The Movement from Monoliths to Component-Based Network Elements

Nils Björkman, Yong Jiang, and Torbjörn Lundberg, Telia Research, Sweden
Alexander Latour-Henner, Service Factory, Sweden
Avri Doria, Nortel Networks, Sweden

ABSTRACT

To be competitive in a rapidly growing market requires rapid upgrades to the performance and functionality of the network. One way to manage rapid upgrades of the network with minimum risk is to deploy equipment using a modular system architecture. Modularity allows a network operator to mix and match best of breed components to achieve the desired system rather than rely on vendors to implement specific technology before making crucial business decisions. This article begins with an overview of the current global movement toward standards that support network elements with modular system architecture. The story begins with university initiatives and the forming of OpenSig and IEEE P1520 more than five years ago, continuing with related and complementary initiatives by the Parlay Group, Softswitch Consortium, Multiservice Switching Forum, and several IETF working groups. Next, special attention is given to the component-based architecture of the Multiservice Switching Forum released in summer 2000. The trend of building network equipment from components with distinctly different functional specialties is described in three examples: media gateways, IP routers, and virtual IP routers. It is envisioned that component-based network infrastructure will spawn new markets for entrepreneurial developers, spurring competition and accelerating the creation of innovative solutions for all facets of global communications. The article concludes with a smorgasbord of new market opportunities.

INTRODUCTION

The evolution of network technology is accelerating. A telephony switch used to depreciate over 30 years or more, whereas asynchronous transfer mode (ATM) switches and IP routers today typically need to be replaced within three years. To be competitive in a rapidly growing market requires rapid upgrades of the performance and functionality of the network. Furthermore, a badly chosen network solution may turn out to be a costly cul-de-sac.

One way to manage rapid upgrades of the network at minimum risk is to deploy networking equipment with modular system architecture. Modularity allows a network operator to mix and match best of breed components and to replace the parts of the system that need either additional functionality or just higher capacity. The Multiservice Switching Forum (MSF) [1], Internet Engineering Task Force (IETF), Parlay Group [2], Softswitch Consortium [3], and IEEE P1520 [4] all contribute to standards that support modular system architectures.

FROM MONOLITHS TO COMPONENT-BASED SYSTEMS

MONOLITHS AND WHY THEY AREN'T A GOOD IDEA ANYMORE

Historically, vendors have provided operators with monolithic equipment built to provide one or several related functions. This was a reasonable approach from the vendor's perspective, since many of them were only in a single business area (i.e., router companies or telephone switch companies). Largely, this was fine with the service providers as well, since they were often in a single industry (i.e., Internet service providers, ISPs, or telephone companies). With the advent of service convergence it is no longer true that service providers provide a single kind of service. As time goes on, the data provider and telephony provider are becoming one, through either acqusition or diversification. Likewise, vendors are beginning to provide equipment for a diverse services. It is difficult, however, for vendors to provide equipment that provides all the services a customer might desire. The permutation of possibilities becomes too

0163-6804/01/$10.00 © 2001 IEEE

large and is very dynamic over time. Instead, the vendor will choose one or two services that fit nicely together and bundle them into its new units. Unfortunately, this often does not provide the service provider with the functionality or versatility required to satisfy their business models. The business model becoming prevalent is one that allows a multitude of services to be rolled out quickly and bundled on a single transmission infrastructure.

Unfortunately, the equipment currently installed in most networks typically has complete control of a set of transmission and other resources. This arrangement does not allow sharing a common physical infrastructure without the use of costly and complex protocol overlay technologies. In cases of monolithic equipment, the logic controlling path determination and traffic flows are bundled into a single physical unit with the switching/forwarding functionality.

