Ericsson Telecom Server Platform 4

Victor Ferraro-Esparza, Michael Gudmandsen and Kristofer Olsson

The marriage of telecommunications and the Internet puts new requirements on equipment. Customers have come to expect the same quality of service as they get from present-day telecommunications networks. At the same time, they expect new services in the multimedia and services domain.

We know from experience that modern telecommunications systems are extremely reliable and provide real-time responses. This level of reliability and response cannot currently be achieved using technologies in the Internet domain. On the other hand, the Internet is far richer in terms of content, where pictures (still and moving) are part of today's experience.

In an ideal world we would reuse technologies from the data communications industry and combine them with those from telecommunications systems. With the Telecom Server Platform 4, Ericsson has taken a giant step in this direction, combining its know-how and experience of telecommunications, reliable systems, and scalable systems with open technologies, such as Linux, CORBA, SNMP, LDAP and other standards, such as cPCI and *de facto* standards. The result is a carrier-class server which is always on, scalable, and open, and which adds value for Ericsson's customers.

Ericsson Review no. 1, 2002 described the Ericsson IP Multimedia solution and the role the Telecom Server Platform plays in this domain. The Telecom Server Platform also plays an important role in delivering services. In this article, the authors describe the Ericsson Telecom Server Platform 4 in greater detail. They also briefly describe the service network framework and give some examples of applications in the services domain.

Product overview

The Ericsson Telecom Server Platform 4 is more robust and fault-tolerant than any comparable server technology. It is extremely reliable with unique, linearly scalable capacity, and real-time characteristics, which means that transmission takes place with minimal and controlled delay.

The telecommunications-grade software is enabled by TelORB clusterware, which runs on top of DICOS and Linux. The software incorporates the very latest in signaling system no. 7 (SS7) and built-in node management, with support for the many protocols needed for interoperability between the Telecom Server Platform and operations support systems (OSS).

In its make-up, the Telecom Server Platform combines Ericsson's long tradition of designing robust and reliable hardware with commercially available components. The result is an open hardware architecture with excellent characteristics: scalable capacity, telecommunications-grade reliability, small footprint and minimal power consumption.

Hardware

The Telecom Server Platform uses commercially available hardware with ample capacity and dependable node performance. This lowers the cost of installation, operation, and maintenance.

The hardware fulfills all telecommunications requirements regarding power consumption, low electromagnetic radiation, reliable equipment practice, and reliable connectors. This ensures trouble-free operation. What is more, the modular design facilitates upgrading.

Some key features of the hardware are redundancy in all hardware components, "hot-swap" hardware replacement, a modular platform for maximum flexibility, standard

BOX A, TERMS AND ABBREVIATIONS

3GPP	Third-generation Partnership Project
AAA	Authentication, authorization and accounting
API	Application program interface
ASP	Application service provider
CAMEL	Customized applications for mobile network-enhanced logic
CAP	CAMEL application protocol
CM	Configuration management
CORBA	Common object request broker architecture
cPCI	Compact PCI
CPI	Customer product information
DICOS	Object-oriented operating system with excellent real-time characteristics. Developed by Ericsson.
E1/T1	PDH transmission frame formats for 2 Mbit/s (E1) alt. 1.5 Mbit/s (T1) transmission rates
FM	Fault management
FPGA	Field-programmable gate array
GEM	Generic Ericsson magazine
GESB	Gigabit Ethernet switch board
GUI	Graphical user interface
HTTP	Hypertext transfer protocol
IIOP	Internet inter-ORB protocol
I/O	Input/output
IPC	Inter-processor communication
IPv4	Internet protocol version 4
ISR	In-system reconfigureable FPGA
J2EE	Java 2 Enterprise Edition
LDAP	Lightweight directory access protocol
MAP	Mobile application part
MIP	Mobile IP
NAS	Network access server
O&M	Operation and maintenance
ORB	Object request broker
OSA	Open system architecture
OSS	Operations support system
PCI	Peripheral computer interconnect
PM	Performance management
RADIUS	Remote access dial-in user service
RAM	Random access memory
RMI	Remote method invocation
RPC	Remote procedure call
SCB	Support & connection board
SCS	Service capability server
SDRAM	Synchronous dynamic RAM
SIGTRAN	Signaling transport
SNF	Service network framework
SNMP	Simple network management protocol
SOAP	Simple object access protocol
SS7	Signaling system no. 7
SSL	Secure socket layer
TCP	Transmission control protocol
TMN	Telecommunications management network
UDDI	Universal description, discovery and integration
UDP	User datagram protocol
W3C	World Wide Web Consortium
WSDL	Web services description language
XML	Extensible markup language

