

AXE 810—The evolution continues

Magnus Enderin, David LeCorney, Mikael Lindberg and Tomas Lundqvist

Thanks primarily to an architecture that was developed to support change, the AXE exchange continues to evolve. Its architecture and modularity have benefited customers—AXE has served as local exchanges and international exchanges, and even in mobile networks to provide mobile switching centers (MSC), home location registers (HLR), and other functions. This has resulted in a total number of about 13,000 exchanges and an all-time-high growth rate.

The modularity of AXE makes it possible to add new functionality in a cost-effective way, but hardware and software R&D must also make the most of new technologies.

This article describes recent adaptations of hardware and software that will prepare AXE for the next generation of networks. The authors focus on a system architecture that will serve as the basis for migration toward a server-gateway architecture for third-generation mobile networks and the next generation of multiservice networks. Adaptations will also enable improvements in existing networks where traffic is growing quickly.

Introduction

In the next few years, networks will evolve toward today's vision of an "all-IP" network. We can see two possible scenarios: one for traditional networks, and one for "next-generation" networks. For traditional networks, the scenario describes what will happen in multiservice networks and in second-

generation mobile networks, such as GSM. The next-generation network scenario describes the development of mobile Internet and fast Internet access in fixed networks.

In traditional networks, evolution is driven by never-ending growth in the need for processing capacity; in mobile networks, by growth in the number of subscribers and usage levels per subscriber. The wireline network is also experiencing a sharp increase in traffic because of Internet "surfing."

In next-generation networks, traditional telephone and data networks will converge to become general networks designed for many different services and modes of access. The convergence of data and telecommunications makes it possible to combine the best of two worlds. Some requirements, such as the need for heightened performance, are fulfilled more easily when development is based on data communications products. Also, the variety of access modes—via second- and third-generation mobile networks, the multiservice networks, and broadband—will necessitate the coexistence of different transmission formats. Thus, gateways will be required at network interconnection points.¹

BOX A, TERMS AND ABBREVIATIONS

AAL	ATM adaptation layer	GCP	Gateway control protocol	RAID	Redundant array of independent disks
ACS	Adjunct computer subsystem	GDDM	Generic datacom device magazine (subrack)	RAM	Random access memory
ALI	ATM link interface	GDM	Generic device magazine (subrack)	RISC	Reduced instruction set computer
AM	Application module	GEM	Generic Ericsson magazine (subrack)	RLSES	Resource layer service specification
AP	Adjunct processor	GS	Group switch	RM	Resource module
APC	AM protocol carrier	GSM	Global system for mobile communication	RMP	Resource module platform
APIO	AXE central processor IO	GSS	Group switch subsystem	RP	Regional processor
APSI	Application program service interface	HDL	High-level data-link control	RPC	Remote procedure call
ASIC	Application-specific integrated circuit	HSB	Hot standby	RPP	Regional processor with PCI interface
ATM	Asynchronous transfer mode	IO	Input-output	SCB-RP	Support and connection board - RP
BICC	Bearer-independent call control	IOG	Input-output group	SDH	Synchronous digital hierarchy
BIST	Built-in self-test	IP	Internet protocol	SES	Service specification
BSC	Base station controller	IPN	Interplatform network	SONET	Synchronous optical network
C7	CCITT (now ITU-T) no. 7, a common-channel signaling system	ISDN	Integrated services digital network	SS7	Signaling system no. 7
CAS	Channel-associated signaling	ISP	Internet service provider	STM	Synchronous transfer mode
CP	Central processor	IWU	Interworking unit	STOC	Signaling terminal open communication
cPCI	Compact peripheral component interconnect	MGW	Media gateway	TCP	Transmission control protocol
CPP	Cello packet platform	MIPS	Million instructions per second	TDMA	Time-division multiple access
DL	Digital link	MSC	Mobile switching center	TRA	Transcoder
DLEB	Digital link multiplexer board in the GEM	MSCS	Microsoft cluster server	TRC	Transceiver controller
DSA	Dynamic size alteration	MSP	Multiplex section protection	TSP	The server platform
DTI	Data transmission interworking	MTP	Message transfer part	TU 11, 12	Typical urban 11 (12) km/hr
ECP	Echo canceller in pool	MUP	Multiple position (timeslot)	UMTS	Universal mobile telecommunications system
ET	Exchange terminal	NGS	Next-generation switch	VCI	Virtual channel identifier
ETSI	European Telecommunications Standards Institute	OSS	Operations support system	VPI	Virtual path identifier
FTP	File transfer protocol	PCM	Pulse code modulation	XDB	Switch distributed board
		PDH	Plesiochronous digital hierarchy	XSS	Existing source system
		PLMN	Public land mobile network	WCDMA	Wideband code-division multiple access
		PSTN	Public switched telephone network		
		PVC	Permanent virtual circuit		

Typical of next-generation network architecture is the separation of connectivity and communication or control services. This new new architecture will appear first in GSM and WCDMA/UMTS core networks, where AXE exchanges will continue to serve as MSCs. For multiservice networks, telephony servers will become hybrid nodes which consist of an AXE exchange that uses an AXD 301 to handle packet-switched data.

Based on the traditional and next-generation scenarios, network evolution will demand increased processing capacity, greater switching capacity, and conversion to packet-switched technology. In addition, new ways of doing business will emerge as virtual operators pay for the right to use telecom equipment owned by other operators. Similarly, more operators are expected to enter the market, putting additional demands on charging and accounting functions.

To succeed in today's telecommunications market, operators must be able to provide their users with new functionality in networks and complete coverage by mobile systems at the same or a lower cost as time progresses. Operators put demands on the return on investment, cost of ownership, footprint (multiple nodes per site), plug-and-play functionality (short lead times in the field), and quality of software packages (quality of service).

This article describes the latest developments made in the AXE platform to meet these requirements, and what new products will be launched. For example, to shorten the software development cycle, a layered architecture was introduced in the early 1990s, resulting in application modules (AM). Current work on application modules focuses on the server-gateway split. AXE hardware is also undergoing a major architectural transformation to further reduce the footprint, power consumption, and number of board types.

Ericsson's goal is to cut the time to customer for AXE by targeting improvements in production, transportation, and installation. Far-reaching rationalization has been achieved through the introduction of the generic Ericsson magazine (GEM), an open, flexible, scalable, high-capacity magazine (subrack). A new distributed, non-blocking group switch (GS890) has also been introduced, as have new devices, such as the ATM link interface (ALI) and ET155-1, which enable AXE to serve as a gateway to an ATM

network. Also, a new input-output (IO) system, called the APG40, has been developed using products from mainstream data communications vendors.

The application platform

The traffic-handling part of AXE has a two-layer architecture: the application platform, which consists of hardware and software; and the traffic applications, which solely consist of software. The application platform can be seen in terms of hardware and software, which present different but complementary views of the architecture.