A consequence of this bundling is that operators do not have flexibility to define new services over existing investment. In the monolithic approach the success of the operator is tightly coupled to the decisions and capabilities of a single equipment vendor. Since the business models and priorities of the vendors and service providers may diverge, this can be frustrating for the service provider. Once a network is in place the operator is seriously constrained with respect to:

- The vendor's internal architecture: If the design chosen by a vendor is not suitable for the addition of functionality without major revision, the service provider will have to wait a long time for equipment that will satisfy current service needs.
- Functionality of new software releases: The more functionality added to a system after its creation, the more complex and time-consuming the testing and release cycle.
- Functionality of new hardware releases: If a vendor needs to create new hardware variations to meet the functionality requirement, the cycle becomes quite long, and quite expensive for both the vendor and the service provider.
- Pricing policy: If a service provider's requirements are fairly specialized, the vendor is justified in charging specialized rates to produce what is required. This is not advantageous to either party.

Often, due to these and other factors, the vendor will not be able to meet the provider's requirements. In these cases, the operator is left with no choice other than to replace the entire network; this is an unfavorable result.

COMPONENT SYSTEMS AND WHY THEY ARE A GOOD IDEA

Several forms of component system architectures are described in this article. One thing is typical to all component architectures: critical interface points are open and defined in detail. These open definitions allow for diverse vendors to build function-specific equipment able to be matched with equipment from other vendors.

There are several advantages of such systems:

- By opening the interface points between planes, multiple services from one plane can share the resources of the plane below as peers.

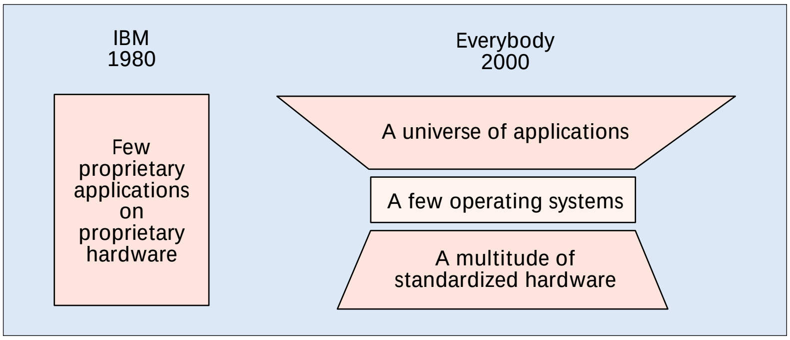

Figure 1. *The PC revolution.*

- Every vendor can concentrate on the part of the system at which they are best.
- Service providers are free to buy equipment from the vendor they think provides them the best capabilities per unit cost.
- When technology changes at one plane, only that plane needs to be upgraded or replaced.
- Service providers are freed from vendor constraints in developing services.

This quick analysis shows that this arrangement is advantageous to the service provider who wants freedom. But is a component architecture also good for the vendor? Many vendors seem to believe they are better off maintaining the service provider in a captive state. A comparison to the history of the PC industry is instructive and enlightening in this respect.

A COMPARISON WITH THE PC INDUSTRY

The decision by IBM to open the PC interface gave Microsoft the opportunity to offer OS support to other PC hardware vendors and established the principle of separating the control software from the computing hardware. Today monolithic mainframes have been replaced by the ubiquitous PC, a few operating system vendors, and a universe of new applications, as depicted in Fig. 1.

The open PC architecture allowed many vendors to participate and compete in the market. This new multivendor environment resulted in increased competition, rapid development of increased performance and new functions, reduction in prices, a larger market, and an explosion of new applications and software. And IBM is thriving in the PC market.

A similar effect can be expected in the communications industry from the introduction and widespread deployment of systems based on open component architectures.

INITIATIVES FOR OPEN INTERFACES

OPENSIG AND OPENARCH

The requirement for network equipment with open interfaces was expressed by network operators as far back as the telephony era. Developers of networking applications have also had a desire to explore new ways of using a network. In 1996 Columbia University [5] held the first OpenSig workshop to promote research on open network control issues. In 1998 the IEEE Communications Society sponsored an expanded OpenSig as part of the IEEE OpenArch conference.