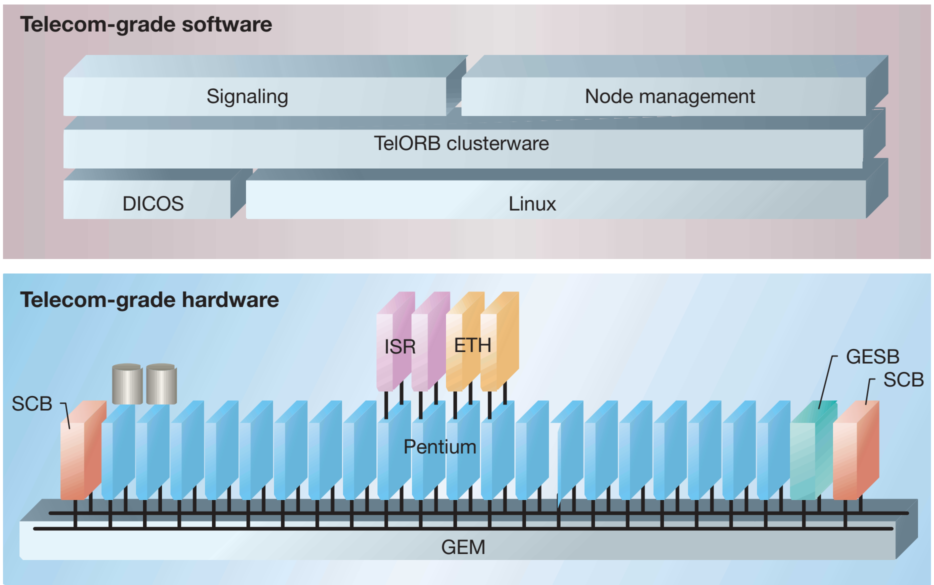

Figure 1
The Telecom Server Platform consists of telecom-grade hardware and software.

hardware interface boards, industry-standard components (creating a future-proof architecture), and optional geographical node redundancy for additional dependability.

The Telecom Server Platform 4 can be duplicated in a redundant standby node. That is, in case of failure, the load is automatically redistributed within the node or, in the case of geographical node redundancy, a remote standby node can take over the functions of the primary node.

The ability to hot-swap hardware allows operators to replace any component at any time without affecting system performance. In a market of converging technologies and standards, the use of industry-standard components creates an architecture that is future-proof, scalable, flexible and cost-effective.

The hardware platform includes processor modules (with or without peripherals), signaling processors, Ethernet switches, power supplies and fans. All processors are standard, commercial, off-the-shelf, single-board cPCI computers (Box B).

The unit is housed in an Ericsson cabinet (BYB 501). An expansion cabinet is avail-

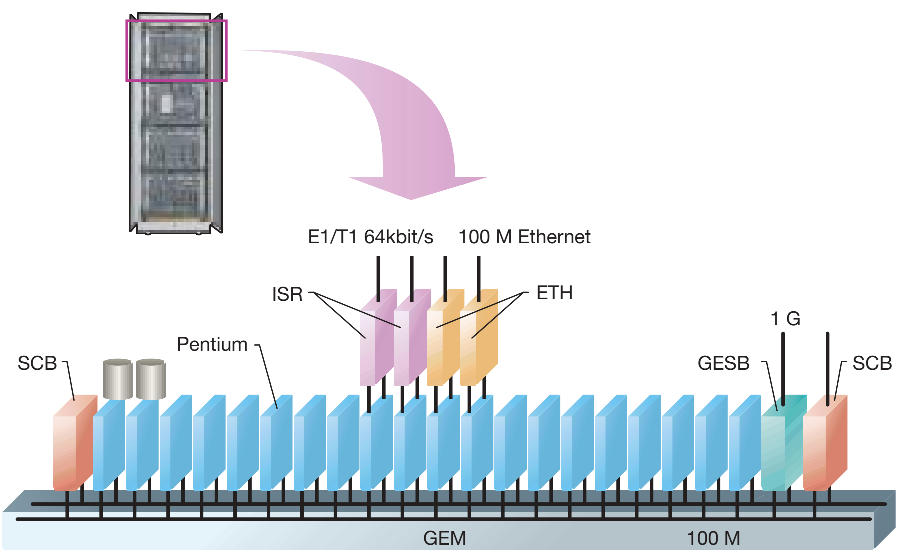

Figure 2
Each subrack consists of *n*+1 duplicated processors connected via duplicated 100 Mbit/s Ethernet. The subracks are interconnected via 1 Gbit/s Ethernet. Every component is at least duplicated.

able if additional processor modules are needed. All hardware can be accessed from the front panel, ensuring easy maintenance and replacement. To ensure high quality and efficient installation, all external cables are connected to a patch panel. Together, the main cabinet and expansion cabinet form a single node.

Software

The software consists of an operating system, clusterware, node management, and network signaling.

Operating system

The architecture of the Telecom Server Platform 4, which has been designed to support

BOX B, HARDWARE COMPONENTS OF THE TELECOM SERVER PLATFORM

Processors	Intel Pentium III 700 MHz cPCI boards with 1 GB SDRAM.
E1/T1/J1 ports	To execute scalable SS7/ITU-T/China/ANSI protocol stacks.
Ethernet ports	100 Mbit/s Ethernet connections used for external IP communication, such as the signaling transport (SIGTRAN) protocol.
Input/output devices	DVD: one drive for input purposes, including initial loading. Tape drive: one 20 GB drive, for example, for backups.
Hard drives	Up to 14 x 18.2 GB storage capacity for every executable unit needed for start-up and backup.
Patch panel	A jack frame to which the cables in the cabinet are terminated, and to which all external cabling is connected.
Alarm collector	Collects alarms from fans and supply voltage supervision at level-1 switches in the magazines.
GEM	Generic Ericsson magazine (GEM)—a standard subrack that provides a cost-effective solution with small footprint. Each GEM has a 100 Mbit/s Ethernet level-1 switch that connects all processor boards. A 1 Gbit/s Ethernet level-2 switch interconnects the magazines. The switches are duplicated for redundancy.
Cabinet	The Ericsson BYB 501 uses forced-air ventilation, allowing heat to dissipate through the bottom and out at the top of the cabinet. Ericsson expansion cabinets are available.
Power supply	Input to the cabinet is -48 volts.