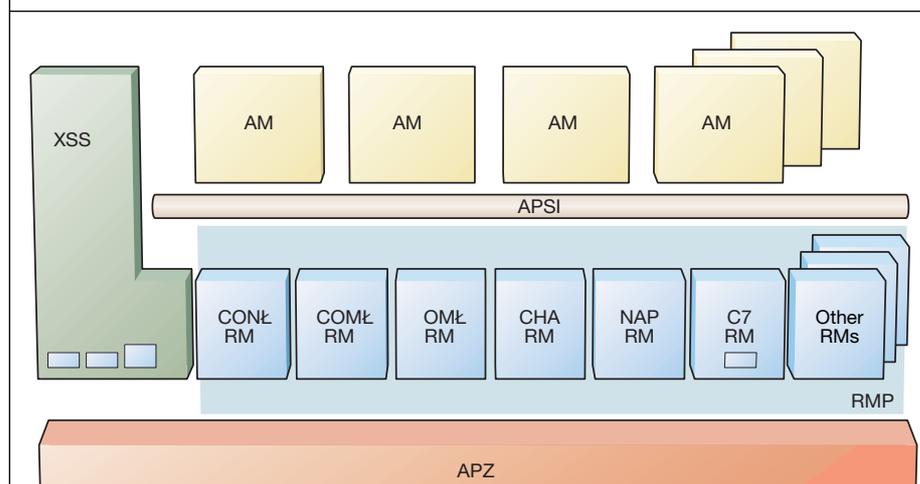
The software view

The AXE application software has continued to evolve since it was first layered and modularized in the early 1990s. Market demand for shorter lead-times was the main force driving the change to layers and modularization.

Application modules

Traffic applications are encapsulated in application modules that use the application-platform software—called the resource module platform (RMP). There is no direct communication between application modules. All communication between them takes place via protocols, where the protocol carriers are implemented in the resource module platform (Figure 1).

Figure 1
AXE 810 application software architecture.



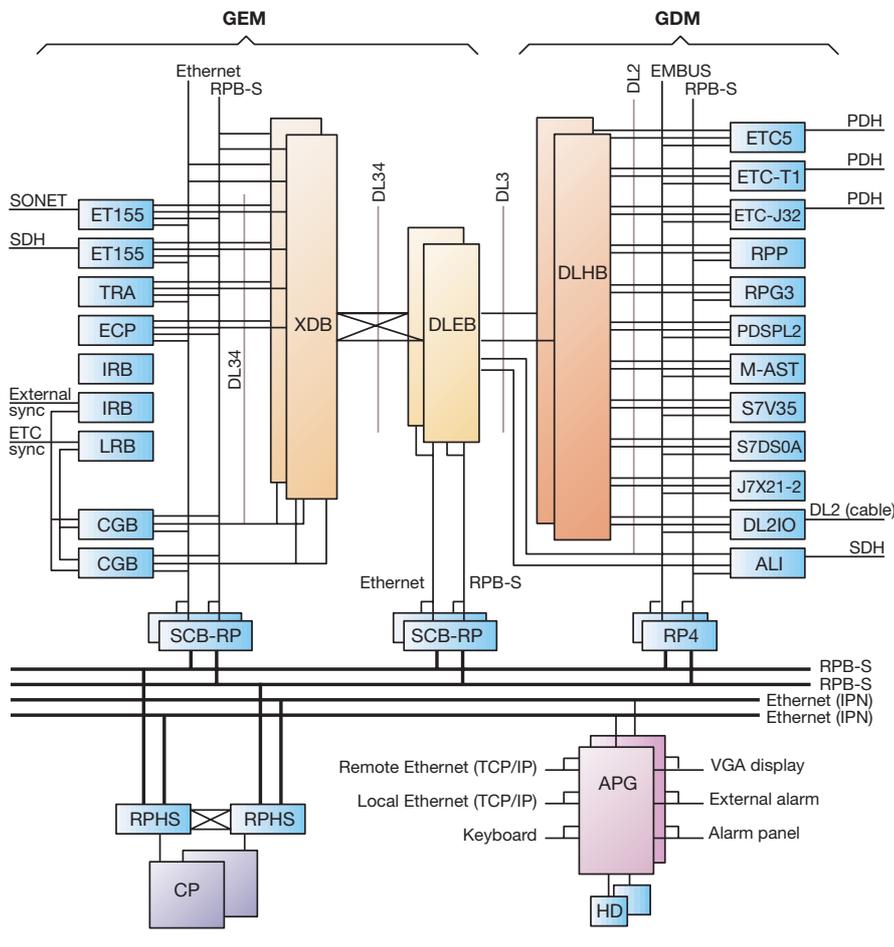


Figure 2
Hardware architectural overview.

BOX B, CURRENT RESOURCE MODULES

- Connection RM (CONRM)
- Communication RM (COMRM)
- Operation and maintenance RM (OMRM)
- Network access products RM (NAPRM)
- Pooled devices RM (PDRM)
- Common channel signaling RM (C7RM)
- Charging support RM (CHARM)

The interface of the resource module platform with the application modules, called the application platform service interface (APSI), is a formalized client-server interface. Subsets of the APSI, or individual interfaces, are called service specifications (SES). Each service specification provides a service for the application modules, some of the most important services being:

- connection service—for setting up a physical connection;
- communication service—for communication between application modules;
- object manager—for operation and maintenance;
- AM protocol carrier (APC);
- charging service; and
- services for signaling no. 7 (SS7), such as the message transfer part service.

The total concept, which consists of the

RMP, APSI, the AMs, and the inter-AM protocols, is known as the application modularity concept. It is largely thanks to this concept that the system lifetime of AXE consistently exceeds expectations—the AXE system is 30 years old, yet we see no end to its potential evolution.

Resource modules

When the RMP was first introduced, it was small in comparison to the large volume of existing software, called existing source system (XSS). Following several RMP development projects, the migration from the earlier architecture to the AM architecture is now virtually complete. The resource module platform, in turn, has become large and complex, leading to a refinement of the AM concept and the division of the RMP into several resource modules (RM). The interfaces provided by the resource modules are formalized as resource-layer service specifications (RLSES).

A priority in the development of RMs and AMs (known collectively as system modules) is to specify the interfaces (SES and RLSES) early in the development process. When the interfaces are “frozen”, the separate system modules can be designed independently of one another, often at different geographical locations, as is common in the Ericsson organization.

The second major advantage for development of applications is that application-platform hardware is now associated with specific resource modules and controlled by them. Application modules simply request services via the APSI. Hardware is “owned” by the software in the respective resource modules.

Hardware

The hardware (Figure 2) revolves around the GS890 group switch, which has 512 K multiple positions, each with a bit rate of 64 kbit/s.

The biggest change in the hardware architecture is the addition of a new exchange interface (the DL34) to existing DL3 and DL2 interfaces. The DL34 interface supports from 128 to 2,688 timeslots to each device board, in steps of 128, via a backplane interface in the generic Ericsson magazine running at 222 Mbit/s. This interface made it possible to improve the devices connected to the group switch, further reducing input power, cabling, the number of board types, footprints, installation time, and other parameters. The high-speed DL34 in-

terface has also facilitated a more efficient version of the group switch itself, also located in the GEM.

Another architectural change (Figure 2) is the interplatform network (IPN). In the first phase, the IPN will be used to speed up communication between the central processor (CP) and adjunct processors (AP). The interface is based on fast Ethernet. In a second phase, it will be upgraded to Gigabit speed.

All devices that support the DL34 interface can be mixed in the GEM more or less without limit. The maximum capacity of each GEM is 16 K multiple-position time-slots (MUP), which corresponds to eight STM-1 ET boards, for example. Physically, there is space for 22 device boards in one GEM.