IEEE P1520

Following the success of the OpenSig and OpenArch conferences, IEEE established P1520 as a standardization group for programmable networks [4]. The goal of this is to make the network as programmable as the PC through a set of standardized application programming interfaces (APIs), while maintaining extensibility and flexibility to accommodate future functionality and proprietary differentiation. Their work is aimed at both ATM and IP networks, and is based on the Reference Model which consists of the following interfaces:

- V interface: User-level APIs that provide access to the value-added services level
- U interface: APIs that provide access to network generic services, such as connection management and directory services
- L interface: APIs that provide access to virtual network devices, allowing manipulation of local device network resource states
- CCM interface: The connection control and management interface is a collection of open protocols to access the state of physical elements.

The P1520 had proof of concept implementations demonstrated in early 1999 and later at Telecom '99. Their current focus for standardization is on the L interface for both IP routers and ATM switches, and the CCM interface for ATM switches.

The goal is to make the network as programmable as the PC through a set of standardized APIs, while maintaining extensibility and flexibility to accommodate future functionality and proprietary differentiation.

THE INTERNATIONAL SOFTSWITCH CONSORTIUM

The International Softswitch Consortium (ISC) [3] was created in May 1999, and has approximately 150 members. Their objective is to promote open architectures, protocols, and APIs. This is expected to facilitate development of services and applications, and to lower barriers to entry for both system vendors and service providers.

The Softswitch architecture distinguishes between five planes: the data, application, control, transport, and management planes.

The Consortium focuses on interoperability and certification of voice and other real-time services. Five standard interfaces have been adopted by the ISC. Protocols for horizontal interfaces include H.323, Session Initiation Protocol (SIP), Real-Time Transport Protocol (RTP), and Real-Time Streaming Protocol (RTSP). Vertical protocols (i.e., protocols that allow decomposition into controlling and controlled components) include Media Gateway Control Protocol (MGCP) and MEGACO.

PARLAY

The purpose of the Parlay Group is to enable enterprises to control a range of network capabilities and access information within the network operator's domain. To achieve this, the Parlay Group has defined a number of APIs allowing direct access to communication facilities.

The Parlay Group foresees a future where smart devices and a multitude of applications will proliferate. These will need integration with communication services enabled by the APIs.

Currently the Parlay API contains service interfaces addressing call control, messaging, security, IP network control, mobility, performance management, and audit capabilities.

The next release, scheduled for 2001, is expected to add:

- Generic charging/billing
- Policy management
- Aspects of service and network management
- Mobile m-commerce/e-commerce
- Subscriber data/user profile/virtual home environment (VHE)

The Parlay Group was started in March 1998 by five companies: BT, Microsoft, Nortel Networks, Siemens, and Ulticom BT. It was broadened in May 1999 with AT&T, Cegetel, Cisco, Ericsson, IBM, and Lucent. This year it was reorganized and is now open to all interested parties.

IETF

IETF has taken many initiatives to allow external control. The activities are mainly in the following areas:

- Media Gateway Control (MEGACO): This protocol is being developed by the IETF in cooperation with the International Telecommunication Union — Telecommunication Standardization Sector (ITU-T). The working group has developed an informational RFC, detailing the architecture and requirements, and a standards track protocol for controlling media gateways from external media gateway controllers. A media gateway is a network element that provides conversion between the information carried on telephone circuits and data packets carried over the Internet or other IP networks [10].
- Common Open Policy Service (COPS): This protocol was developed as a resource allocation protocol. COPS extensions for policy provisioning and traffic engineering are being defined. COPS has received much interest from both vendors and carriers.
- General Switch Management Protocol (GSMPv3): GSMP provides an interface that allows a routing logic component to control a label switch [9].
- Simple Network Management Protocol (SNMP): Recent work has focused on incorporating security, policy, and other improvements, and outlining the most effective methods for using the framework to accomplish configuration management.

MULTISERVICE SWITCHING FORUM

The Multiservice Switching Forum was founded by Cisco, Bellcore, and MCI WorldCom in November 1998 and currently has 57 members. MSF has gathered strong industrial support from carriers, manufacturers of datacom and telecom equipment, and several hardware and software component vendors.

The MSF's mission is to accelerate the deployment of open communication systems using the flexible support of a full range of network services over multiple infrastructure technologies. The focus is on development of architectures and industry agreements.