Capacities	**Maxi**	**Midi**	**Mini**	**Micro**
E1/T1 connections	24	16	16	8
ITU No. 7/SS7 signaling channels	192(E1)/144(T1)	128(E1)/96(T1)	128(E1)/96(T1)	64
Ethernet ports	32	28	22	12
Processors	42	31	21	10
RAM data storage	42 GB	31 GB	21 GB	10 GB
DVD	1	1	1	0
Tape streamer (20GB)	1	1	1	0

Environmental requirements

Recommended minimum ceiling height:	210 cm
Relative humidity (min-max):	20-80%
Temperature (min-max):	+5° to 40°C*
Temperature (normal operation):	+20°C

* The tape streamer/tapes are a limiting factor when it comes to average temperature and humidity range. If the temperature reaches +35°C, the tape streamer tapes stretch, which leads to loss of data.

Agency approvals

The hardware has been designed to comply with NEBS level 3.
Seismic vibration, EN 300 019-2-3 and GR-63-CORE zone 4.
EMC - EN 300 386 class B,
Part 15, Subpart B, Class B/Federal Communications Commission (FCC) according to GR-1089-CORE.
Product safety - EN 609 50, IEC 609 50 and ANSI/UL 1950, third edition.

Design for environment

The Telecom Server Platform complies with Ericsson's policy to avoid the use of banned and restricted substances.

End-of-life treatment

Ericsson offers recycling services for old Ericsson products. The materials are taken care of by approved recycling companies in compliance with EU or other national legislation.

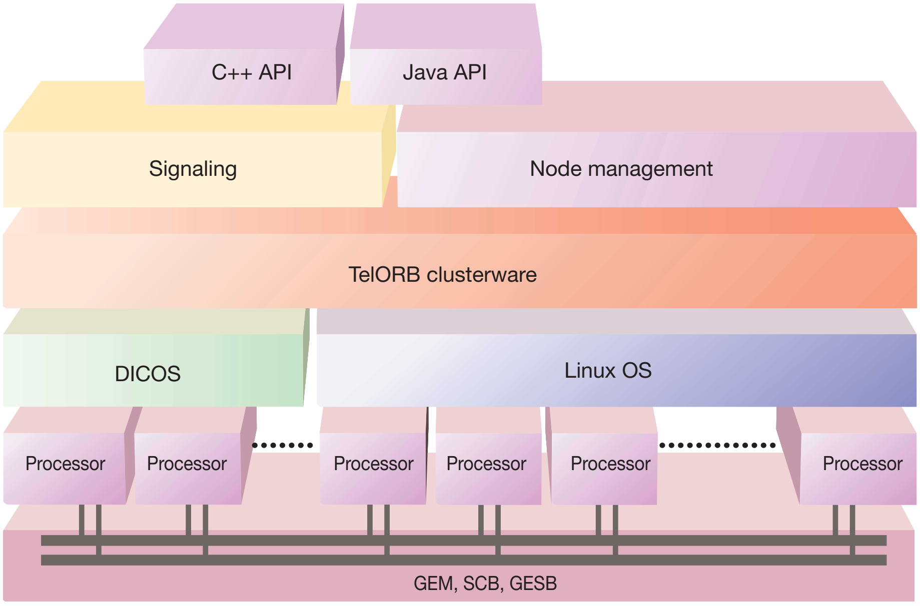

Figure 3
Software architecture of the Telecom Server platform. The TelORB clusterware also controls the Linux processors in the cluster.

different commercial hardware and multiple operating systems, provides the most appropriate executing environment for processes. The current version of the Telecom Server Platform uses the Linux operating system for UNIX-like performance, and DICOS for real-time, mission-critical tasks.

DICOS, which is based on queuing theory, offers soft real-time response. This kind of real-time performance is suitable for telecommunications and data communications applications, especially for database clusters.

TelORB clusterware

The communication layer of the TelORB clusterware connects the different processors to each other to enable inter-processor communication (IPC). Internally, the TelORB clusterware uses the lightweight protocol IPC to manage duplicated Ethernet backbones. When communicating to the external world, it uses the transmission control protocol (TCP), user datagram protocol (UDP), and IPv4.

A built-in object-oriented database provides persistent storage in RAM. The database is thus always held in primary memory and distributed over the processors. All data is replicated and stored on more than one processor. Should a subroutine or processor fail, the entire database remains available to the software.

The software management layer automatically configures the executing software to ensure that it runs efficiently on any of the processors available to the TelORB cluster. It provides support for software upgrades and monitors and manages software components to ensure high availability. It also supports binaries from third-party vendors. The node-management function includes an element manager that configures and manages every managed object of the TelORB clusterware and operating system.