If a switch is needed that is larger than 16 K MUPs, or when the number of devices exceeds 22, additional GEMs must be added. These can be configured without interruption while the system is processing traffic. An exchange based on the GEM is linearly expandable. The maximum switch size is 512 K MUPs at normal rate (64 kbit/s) or 128 K MUPs at a subordinate rate ($n \times 8$ kbit/s).

By inserting a pair of digital link multiplexer boards (DLEB) in the GEM, we can convert the DL34 backplane interface into four DL3 cable interfaces. This allows all of today's devices in the generic device magazine (GDM), generic datacom device magazine (GDDM), and other formats to be used with the new switch core.

Software evolution

The original AXE concept has enabled ongoing evolution and modernization. Legacy software can be reused or modified, subsystems and system modules can be modernized or completely replaced, interfaces can be left unchanged or redesigned to meet new needs, and new architectures can be introduced. Yet AXE remains fundamentally AXE.

Several software development programs are under way in parallel. The APZ processor is being adapted to serve as a multi-processor CP to multiply call-handling capacity. Cost of ownership is being reduced. Ericsson is improving its internal management of software to reduce time-to-market and time-to-customer. These activities are in addition to the four programs described below.

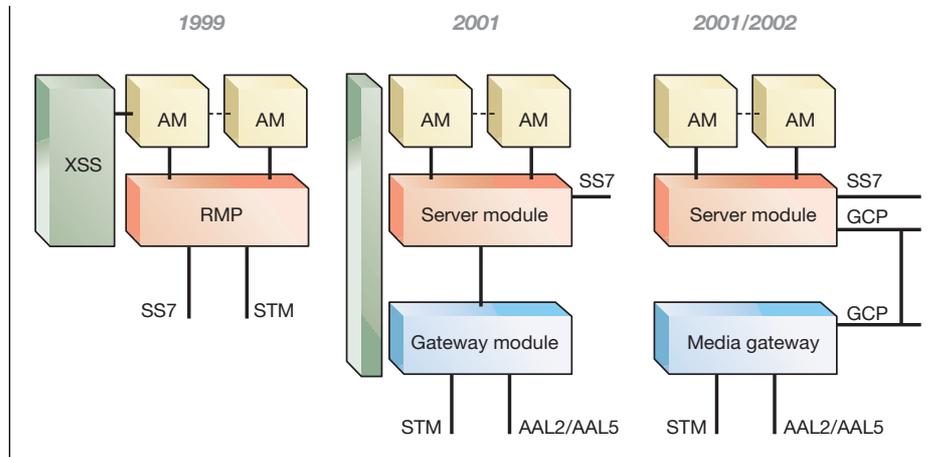
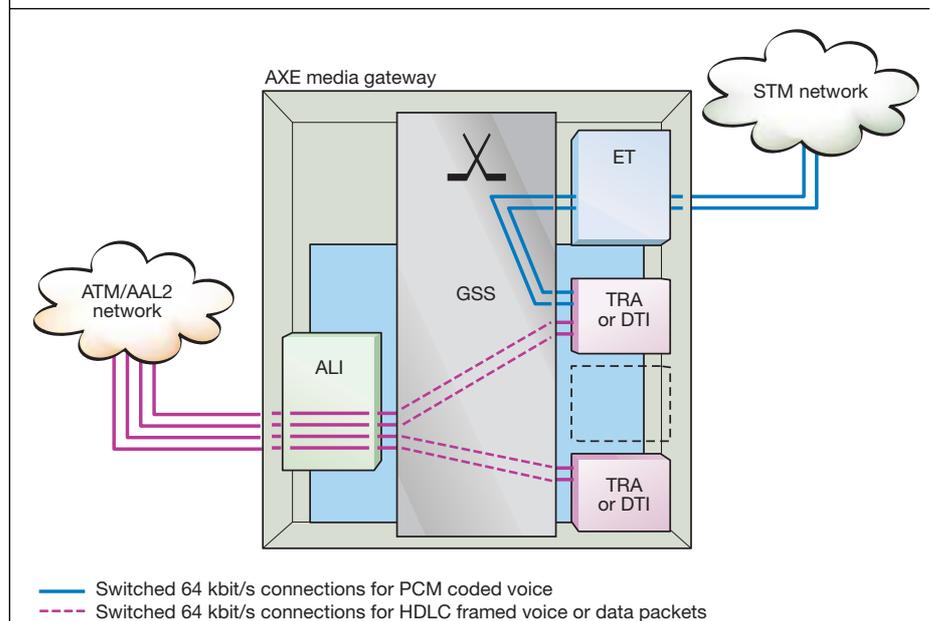


Figure 3
Software evolution toward the server-gateway split.

Software migration

The network architecture is changing. In particular, we see two new types of node: the server (or media-gateway controller) and the media gateway itself, resulting in what is known as the “server-gateway split.” In order to meet the demands of the new architecture, the software in the application platform is being migrated from a traditional telecom environment—which is STM-based, mainly for voice traffic—to-

Figure 4
System architecture, ATM interworking function.



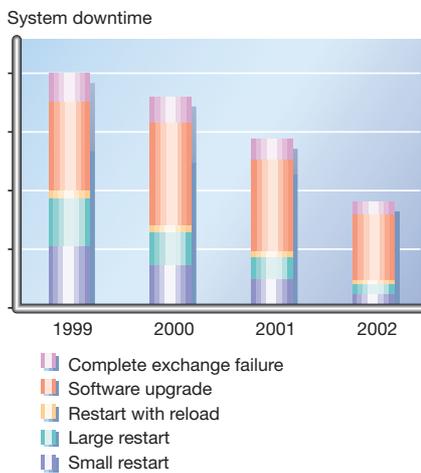


Figure 5
Reduction in system downtime.

ward a packet-based environment, initially for ATM, and subsequently for IP.

In this new architecture, the server controls the call or session, whereas the media gateway provides the bearer service, transmitting information over the connectivity layer. An advantage of separating call control from bearer control is that bearer technology can be upgraded from STM via ATM to IP without affecting the call control.

The original AM architecture has proved itself to be future-proof and is well suited to accommodate the server-gateway split. The split has been implemented in system modules (new and modified AMs and RMs) without any fundamental changes to the software architecture. For example, new RMs for bearer-access service and for TDM access are currently being developed for the WCDMA/UMTS program.

Two separate migration tracks are currently being followed. For the WCDMA/UMTS program, AXE is the natural choice of technology for the server, thanks to its high capacity and high level of functionality. AXE is also available as a combined server/media-gateway node, where the server part controls its own gateway part. By contrast, the Cello packet platform (CPP) is the preferred choice of technology for ATM-based media gateways. In such a network, the server nodes communicate using bearer-independent call control (BICC). The server controls the gateway using the gateway control protocol (GCP). ATM-based gateways communicate using Q.AAL2.

The second migration track is for the next-generation switch (NGS) program for the multiservice network. The server node is a hybrid node made up of co-located AXE and AXD 301 systems. The gateway node consists of an ATM switch (also an AXD 301) for AAL1 connections, and is controlled by the server using the gateway control protocol.

In-service performance

Robustness, or in-service performance, has long been a priority in AXE development, and considerable improvements have been and are still being made. No software is 100% reliable, but robust systems can be achieved by means of recovery mechanisms that minimize the effects of software malfunctions or program exceptions.