One of the fundamental insights forming the basis of MSF involves recognizing the architectural similarities shared by IP routers with frame relay and ATM switches. The forwarding hardware becomes a common platform for switching data packets of any format. This allows a wide range of services to be supported on a common platform. This in turn enables carriers and service providers to specialize their offerings.

THE MSF ARCHITECTURE

FOUR CORNERSTONES

Mapping the principles established in the computer industry to the communications industry implies:

- Division of the monolithic switch into specialized components
- An open standardized interface between these components
- Logical partitioning of hardware to allow the execution of several controlling software instances in parallel
- Coordinated, hierarchical management of physical and virtual network elements

THE COMPONENT APPROACH

Figure 2 shows all the functions and reference points of the MSF architecture [5]. A small set of standardized protocols is required to implement all open interfaces across a subset of the reference points in the architecture.

Figure 2 shows the multiplanar system model chosen. Starting from the bottom of Fig. 2, the adaptation plane supports the physical interface to a user or another network element. The switching plane supports the actual switching fabric by which physical interfaces are connected. The control plane provides the generic capability to manage network service events, and provides control over both the adaptation and switching planes.

Standard protocols are used in communicating between the control plane and the currently undifferentiated switching and adaptation planes. An opening of an interface between the adaptation and switching planes is expected to lead to an additional increase in flexibility.

The application plane provides services that use the capabilities of the control plane while also providing enhancements to the services realized within the control plane.

The planes aggregate functions that interact to realize the generic behavioral model of an MSF-compliant system. These are defined below.

Logical Port Function — The LPF provides media mapping and service-specific multilayer adaptation functions for the incoming media stream.

Virtual Switch Function — The VSF is an arbitrary subset of switch resources that can be controlled as a unit. Switching resources are responsible for switching media streams from one logical port to another. The switching resources may provide packet switching, frame switching, cell switching, and so on.

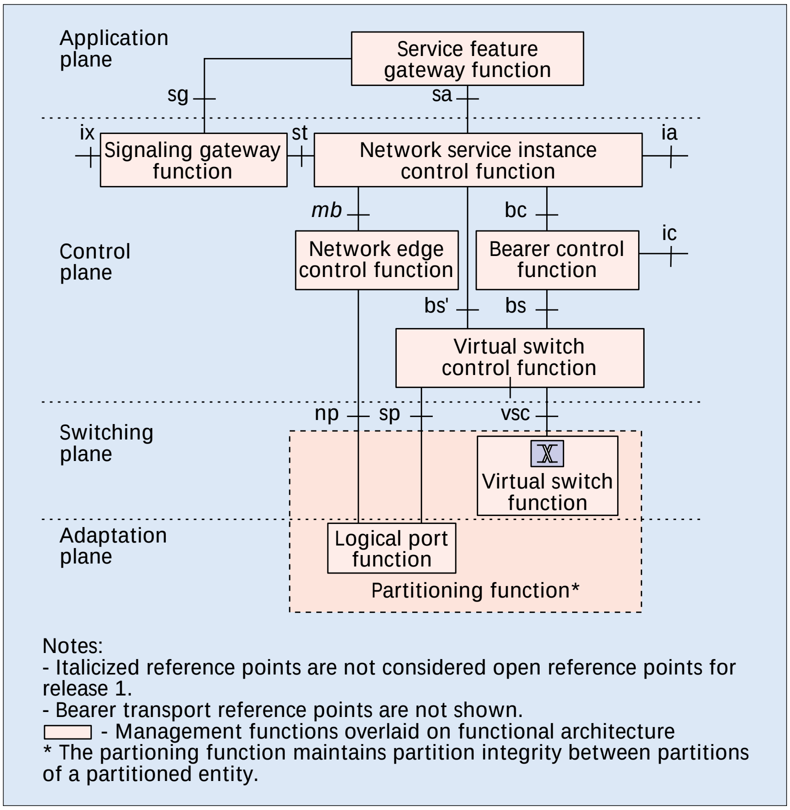

Notes:
- Italicized reference points are not considered open reference points for release 1.
- Bearer transport reference points are not shown.
▭ - Management functions overlaid on functional architecture
* The partioning function maintains partition integrity between partitions of a partitioned entity.