An object request broker based on the common object request broker architecture (CORBA) has been incorporated to allow the TelORB system to communicate with other systems and the graphical user interfaces used to manage them. It supports

Figure 4
The architecture of the TelORB clusterware.

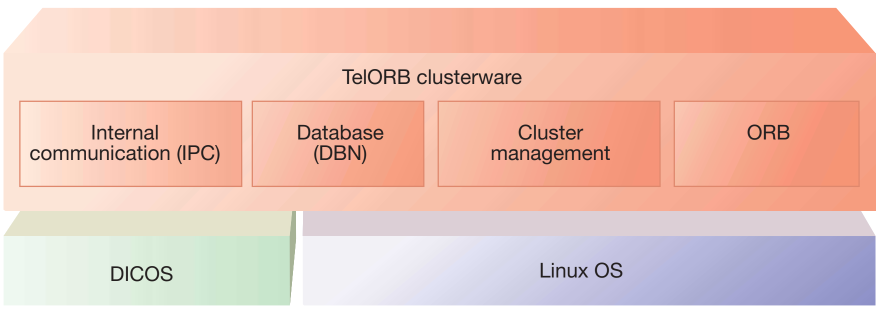

Figure 5
The TelORB clusterware handles hardware and software faults. All data is duplicated (1 & 2).
When a software fault occurs, the processor is restarted and the processes and data are restored (3a & 4a).
When a hardware fault occurs, the processor is taken out of service and new duplicates of the data are made in the cluster (3b & 4b).

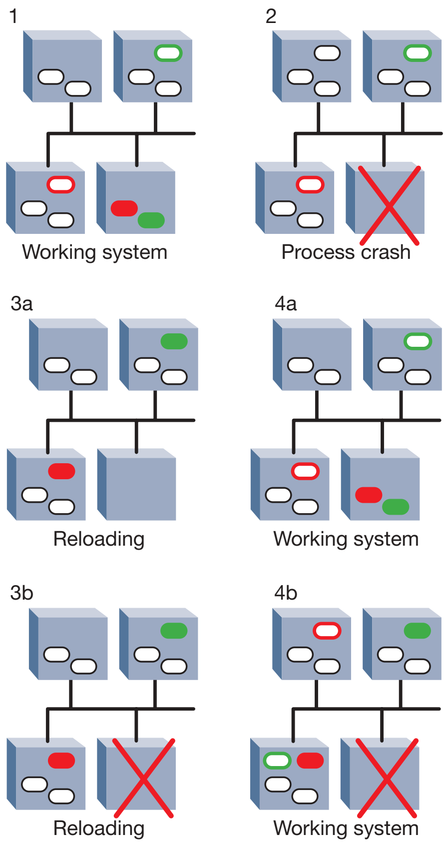

Figure 6
Node management in the Telecom Server Platform.

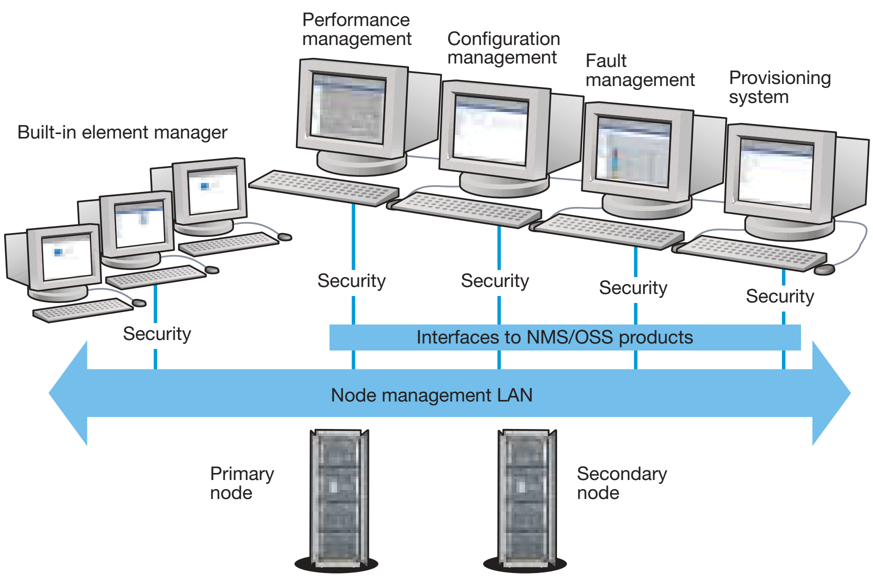

IIOP1.1 to the Internet inter-ORB protocol (IIOP) gateway, Java remote method invocation (RMI) over IIOP/IPC, and secure CORBA through SSLIOP.

Simply put, high availability means that regardless of what happens to the node, the services supplied by the network are not disturbed. The TelORB clusterware provides high availability to the node thanks to inter-node as well as intra-node redundancy.

TelORB enables applications to be divided into a large number of mutually independent resources and execution tasks that are evenly distributed among available processors in the cluster. This is how the Ericsson Telecom Server Platform succeeds at scaling capacity in a completely linear fashion.