One of the most powerful robustness mechanisms is known as *forlopp*, which is a recovery mechanism at the transaction level. (Forlopp is an anglicization of the Swedish

word *förlopp*, roughly the “sequence of events.”) The forlopp mechanism is a low-level recovery mechanism by which only the transaction affected by the software malfunction is released. This minimizes the disturbance to the overall system. The forlopp mechanism, which has been refined in several stages and is now applied to all traffic handling in AXE and to many operation and maintenance functions, has significantly reduced the number of recoveries at the system level.

System restart is used as a recovery mechanism for certain faults and for system upgrades. The restart mechanism itself has been optimized in several ways, so that in the few cases that still require restart, its duration is as short as possible.

Activities for restarting software in the regional processor (RP) have been especially improved. For example, regional processors are restarted through minimal restart by default—that is, with the suppression of unnecessary actions. Complete restarts are performed on regional processors only when necessary. Regional processors are restarted asynchronously, which means that no regional processor has to wait for any other regional processor to become ready for a restart phase. These improvements have significantly reduced the time consumed when restarting application hardware.

If necessary, AXE can, via the IPN, be reloaded from the backup copy of the system on the APG40 file system. This approach is ten times faster than the design that applied before the APG40 and IPN were introduced. Also, the time it takes to make a backup copy of the system on APG40 is one-fifth of what it was before.

The restart duration for a system that is upgrading to new software has also been reduced. The most recent improvements involve RP software, which is now preloaded prior to an upgrade, instead of during the upgrade, thus saving restart time.

Major improvements have been made to the function-change mechanism used for system upgrades. Inside the central processor, the time needed for data transfer prior to a changeover to new software has been cut from hours to typically ten minutes. (The data transfer does not disturb traffic handling, but exchange data cannot be changed during data transfer.)

Another improvement applies to the retrofit of new CPs that replace old ones. The bandwidth of the connection between the processors has been expanded by using re-

gional processor with industry-standard PCI interface (RPP) and Ethernet connections, thereby reducing data-transfer times ten fold. As a consequence of these many improvements, system downtime has been reduced significantly in recent years, a trend that is sure to continue (Figure 5).

Software licensing

The main force driving software licensing is customer preferences to pay for access to specific functions, features, and capacity. Thus, “software licensing” is really the “licensing of features and capacity.” The second driving force is Ericsson itself, because it is in the company’s interest to deliver standard nodes (those that contain standardized software and hardware).

With software licensing, Ericsson delivers a standard node for a particular market or market segment, such as an MSC for a cellular system using time-division multiple access (TDMA). This standard node consists of a complete software configuration with a standard hardware configuration that is deliberately over-dimensioned. Ericsson then limits the call capacity and functionality of the node by means of software keys. When a customer requests more call capacity, Ericsson personnel execute a password-protected command to increase the capacity in the node to the new level.

This method of increasing capacity, referred to as “traffic-based pricing” or “pay-as-you-grow,” is much simpler than the traditional method, in which Ericsson personnel would deliver, install, and test new hardware on a live node. The commands can also be executed remotely from a maintenance center, making a visit to the site unnecessary. This arrangement benefits customers and Ericsson, especially in the mobile market where growth can be rapid and unpredictable.

Similarly, in the future, software will contain functions and features for which customers can purchase licenses on an *ad hoc* basis. The software for these functions and features will be unlocked using password-protected commands. This method of delivering software is much simpler than the traditional method, by which new software is delivered in the form of software upgrades.

The licensing of call capacity is already available on AXE nodes, and the introduction of a general mechanism to handle all software licensing of functionality and capacity is planned. Some of the characteristics of this general mechanism are as follows:

- Software licensing will be common to AXE 810, TSP, CPP, and other Ericsson platforms.
- Ericsson will maintain a central license register, from which new licenses can easily be issued.
- License keys will be distributed electronically, in encrypted form.

Improved handling

Several improvements to the operational handling of AXE have recently been made, such as parameter handling and hardware inventory. Other handling improvements are being planned, such as plug-and-play functionality. These improvements reduce operational costs and also often have a positive impact on the ISP, since they reduce the risk of human error.

One of the most important improvements to handling is called dynamic size alteration (DSA). The size of many data records in AXE is traffic-related, since the number of individuals, or instances, in the record varies from node to node, as well as over time, in accordance with the capacity demands put on the node. Thanks to dynamic size alteration, the number of individuals in the record is increased automatically, up to a preset limit, without any intervention from the maintenance technician. However, as a warning, AXE issues an alarm when, say, 75% of the reserved data capacity has been used up, so that the technician can raise the limit as necessary.

Dynamic size alteration also simplifies the handling of this kind of alteration, since the alarm (referred to above) specifies an action list of the data records whose size needs to be altered (increased). The technician merely enters a single command to increase the sizes of all these data records.

Hardware evolution

Generic Ericsson magazine

Another far-reaching improvement is the GEM, a high-capacity, flexible, and scalable magazine (subrack) that anticipates future developments (Figure 6). GEM-based nodes will be smaller, dissipate less power, and have greater maximum capacity. Their implementation will dramatically improve the cost of ownership and cut time-to-customer for AXE.

In previous versions of AXE, each function was located in a separate magazine, and

Figure 6
The GEM, a high-capacity, flexible and scalable magazine (subrack).



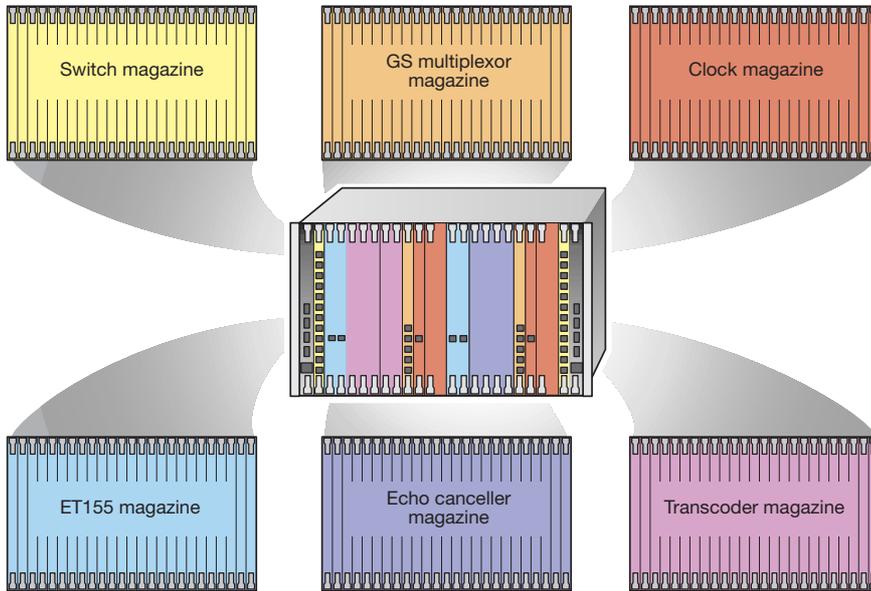
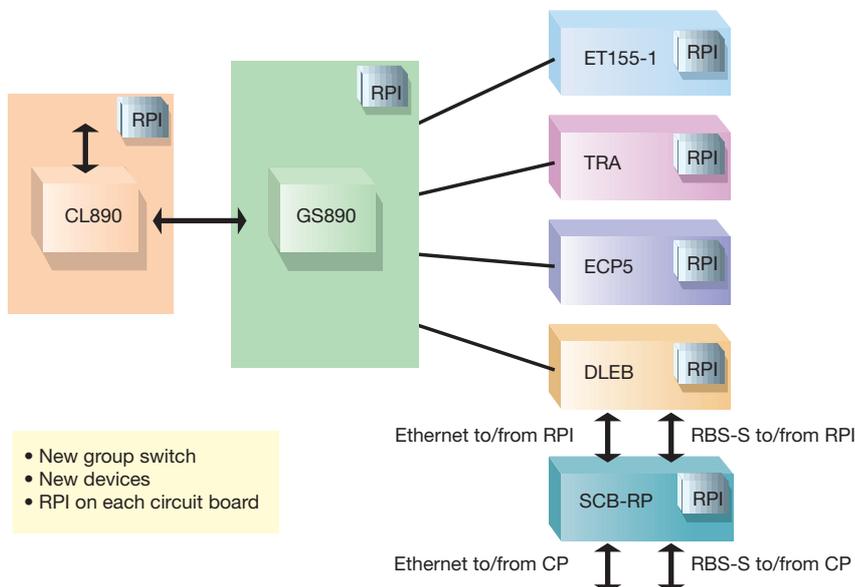