Figure 2. *The MSF reference architecture.*

Network Edge Control Function — The NECF is responsible for sending and receiving control information to and from an LPF. The scope of control encompasses all protocol layers covered by the adaptation in an LPF. Some examples are checksum verification and encryption.

Virtual Switch Control Function — The VSCF controls and monitors the VSFs and LPFs within a partition. The VSCF provides the required cross-connect information, including traffic and quality of service (QoS) information, across the VSF from one LPF to another.

Bearer Control Function — The BCF establishes, modifies, and releases end-to-end bearers between end point(s) of a bearer connection. Some examples are user-network interface/private network-network interface (UNI/PNNI) signaling for ATM connections or intermediate system-intermediate system/Open Shortest Path First (IS-IS/OSPF) for intradomain routing.

Network Service Instance Control Function — The NSICF establishes, maintains, modifies, and releases network service instances. Some examples of network service instances include circuit-switched calls and interdomain Border Gateway Protocol v. 4 (BGP-4) routes.

Signaling Gateway Function — The SGF processes signaling. It maps, relays, or tunnels the signaling between networks to produce an end-to-end bearer connection.

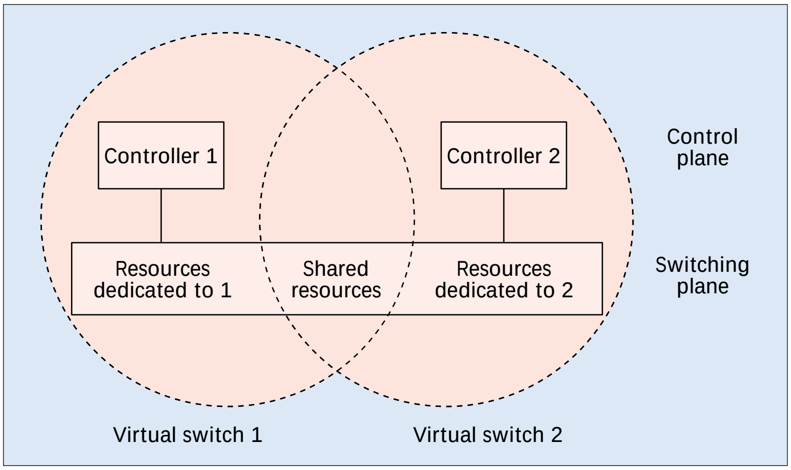

Figure 3. *Statistical sharing of switching and adaptation resources between two controllers.*

Service Feature Gateway Function — The SFGF allows access to intelligent network services and other network-provided applications. It also allows directly signaled services in the application plane to access the control plane functionality.

VERTICAL OPEN INTERFACES

The MSF architecture encourages new innovations and new standards. A key objective is to allow the technologies of forwarding and control to evolve independent from each other.

The MSF distinguishes between intrasystem and intersystem interfaces. An intrasystem interface defines the information exchange between the forwarding and control components within a network element. Intersystem interfaces handle the interaction across the network between components of the same plane. System integration with components from different vendors who specialize in different component types is made possible through the introduction of open intrasystem interfaces.