Node management

The Telecom Server Platform offers a node-management solution based on the telecommunications management network (TMN) model. To better align with operator requirements and industry trends, the node-management system also incorporates other standards, such as CORBA, lightweight directory access protocol (LDAP), hypertext transfer protocol (HTTP) and the simple network management protocol (SNMP).

The node-management system implements a manager-to-agent architecture that integrates the node into external network-management and customer-administration systems using standard protocols to communicate with external systems.

In the Telecom Server Platform 4, the node-management function incorporates fault management (FM), configuration management (CM), provisioning support, performance management (PM), logging service (and log querying), an XML exporter, an element manager, and license management.

Network signaling

The signaling solution of the Telecom Server Platform can handle SS7 and Internet protocols. The SS7 protocol is still the most commonly used protocol for data intercommunication in telecommunications networks. The SS7 signaling stack of the Telecom Server Platform 4 is the latest in SS7 technology and features a scalable implementation that can be distributed over the processors used for traffic handling. The design includes a front-end process for terminating the E1/T1 interfaces, and a back-end process for the upper layers of the stack.

Operator benefits

The Telecom Server Platform is the foundation for several revenue-generating applications for mobile and wire-line operators.

High availability and reliability

High availability is a must in today's modern networks. The Telecom Server Platform achieves high availability through the TelORB clusterware and highly reliable hardware. The main characteristics of the TelORB clusterware can be summarized as follows:

- always-on operation;
- automatic software recovery;
- data replication;
- overload control;
- software updates during operation;
- upgrade of operating system allowed during operation;
- online backup;
- hot-swap hardware replacement; and
- optional geographical node redundancy.

Scalability

Besides high network availability, operators also want scalable network nodes. The Telecom Server Platform has been designed with scalability in mind, ensuring that the initial investment suits customer needs, and that the system can grow with the business. Also, because it is built on commercial hardware, the Telecom Server Platform benefits from continuous enhancements in the processor industry. The modular architecture facilitates rapid expansion of capacity when and where needed. This can be achieved through predefined expansion paths between the configured cabinets.

Cost reduction

The Telecom Server Platform uses commercial, off-the-shelf hardware and software. Operators can thus employ best-in-class hardware and software, and benefit from economies of scale. Because it is scalable and supports co-location, the Telecom Server Platform allows operators to add new applications on a single node in accordance with market demand.

The Telecom Server Platform node management uses a Web-based graphical user interface (GUI) that enables operator staff to access all O&M tasks online—on- and off-site (Figure 7). This functionality reduces delays and travel costs, which, in turn, reduces the total cost of ownership.

Compared to other solutions on the market, the Telecom Server Platform is a total-

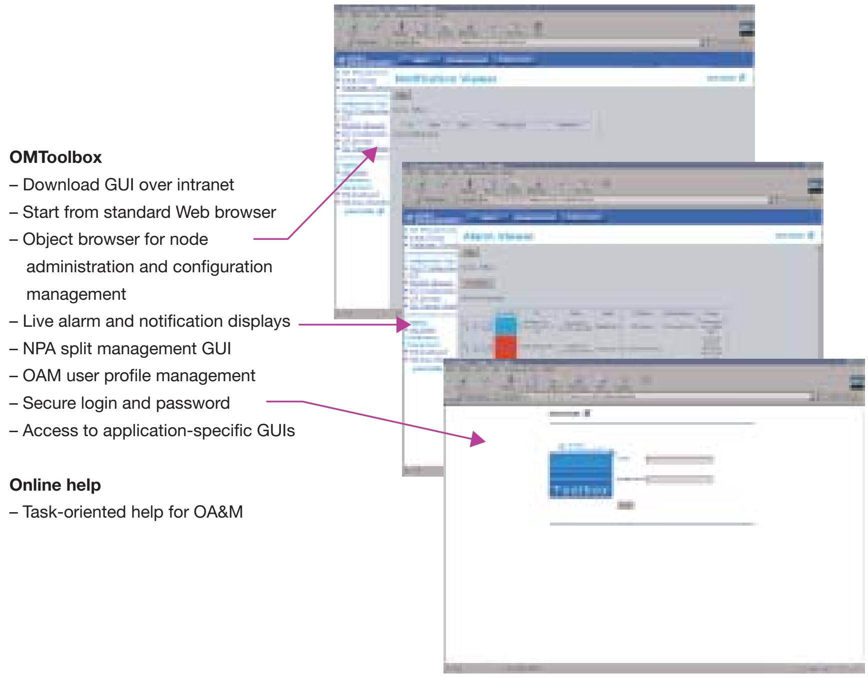

Figure 7
A built-in Web-based node management GUI enables staff to access all O&M tasks online.

Figure 8
The Telecom Server Platform scales linearly. Customers can thus increase capacity as traffic increases.