Figure 7
Many different functions can be mixed in the GEM.

the associated regional processors were on separate boards in the magazine. The GEM and the boards developed for the GEM represent fundamental change.

Several functions are mixed in a single magazine, and the RP function is integrated on each board. This makes the most of

Figure 8
The GEM concept can be used to build anything from an extremely small node to an extremely large one, using the same boards and magazine.



advances in technology to modify the architecture rather than merely shrink the hardware. One advantage of this change is that considerably fewer magazine and board types are needed in each node. A second advantage is that many internal AXE interfaces have been moved to the GEM backplane, reducing the need for cables. As shown in Figure 7, many different functions can be mixed in the GEM.

Boards developed for the GEM have high capacity, and because more functions are concentrated on each board, the total number of boards can be reduced. Mixing these high-capacity boards in a single magazine and removing many cables by moving the interfaces to the GEM backplane produces a more compact system.

The GEM concept can be used to build anything from an extremely small node to an extremely large one, using the same boards and magazine (Figure 8). The smallest node will consist of one GEM, and the largest, 32 GEMs. As capacity requirements grow, the node can be expanded smoothly and without interruption of traffic by adding one or more GEMs with a suitable combination of devices. A GEM can house

- two SCB-RPs, providing an Ethernet switch, control function, maintenance support, and power distribution;
- two switch boards (XDB), providing a 16 K plane duplicated group switch;
- up to 22 device boards with 15 mm spacing, such as ET155, ECP, or TRA;
- pairs of DLEBs, providing multiplexing functionality for DL3 cable interfaces, placed in device slots; and
- CL890 clock modules placed in device slots.

The GEM, which has been developed for use in the BYB 501 equipment practice, provides several infrastructure functions via the backplane for the boards housed in the magazine:

- duplicated power distribution;
- duplicated group switch connection;
- duplicated serial RP bus;
- maintenance bus for hardware inventory and fault indication;
- duplicated Ethernet, 10 or 100 Mbit/s; and
- extremely robust backplane connectors.

The option to use Ethernet as the control protocol makes the GEM an open magazine prepared for future developments. Another advantage of the GEM is that it is generic and can be used in non-AXE products from Ericsson.

The GS890 group switch

The GS890 fully supports the GEM concept and its goal of dramatically improving cost of ownership and time-to-customer. The figures in Table 1 show an extraordinary reduction in power dissipation and number of boards and cables. This was achieved through a combination of recent advances in ASIC technology and high-speed interfaces with architectural modifications.

The maximum size of the GS890 has been increased considerably, making it possible to build larger nodes. For many network configurations, it would be better to build and maintain a small number of large nodes. The switch is also strictly non-blocking, so many limitations on the configuration of the node and the surrounding network have been removed.

The subrate switch function in AXE, which enables switching at 8 kbit/s, has also been modified. Primarily used in GSM networks, this function makes efficient use of pooled transcoders and transmission resources surrounding the base station controller (BSC). The subrate function was previously implemented as a pooled function, but in GS890 (Box C) it has been integrated into the switch core and increased to 128 K. The advantages are that much larger BSCs are feasible, less equipment is needed, and the traffic selection algorithms are easier and faster because the pooled switch has been removed.

GS890 group switch connection

All existing group switch interfaces are supported, making it possible to connect all current devices to the GS890. The DL2, which is a cable and backplane interface, carries 32 timeslots. Similarly, the DL3, which is a cable interface used by the ATM inter-

TABLE 1, IMPROVEMENTS IN THE GROUP SWITCH

Characteristics	GS12 (Max 128 K)	GS890 (Max 512 K)	Reduction
Power dissipation in 128 K MUP group switch core	1200 W	250 W	81%
Internal cables in a 128 K MUP group switch core, and including clock distribution	1,024 (12-pair cables) and 48 (4-pair cables)	88 (4-pair cables)	92%
Number of boards in a 128 K MUP group switch core, excluding clock boards	320	16	95%
Equivalent figures for a 512 K group switch	--	1,000 W, 448 (4-pair cables) and 64 boards	--

working unit (IWU) and GDM, carries 512 timeslots. The new DL34 backplane interface supports new high-capacity devices developed for the GEM.

GS890 hardware

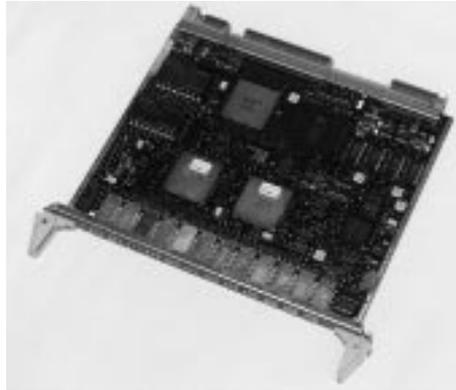
The main functionality of the GS890 is implemented in two new ASICs: one core ASIC, which handles the switching; and one multiplexer ASIC, which concentrates traffic into the core. Both ASICs have been designed using 0.18 µm CMOS technology.

The core ASIC contains 4 Mbit of RAM, runs at up to 333 MHz, has 500 K gate logic and 20 high-speed interfaces running at 666 Mbit/s, allowing the ASIC to handle up to 13 Gigabits of data every second. In a fully expanded switch matrix, 64 of these ASICs will be needed per plane. Standard four-pair cable is used to interconnect the switch matrix elements located in different magazines. The cable can be up to 20 meters long and carries 1.3 Gigabits of bidirectional data per second, two 666 Mbit/s interfaces in each

BOX C, MAIN CHARACTERISTICS OF THE GS890

- Maximum size 512 K MUP, equivalent to 524,288 (64 kbit/s) ports
- Strictly non-blocking architecture, regardless of traffic type
- Fully distributed group switch, with switch matrix distributed among up to 32 GEMs
- Integrated subrate switching capability up to 128 K MUP, equivalent to 1,048,576 (8 kbit/s) ports
- Improved maintenance support by logic BIST and RAM BIST in all ASICs, reducing maintenance time and the time that the system runs on one plane
- Hardware support for fast dynamic fault isolation
- New and flexible high-speed-device interface
- Device protection support with no wasted capacity (used for ET155-1)
- New high-speed, internal group switch interface
- Rear cabling removed

Figure 9
Photograph of the switch distributed board (XDB).



direction. The new hardware requires two new board types:

- the switch board; and
- the multiplexing board, which makes it possible to connect existing devices.