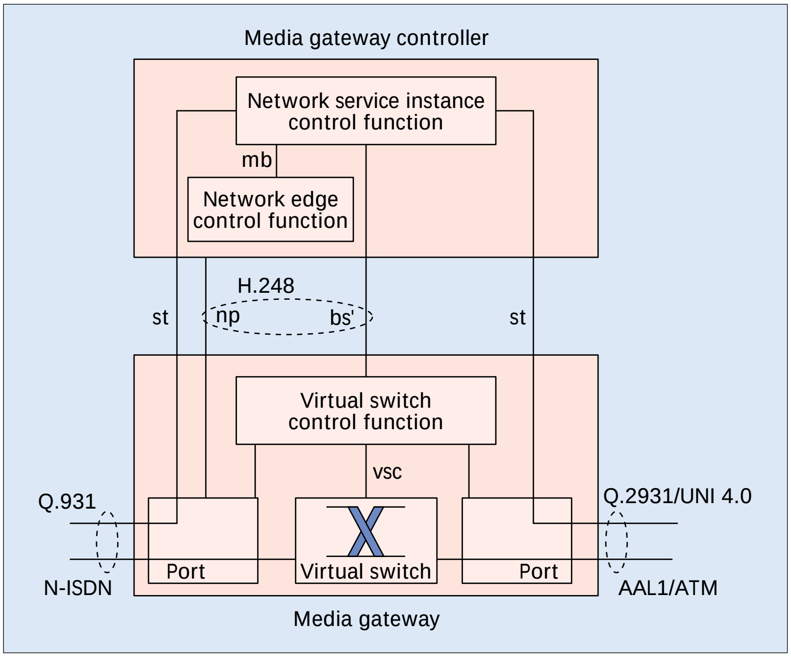

Figure 4. *A physical model of a decomposed media gateway.*

SWITCH PARTITIONING

The functions in the control plane can share the resources of the switching and adaptation planes through the partitioning function, the VSF.

Management creates a virtual switch by specifying the switch resources to make up the partition. These resources can include physical port resources such as bandwidth and buffer space, and physical switch resources such as forwarding table space.

It is possible to combine virtual switches in such a way that they interact with each other only in a predictable and controlled manner. The sharing of resources between virtual switches may be either deterministic or statistical.

The first application of partitioning is expected to provide multiple instances of services to separate groups of users. These user groups expect strict isolation from each other's activities. This implies that the principle of resource sharing *between partitions* should be *deterministic*. Note that the control logic of a partition still requires an ability to use statistical allocation principles for the resources *within* its partition.

More efficient utilization may be achieved if virtual switches statistically share pooled resources among several otherwise noncooperating control logic instances, as illustrated in Fig. 3. The concept of *statistical* resource sharing between partitions is more difficult to define stringently than the deterministic case. The challenge lies, as previously (ATM Forum service classes, IETF DiffServ/IntServ, etc.), in:

- Defining a control mechanism together with a behavioral conformance definition to maintain a utilization contract with each instance of partition control logic
- Defining a performance target that governs the allocation of partitioned resources so that additional requests for allocation are blocked when the probability of failing to maintain the performance is unacceptably high

The terms deterministic and statistical are not to be confused with static and dynamic configuration. The term *static configuration* means that the virtual switch partitions are established prior to the switching system as a whole being brought in service. A static configuration can be changed only after the switching system as a whole is brought out of service. Dynamic configuration allows for changes without bringing the switching system as a whole out of service.

THE MANAGEMENT HIERARCHY

The MSF management plane complements existing standards with functions, management information bases (MIBs), and interfaces required by the MSF architecture. These new requirements stem mainly from the fact that the MSF reference architecture:

- Allows components to be "mixed and matched"
- Defines the VSF, responsible for partitioning of switching resources into virtual switches

The first bullet leads to a requirement for management of each component separately. The management interface and management information for each component need to be standardized in the management hierarchy. As in PCs,