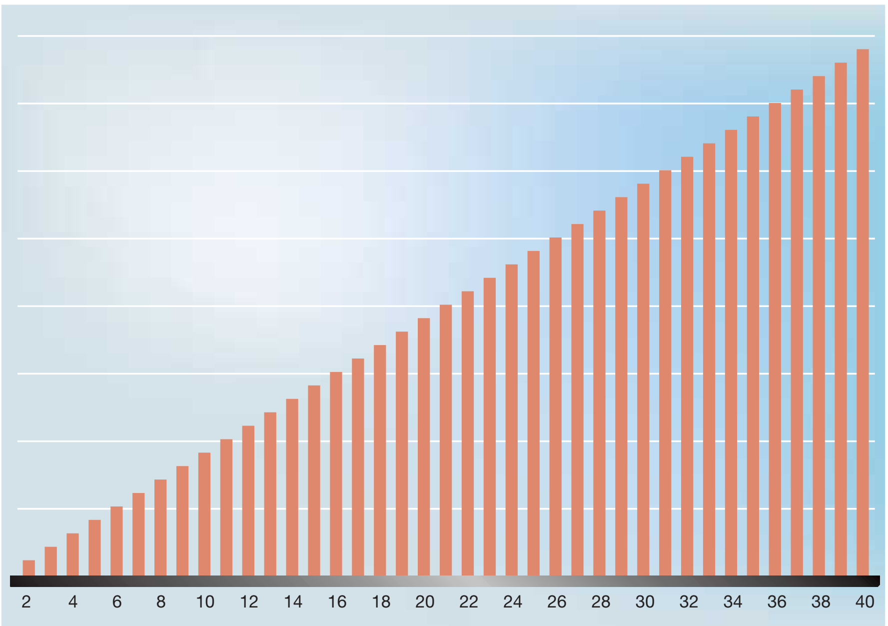

ly integrated solution in a very compact node with small footprint.

The extensive set of O&M and provisioning protocols makes it easy to integrate the Telecom Server Platform into existing O&M infrastructures.

Service network framework

The service layer represents the top layer in Ericsson's three-tier architectural model (Figure 9). The features and functions offered in the service layer pertain both to end-users and the operator of the network infrastructure. They are often offered in the form of an XML Web service interface.

Architectural and design decisions for products and solutions in the service layer are guided through the service network framework (SNF), which is Ericsson's textbook on how system development can be shifted toward horizontally layered systems. The SNF, which provides standards and guidelines for the system structure, capabilities, and interaction, is an architectural framework that consists of reusable designs for products and solutions in the service layer.

The SNF also forms the foundation for reuse. It has always been difficult to reuse objects from different sources, due to a lack of standards that define how objects should be constructed, what properties they should have, and how they should interact. The SNF provides these standards.

Similarly, the SNF provides specifications (or references to specifications) for common management, provisioning, charging, and shared services. Its service specifications are based on open-standard protocols. Apart from the implied need for IP connectivity, the SNF does not stipulate any specific operating environment.

In addition to the actual specifications, the SNF delivers an architectural framework that is designed to accommodate extensions and grow in the future, yielding a tool for expressing a common architectural direction for the service layer. The SNF is influenced by the following mainstream technologies:

- The application server—considerable Internet application development is conducted according to the Web application paradigm in which the application server deploys the application and provides the run-time and characteristics model.
- The directory server—enterprise and Internet computing domains tend more and more to use LDAP-capable directory servers as repositories of user and resource data. Many products and solutions in these domains interoperate easily with LDAP-capable directory servers.
- The distributed computing paradigm—the industry predominantly employs networks of computers to meet the challenges of scalability and availability. Distributed computing patterns, such as *n*-tier architectures, are used in combination with distributed computing technologies, such as CORBA and COM+, to provide system architectures that lend themselves to distribution over large networks of computers.
- The Java language and the Java 2 Enterprise Edition (J2EE) platform—in recent years we have seen a large gain in server-side Internet and enterprise developer momentum for the Java language with its associated server-side platform specification, J2EE.
- The Web services paradigm—there is a growing trend in distributed computing technology to expose traditional remote procedure call (RPC) services via universally accessible extensible markup language (XML) interfaces using methods standardized by the World Wide Web

Figure 9
Ericsson's horizontally layered network architecture.