In addition, two more boards have been developed to support a smooth migration from the GS12 to the GS890. The migration consists of replacing one board in each GS magazine, thereby converting it into a multiplexing magazine. The concentrated traffic is then connected to the new GS890 core. These boards help safeguard investments already made in AXE. The GS890 (Figure 9) supports the following types of traffic:

- normal rate, 64 kbit/s;
- non-contiguous wideband, $n \times 64$ kbit/s up to 2 Mbit/s;

- subrate, 8 kbit/s or 16 kbit/s;
- wideband on subrate, $n \times 8$ kbit/s up to 256 kbit/s; and
- broadband connections up to 8 Mbit/s.

The CL890 synchronization equipment, also developed to support the GEM concept, consists of

- duplicated clock modules;
- up to two highly accurate reference clock modules; and
- up to two incoming clock reference boards for connecting additional clock references.

Devices

RPG3

A new version of the RPG2 with slightly improved CPU performance has been designed that only occupies one slot instead of two. The RPG3, which is still located in the GDM, is mainly used for different kinds of signaling application.

ET155-1

The ET155-1 is an SDH exchange terminal that will occupy one slot in the GEM and will be available for ETSI and ANSI/SONET standards. The ETSI variant will carry up to 63 TU12s and support both optical and electrical line interfaces. In ANSI/SONET mode, the board will be able to handle up to 84 TU11s.

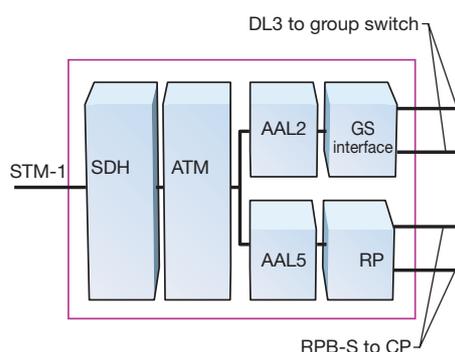
The board fully supports channel-associated signaling (CAS) in both standards by extracting the signaling information from the incoming PDH frames and sending it through the switch to pooled signaling terminals. The ET155-1 can work either as a non-redundant single-board exchange terminal or in tandem, supporting MSP 1+1 and equipment protection.

One motherboard, which contains the main functionality, and two generic line interface modules, containing the optical respectively electrical line interface, have been developed. An ET155-1 board is produced by mounting the required module on the motherboard. The line interface modules are generic and will also be used in non-AXE products from Ericsson.

Transcoders and echo cancellers

A big improvement has been made in the transcoder area. A new board designed for the GEM supports 192 channels for all the codecs used in GSM and TDMA systems. The same board is also used for echo cancellation, in which case it supports 128 channels.

Figure 10
Block diagram of the ATM link interface (ALI).



A fully equipped GEM with transcoders supports 4,224 transcoder channels; when fully equipped with echo cancellers, the GEM handles 2,816 channels.

Future GEM devices

The devices described here were designed for the first release of the GEM. Future releases will include more devices ported from today's GDM and GDDM. Technology allows more traffic to be handled by each board in AXE, so the boards can efficiently take advantage of the configurable bandwidth of the DL34.

ATM interworking function

With the introduction of WCDMA/UMTS, AXE will be used as a combined server and media gateway and purely as a server. In the downlink direction (to the radio access network), AXE will be attached to the *Iu* interface, a new interface based on ATM. ATM will also be used in the core networks.

An ATM interworking unit has been implemented to add an ATM interface to new and legacy AXE nodes. The interworking unit enables existing AXE-based MSCs to be upgraded easily to handle signaling and the user plane in ATM networks. The interworking unit is connected directly to the group switch interface for the user plane, and the signaling information is transferred to the central processor via the RP bus (Figure 10).

The interworking unit provides a 155 Mbit/s ATM/SDH interface to AXE. New hardware and software for AXE enable it to perform as a gateway to an ATM net-

work capable of handling voice, data, and signaling. The first release of the ATM interworking unit supports three types of payload in the ATM network:

- coded voice on AAL2;
- circuit-switched data on AAL2; and
- signaling on AAL5.

The device that implements the interworking unit is called the ATM link interface (ALI). This hardware unit is divided into three boards that are mounted together on a plug-in unit which can be located in a full-height GDM that provides RP bus, -48 V and maintenance bus via the backplane. The unit is then connected via DL3 cables to the group switch. The fiber for the STM-1 interface is connected to the external equipment.

Adjunct processor platform

Statistics from the telephony or radio parts of a network must be collected to optimize network efficiency. At the same time, virtual operators are willing to pay for the right to use telecom equipment owned by other operators, and network operators are demanding fast, almost instant billing (such as with prepaid cards). These requirements demand processes with close to real-time execution.

In terms of operation and maintenance, it is better to have one big node in a network instead of several small ones, although this demands greater switching capacity. To reduce operators' cost per subscriber, big switches are being planned. This has spurred the development of new high-end AXE switches. There are numerous methods of

BOX D, ATM VIA AXE

To transmit coded voice over ATM, a permanent virtual circuit (PVC) is set up between the ALI and the remote end. The PVC is specified by its VPI/VCI address. The first ALI release allows a total of eight PVCs to be set up in the user plane.

Within each PVC, AAL2 connections (based on the I.363.2 standard) can be set up and released using Q.2630.1 (Q.AAL2) signaling. Up to 248 AAL2 connections can be set up in each PVC.

The C7 signaling links for supporting Q.AAL2 and other user parts are configured as another type of PVC. For each ALI, up to 64 signaling PVCs can be configured to carry adaptation layer AAL5, based on the I.363.5 standard. Each of these AAL5 channels carries a C7 sig-

naling link using SSCOP (Q.2110) as the link protocol.

Adaptation layers for both ETSI (SSCF-NNI, Q.2140) and Japanese C7 standards (JT-Q.2140) are provided.

The architecture for supporting circuit-switched data is similar to that for coded voice except that the payload packets are routed to a data transmission interworking (DTI) unit instead of a transcoder. The first release of ALI supports a total of 32 data channels.

In the future, the ATM interworking concept might be used to implement new functions such as 64 kbit/s voice on AAL1 or on AAL2 based on I.366.2. Another function in the works is "soft PVCs," which can be operated using B-ISUP signaling.