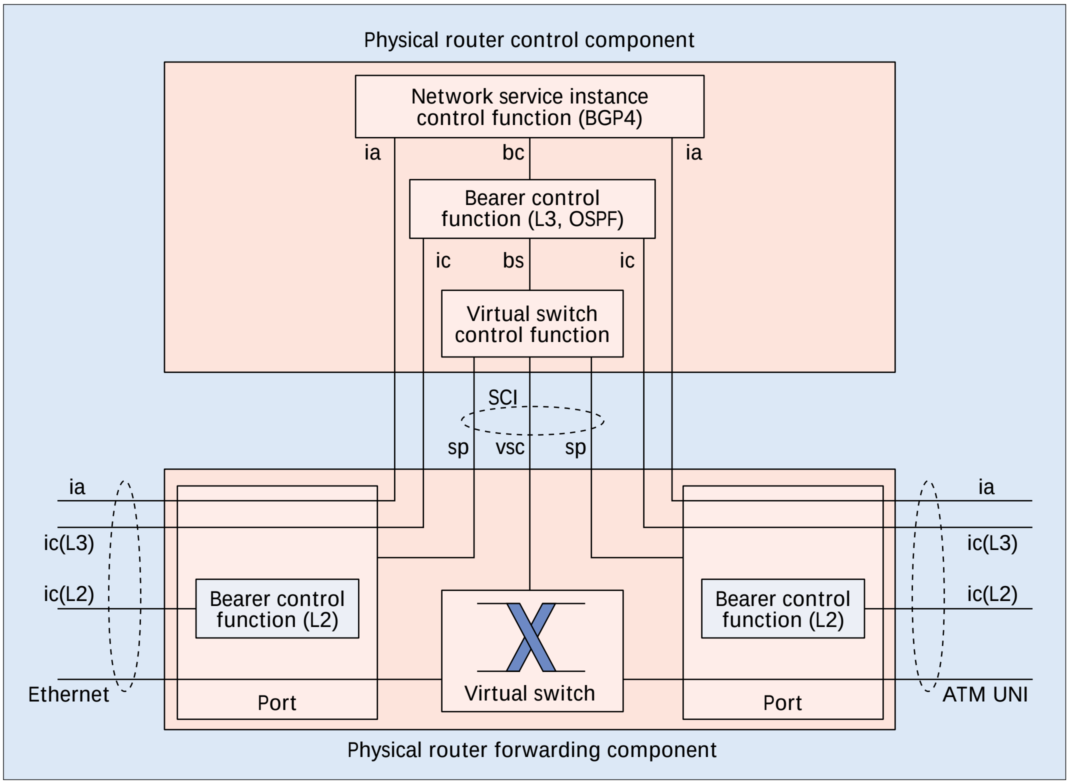

Figure 5. *A physical model of a classical IP router.*

The logical separation of control from forwarding will increase the speed and decrease the cost of introducing new services and features, a definite advantage over the traditional monolithic approach.

plug-and-play functionality needs to be provided for all components in the architecture.

The second bullet allows virtual switches to be created on some set of switches in the network to form a virtual network. Each virtual switch can be controlled by an independent controller as though it were a physical switch. Similarly, each virtual network can be managed by a separate network manager as though it were a physical network. In this environment, management functions are required on two levels:

- Superordinate management that manages the physical resources and allows the creation, modification, and deletion of virtual entities
- Subordinate management of the virtual entities

REALIZATION EXAMPLES

The trend of building network equipment from components with distinctly different functional specialties is described by three recent examples: media gateways, IP routers, and virtual IP routers.

TELEPHONY GATEWAY (MGC/MG)

This example, depicted in Fig. 4, decomposes two distinctly different functional domains into two separate components: a media gateway (MG) and a media gateway controller (MGC). This system provides conversion between the information carried on telephone circuits and the data packets carried over an IP network.

The MGC component is responsible for protocol interactions with the service logic leading to establishment of a call. It is important to note that for the telephony service, the operator profile is mainly established through specific implementation of the service logic exercising this protocol interaction.

The MG component is responsible for adapting the media streams as packaged at the network boundaries.

This decomposition will deliver freedom of choice to network operators for deployment of MGs and MGCs from different suppliers. The logical separation of control from forwarding will also increase the speed and decrease the cost of introducing new services and features; this is a definite advantage over the traditional monolithic approach.

CLASSICAL IP ROUTER

Figure 5 describes the application of the MSF architecture to best effort IP forwarding capability as defined in RFC 1812, "Requirements for IPv4 Routers." The purpose is to illustrate how a router is optimally decomposed into two separate "boxes."

An MSF-compatible box is composed of some or all functions described in the MSF architecture. The boxes are connected over an interface traversing one or more MSF-defined reference points. To emphasize that the examples represent a physical implementation, and not just a set of logical relationships, all traffic enters/leaves the boxes through physical ports.

