing board, MPB—handles protocols on level 3 and above.
- interface board, IB—mainly handles interface-specific protocol termination. Many variants are available, such as STM-1, E1 and T1.
- power and Ethernet board, PEB—supplies power and handles internal communication.

The middleware consists of Ronja⁶, a type of portable Java middleware for telecom applications, and an integrated distributed processing environment (DPE).

CSCF server

The CSCF server, which is under development on the APZ 303 platform, is based on the open Solaris operating system that executes on commercially available SPARC CPUs. The middleware consists of Ronja and an integrated distributed processing environment.

HSS

The HSS is under development on the APZ 302 platform. It contains a highly scalable, real-time database cluster for handling high-performance, transaction-oriented applications. The processing hardware is a cluster of loosely coupled CPUs, which provide excellent scalability from very small systems up to large clusters of several dozen CPUs. The main building blocks of the HSS are

- equipment shelves with a switched Ethernet backplane;
- processor boards with or without hard disks and loading media;
- operating system and support software;
- high-performance, real-time transaction software and database; and
- O&M functionality.

Conclusion

The horizontally layered architecture allows network operators to evolve their systems with the highest flexibility—a key benefit now, when the industry's future seems more and more difficult to foresee.

The network control layer, which is the central layer in the horizontally layered architecture, houses different domains; the control servers are specific to the “controlling” domains. As new technology becomes available, networks can be upgraded with, say, a control domain for packet switching, or for SIP-based IP multimedia.

All control domains reuse a common transport mechanism that optimizes the use of transport resources. Reuse allows for different transport technologies—both existing and new—to be deployed without affecting any layer other than the ones on which they are applicable.

All control domains share a generic interface to the application layer. Services and applications can be developed and implemented once and then reused by all the control domains, independent of the connectivity technology that may be applied.

TRADEMARKS

Solaris is a trademark of Sun Microsystems, Inc. in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

Pentium is a registered trademark of Intel Corporation in the U.S. and other countries.

REFERENCES

- 1 Dahlin, S and Örnulf, E.: Network evolution the Ericsson way. Ericsson Review Vol. 76(1999):4, pp. 174-181.
- 2 Architectural Principles for Release 2000, 3GPP technical report TR 23.821
- 3 Hennert, L. and Larruy, A.: TelORB—The distributed communication operating system, Ericsson Review Vol. 76(1999):3, pp. 156-167.
- 4 Fyrö, M.; Heikkinen; K.; Petersen, L-G. and Wiss, P.: Media gateway for mobile networks. Ericsson Review Vol. 77(2000):4, pp. 216-223.
- 5 IP Multimedia Subsystem Stage 2, 3GPP technical specification 3G TS 23.228
- 6 Karlson, M.: Ronja—A Java application platform. Ericsson Review Vol. 77(2000):4, pp. 244-247.