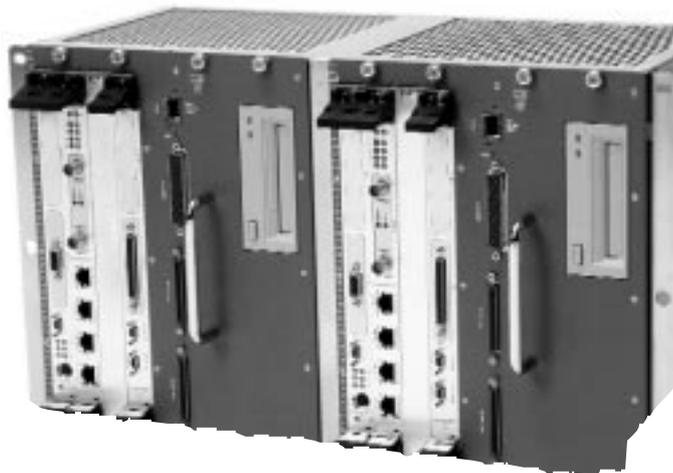


Figure 11
Photograph of a large APG40 cluster.

enhancing performance, since AXE processing power is distributed among several processors (CP, RP, and AP), and research and development can be pursued along different lines in parallel.

APG40 platform for near real-time applications

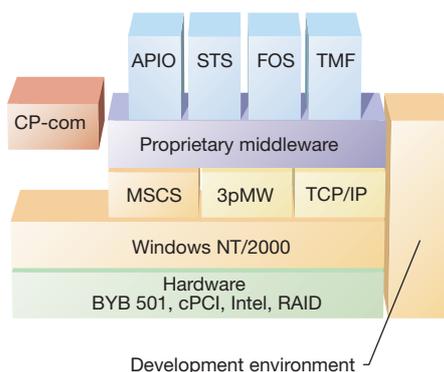
The most widely deployed AXE IO system, IOG20, was introduced in 1997. The IOG20 is a proprietary product for the CP file system, man-machine communication, boot server, and statistics functionality. In

1998, the first adjunct processor—the APG30—was released, primarily as a billing system. The APG30 is widely deployed in TDMA systems as a platform for billing but also for operation and maintenance of AXE as an IOG replacement. The APG30, which is based on a commercial computer from an external vendor, uses MIPS RISC microprocessors and NonStop-UX O/S.

The APG40, successor to both the IOG20 and APG30, is a platform for near real-time applications. Three needs have driven the development of the APG40 (Figure 11). One need is to offer a platform for new data-handling tasks and applications that are best performed at the network element level by a standard computer system. Another need is for improved performance, as a result of increased traffic per switch and new types of data extracted from each switch. The third need is to reduce time-to-market and time-to-customer.

The APG40 is a platform for the AXE central processor IO functions (APIO) that were inherited from the IOGs. It is also a platform for billing (for example, FOS) and statistic data collecting (for instance, STS), storage, processing and output from the AXE switch. For example, the APG40 can be a platform for collecting data related to in-service performance from the central processor. It can also format that data for distribution to an operations management

Figure 12
APG40 building blocks and development environment.



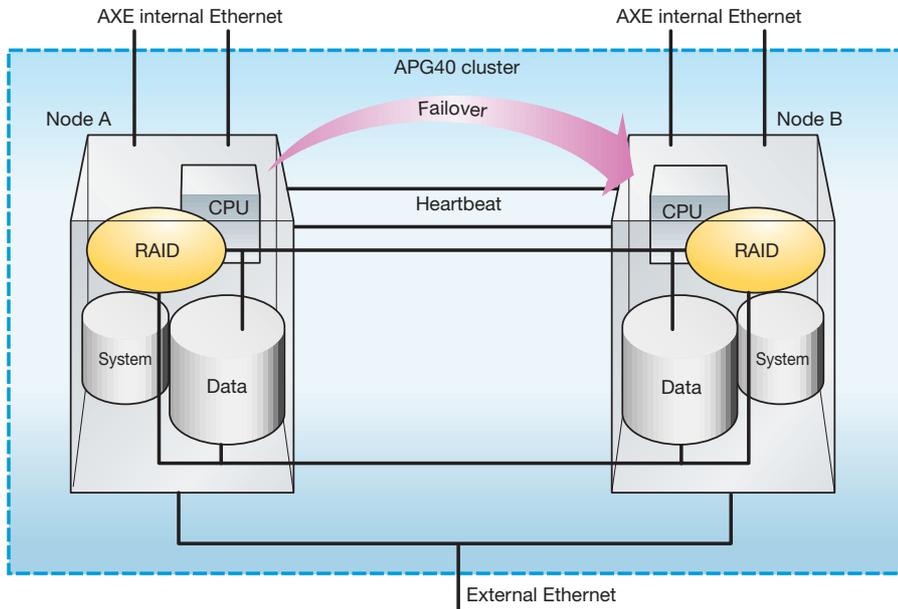


Figure 13
The RAID hardware system.

center, such as the operations support system (OSS), for preventive measures or for sending an alarm when a critical level is reached.

The applications that run on the APG40 do not communicate with one another and are not interdependent. The platform provides basic functionality, such as highly available processing capacity, safe storage, and data output mechanisms.

The adjunct computer subsystem (ACS) (proprietary middleware in Figure 12) provides the software platform for applications in the APG40. The ACS offers Windows NT, Microsoft cluster server (MSCS), and communication protocols, such as FTP, RPC, TCP/IP, and Telnet.

In-service performance

Operators require reliable input-output systems that store AXE system backup copy safely and perform reloads fast. A switch must always be available for operation and management from an operations support system. To ensure the in-service performance of the APG40 itself, and AXE as a whole, concepts introduced in the IOGs and APG30 have been brought to this new platform. In the APG40, mainstream technologies, such as MSCS and RAID hardware, form the basis for telecommunications-grade in-service performance.

The APG40 is a cluster server (Figure 13). In its first release, a cluster will consist of

two nodes. Each node has its own memory and runs the operating system. All applications running on an APG40 cluster are active on one node only—the executive node; the other node is a standby node. Thus, if a major fault occurs, the standby node starts and takes over the application. A watchdog in the MSCS ensures that the adjunct processor software does not simply stop working. When an error is detected, either in hardware or software, the cluster server can restart the failed resources in the other node. This process is called *failover*.

The physical size of the APG40 is a half magazine for a small cluster and one full magazine for a large cluster. The small APG40 has one 18 GB disk per node, while the large APG40 can house up to three 18 GB disks per node. Later on, it will be possible to introduce disks with a capacity of 36, 72, or 144 GB.

Thanks to the RAID hardware system (Figure 13), the APG40 has a higher disk-integration speed than the IOG20. The RAID hardware system is used for safe storage. Only the executive node can write data to the data disk. To avoid loss of data due to a disk error, the executive node writes data to a Windows NT file system, and the RAID system writes to the data disk in both the executive node and the standby node. The executive node thus owns all data disks in both nodes. In the IOG20, on the other hand, each node owns its own set of data