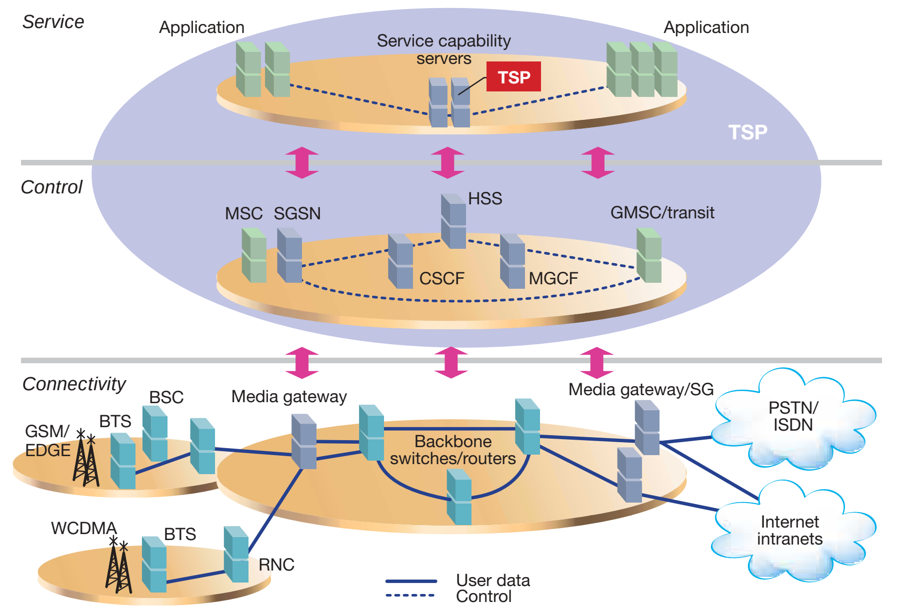

Consortium (W3C) for universal description, discovery and integration (UDDI), the Web services description language (WSDL), and the simple object access protocol (SOAP). The widespread availability of Web services will greatly increase the variety of applications that can be created and deployed in a cost-effective way. The paradigm is simple, universally interoperable and pervasive.

Like the SNF, the operating environment of the Telecom Server Platform also supports and employs these technologies. The products and solutions that leverage the service network framework of reusable designs take a decisive step toward becoming highly scalable, manageable, standards-oriented, open and interoperable, secure and modular. By default, these services are provided in the operating environment of the Telecom Server Platform and are thus facilitated in products and services designed for the service layer.

Collections of products and solutions (systems) that use the technologies and provide the qualities defined above can be deployed to create services networks. Services that target operators and end-users are offered by the services of individual systems. This means that the size of a service network (in terms of the number of constituent systems) can vary and is related to the needs of the business it serves and the nature of the surrounding technical environment in which it is deployed.

Benefits to applications

Ericsson AAA Server

The Ericsson AAA Server provides the authentication, authorization, and accounting (AAA) functions that network operators and service providers need to provide Internet access. The key roles of the server are

- authentication—verification of the identity of an entity (user or application) prevents unauthorized actions and use of resources and services;
- authorization—end-users gain access to the network services and resources that match their profiles; and
- accounting—accounting data is collected and consolidated at the end of an IP session, thereby enabling the service provider to charge end-users on virtually any basis (connection time, amount of data transferred, services accessed, and so on). Billing records can be provided and distributed in customized and flexible formats.

The Ericsson AAA Server provides all the basic AAA functions defined in the RADIUS standard (RFC 2865, 2866). It also supports the DIAMETER protocol, which consists of the base protocol and certain specific DIAMETER applications, such as extensions for network access servers (NAS), Mobile-IP (MIP), strong security, and resource management.

For services to work properly, the Ericsson AAA Server must be available at all times. This puts stringent requirements on the reliability of hardware and software. Lab trials based on a typical traffic model show that the Ericsson AAA Server easily scales to support from 250,000 to more than 5 million subscribers. Moreover, at 50% load it performs more than 2,000 transactions per second.

The inherent characteristics of the Telecom Server Platform are important for the Ericsson AAA Server: it enables the Ericsson AAA Server to scale from small to large systems, gives real-time responses to requests, and is always available.

Ericsson SCS

One enabler in the service network is the Ericsson Service Capability Server (SCS), which uses the capabilities of the Telecom Server Platform for real-time execution of services, open interfaces and scalability. The

Figure 10
Examples of applications that use the Telecom Server Platform in the core network.

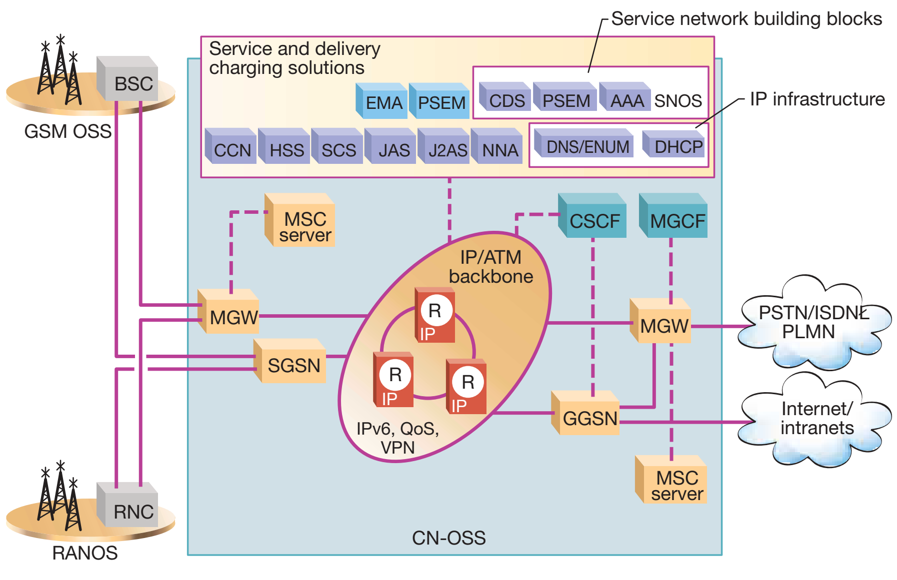

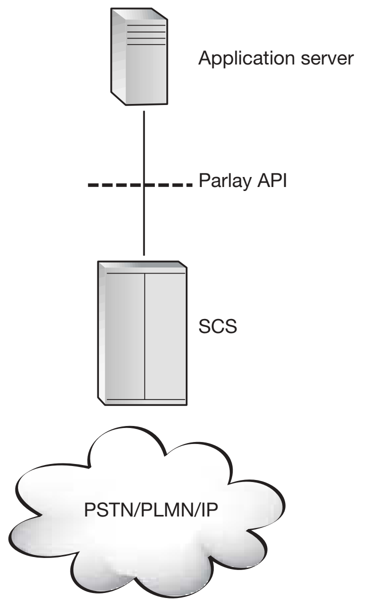

Figure 11
Application service providers (ASP) access the network via the service capability server (SCS) using the standardized Parlay interface.

aims of the SCS are to encourage the development of innovative services and shorten time to market for new applications. It provides open application program interfaces (API) for call control, user interaction and user status and location according to the Parlay/3GPP-OSA standards. Applications and application servers in the network can thus access network service resources regardless of platform and underlying network technology.

The Ericsson SCS is a Parlay gateway that gives various application services transparent access to network capabilities. It gives applications designed according to the Parlay API specification access to network resources in a controlled and secure way. It also functions as a firewall, ensuring that access to the telecommunications network is securely managed. This means that new players, such as application service providers (ASP) and virtual operators, can enter the market and offer services on top of the telecommunications network as a complement to the services of the network operator.

The Ericsson SCS map translates Parlay commands to different network protocols (such as CS1, CS1+, CAPv2, and MAP). In general, the Parlay API is independent of the underlying network. The vendor of the Parlay gateway decides which network protocols will be implemented.

One important role of Parlay is to hide the network complexity from the designer. Provided he complies with the Parlay specification, the designer should not be required to have in-depth knowledge of the underlying networks. Since Parlay is independent of the network and the programming language, the entire market for application development can instead concentrate on developing services. This means that the applications can be developed more quickly

TRADEMARKS

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

UNIX is a registered trademark of the Open Group.

and at less cost. It also means that more innovative applications can be released to the market, helping operators to enrich their service offering, reduce churn, attract more subscribers and increase the use of airtime and revenues.

Ericsson supports application developers through its third-party program in Ericsson Mobility World. This support includes a website with a simulator, and the certification of applications. Where Parlay/OSA is concerned, Ericsson is currently partnering with Solomio, Appium and Wirenix, plus a number of associates.

A customer demonstration center strengthens and supports the sales process by showing that the Ericsson SCS and third-party applications actually work on the network. The customer demonstration center is also used to verify and test third-party applications.

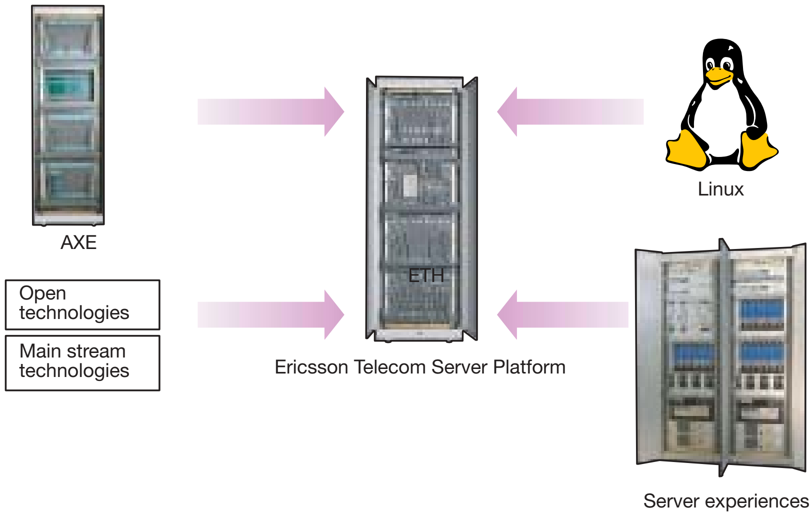

Figure 12
The Ericsson Telecom Server Platform combines the best of the telecommunications world with open Internet technologies.

Conclusion

Ericsson's Telecom Server Platform 4 joins the best of two worlds by integrating IT technologies into a telecommunications-grade server. Many of the applications that use the Telecom Server Platform benefit from its inherent characteristics.

Besides offering the excellent characteristics required for telecommunications systems, the Telecom Server Platform 4 incorporates the latest in open technologies, for example, Linux, which is an important enabler. Ericsson is working actively with Linux standardization bodies, such as the Open Source Development Lab, to make certain that a Linux standard is made available and distributed to the telecommunications industry. The addition of Ericsson's TelORB clusterware to the Linux software gives the system characteristics needed in a telecommunications environment.

REFERENCES

1 Hameleers, H. and Johansson, C.: IP technology in WCDMA/GSM core networks. Ericsson Review Vol. 79(2002):1, pp. 14-27.
2 Granström, P., Olson, S. and Peck, M.: The future of communication using SIP. Ericsson Review Vol. 79(2002):1, pp. 28-35.
3 Ahlforn; G. and Örnulf, E.: Ericsson´s family of carrier-class technologies. Ericsson Review Vol. 78(2001):4, pp. 190-195.
4 Hennert, L. and Larruy, A.: TelORB—The distributed communications operating system. Ericsson Review Vol. 76(1999):3, pp. 156-167.