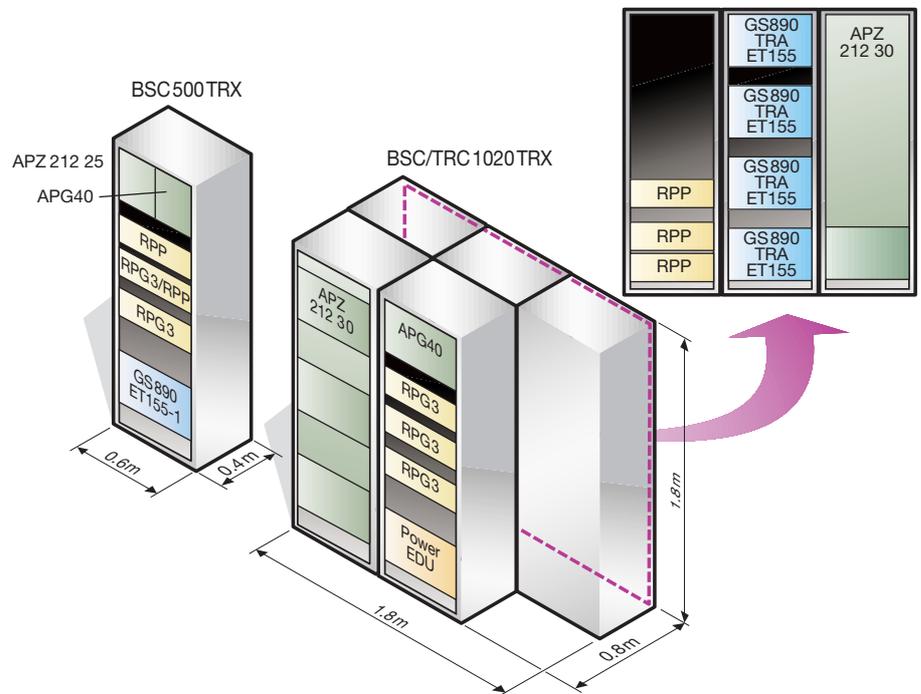


Figure 14
Node layout for BSC 500 TRX and
BSC/TRC 1020 TRX.

disks, which must be re-integrated each time a failover occurs. Re-integration in the APG40 occurs when a node is replaced or the RAID system fails.

The AXE IO platforms listed in Table 2 are based on three different high-availability concepts. In the IOG20, the passive node is *hot-standby*, whereas in the APG40, the passive node is *warm-standby*. In this context, the main difference between hot- and warm-standby is that a larger part of the software is up and running in hot standby. In the APG30—with its fault-tolerant concept—identical processes are executed in parallel on separate processor units. If a hardware failure occurs, an immediate failover is initiated from the failing processor unit.

Behind the different availability concepts in Table 2 lie different assumptions about the most probable source of failure. If a hardware failure happens more often than a software failure, the fault-tolerant technology gives better in-service performance. A failover in the APG30 executes in an instant, because the same processes, including the applications, are running in parallel on both sides.

Nevertheless, because the risk of hardware failure is considered extremely low, the case for dimensioning has been based on software failure, and the high-availability concept used in the APG40 has been introduced as the preferred solution. Given the risk of hardware faults, application software faults and system software faults, the failover period has been held to a minimum. On the APG40, these three kinds of fault are well below the minimum CP buffer capacity of two minutes, which means that no data is lost.

Performance and functionality enhancements

The demands put on the adjunct processors are not the same as those put on the central processor. The central processor provides non-stop computing, so the adjunct processor concept is based on and provides high-availability computing. Consequently, new processor systems can be obtained from mainstream data communications vendors. The first generation APG40 is based on Intel Pentium III 500 MHz processors and Microsoft NT4 clustering technology. By adapting these mainstream products to spe-

TRADEMARKS

MIPS is a registered trademark of MIPS Technologies, Incorporated.
NonStop-UX is a trademark of Compaq Information Technologies Group, L.P. in the United States and/or other countries.
Windows and Windows NT are registered trademarks of Microsoft Corporation.
Intel and Pentium are trademarks or registered trademarks of Intel Corporation.
Rational and ClearCase are trademarks or registered trademarks of Rational Software Corporation in the US and in other countries.

cific needs, Ericsson benefits from ongoing research and development conducted by others. This approach also simplifies the rapid introduction of new features and applications using Microsoft and Rational development tools integrated in an adjunct processor design environment. Examples of this are debugger and Rational test tools integrated in Microsoft Visual C++, Rational ClearCase, and HTML online help.

Other examples of mainstream technology being introduced full-scale in the AXE by the APG40 are the TCP/IP, FTP, and Ethernet (for external communication and internal AXE communication), and the cPCI building practice. By introducing TCP/IP and Ethernet in the IPN for CP-AP communication, the central processor reload speed is 10 times that obtained when using the APG40 with signaling terminal open communication (STOC).

The introduction of TCP/IP as the external interface to the OSS is a major step toward replacing the proprietary interfaces based on MTP and X.25. The external Ethernet board will enable a TCP/IP connection of up to 100 Mbit/s.

Telnet has been introduced for handling commands. Similarly, FTP and RPC are used for handling files. Support for a component object model-based (COM) OSS might be introduced in the future.

For local element management of the AXE 810, WinFIOL and Tools 6.0 (and later versions) have been adapted to the new APG40.

Results

Best-in-class hardware embodies compact layouts and easy extensions, exemplified in Figure 14 by one BSC and one combined BSC/TRC for the GSM system.

The BSC, which has a capacity of 500 transceivers, occupies only one cabinet. At the top, the CP (APZ 212 25) and APG40 share one shelf. Below this are three magazines with a mix of RPPs and RPG3s. The C7, transceiver handler and PCU applications run on this hardware. At the bottom of the cabinet, one GEM contains ET155, synchronization, DL3 interface boards (DLEB) and the group switch.

The combined BSC/TRC has a capacity of 1,020 transceivers and approximately 6,000 Erlang. The APZ 212 30 is needed for this larger node, occupying two cabinets. Another cabinet houses four GEMs, each of which contains a mix of transcoders,

TABLE 2, THREE AXE IO SYSTEMS

	HSB (IOG20)	FT (APG3x)	HA (APG4x)
Hardware fault	Software failover 90 sec.	Hardware failover < 1 sec.	Software failover < 60 sec.
Application software fault	Software failover 90 sec.	Process restart 22 sec.	Process restart < 1 sec.
System software fault (major)	Software failover 90 sec.	Reboot 80 sec.	Software failover < 60 sec.

ET155, synchronization, DL3 interface boards and the group switch. The next cabinet contains one APG40, the RPG3s (for transceiver handling) and C7. The power equipment is located at the bottom. PCUs can be included in the last cabinet if the BSC/TRC is intended to handle GPRS traffic.

Conclusion

The AXE architecture allows us to continue to improve the different parts of the system as required by network evolution. This ensures that operators always have a network with the best possible cost-of-ownership. A limited number of high-capacity nodes with a limited footprint keep site costs down. State-of-the-art hardware with reduced cabling also ensures quick, high-quality installation procedures. Operation and maintenance is kept economical through low power consumption, an interface consistent with the OSS, and a high-capacity, user-friendly extension-module platform. The software architecture allows migration toward next-generation networks, whether fixed or wireless.

REFERENCES

- 1 Ekeröth, L. and Hedström, P-M.: GPRS support nodes. Ericsson Review Vol. 77 (2000):3, pp. 156-169.
- 2 Hägg, U. and Lundqvist, T.: AXE hardware evolution. Ericsson Review, Vol. 74 (1997): 2, pp. 52-61.
- 3 Hallenstål, M., Thune, U. and Öster, G. ENGINE server network. Ericsson Review Vol. 77 (2000):3, pp. 126-135.
- 4 Andersson, J-O. and Linder, P. ENGINE Access Ramp—The next-generation access architecture. Ericsson Review Vol. 77 (2000):3, pp. 146-147.