The bundling applied in this example separates the IP routing control component from the IP forwarding component. This supports a natural specialization and the independent scaling of each component. The upper control part maximizes reachability while maintaining stability. With the volume of traffic in the Internet growing exponentially, the key feature of the forward-

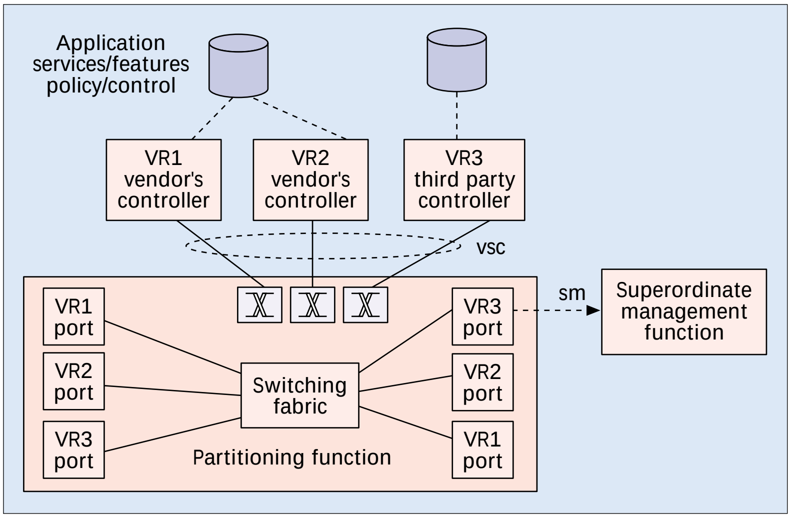

Figure 6. *A partitioned IP router.*

ing system is the ability to forward an ever-increasing volume of traffic.

The example illustrates a minimal split into two physical elements, thus requiring the smallest number of open interfaces. In the future it may be feasible/required to introduce finer modularity in the physical model.

VIRTUAL ROUTERS

The concept of partitioning enables the operation of multiple controllers over a single forwarding engine. Figure 6 illustrates this concept for an IP router partitioned into multiple virtual routers. This would involve creating a model of a virtual router that could be instantiated as a data model in each of the partitions. The data model lists router resources such as routing tables, bandwidth, buffer space, labels, and CPU control resources.

Issues to be addressed include dynamic control of resource allocation, and deterministic and statistical partitioning of resources. This can be done through open interface control protocols, such as GSMP or COPS.

NEW MARKET OPPORTUNITIES

Efficient operation of a network requires economies of scale. An operator's dilemma is that standardized solutions make product differentiation more difficult. The MSF has adopted the concept of virtual switches, combining the economies of scale for transport and switching while allowing differentiation of service components. The network control logic and switching/forwarding functionality are allowed to evolve independent of each other. This leads to flexibility for operators who no longer have to make all-or-nothing decisions when building networks. Instead, the networks can be allowed to evolve flexibly with vendors and products being substituted without the need to rebuild the whole network.

The potential for market growth is in fact tremendous for both vendors and operators. The MSF open architecture with open intrasystem interfaces will bring about:

- Best-of-breed components from multiple vendors. Open interfaces enable greater participation in the development of future products, enabling a larger variety of services and network operator profiles.
- Widen the geographical market for operators without physical presence by offering remote control of a logical partition. Thus, a network operator can offer a leased virtual switch service.
- Lower the threshold for introduction of new types of network control while allowing gracious phaseout of network control types with declining demand. No more forklift replacement!
- Shorten the time from the original idea to market for nonstandardized services. The concept of virtual switches allows prestandard products to be introduced on a small scale while waiting for a standard to settle.

VISION

This article describes a model of global connectivity in which network elements are interconnected and intermixed with relative freedom. This leads to a next-generation network infrastructure that will spawn new markets for entrepreneurial developers of network elements. This will spur competition, and in turn will accelerate the introduction of innovative solutions for problems and opportunities not yet even identified.

Both large and small providers of telecommunications services and equipment are aligning behind this vision because it will enable rapid deployment of new and enhanced services, without necessitating new investments in switching and transmission resources. Such an infrastructure will support quick deployment of innovative, even experimental, services, controlled by software that can be adapted to meet the new and evolving requirements of end users.

This new view is exemplified and motivated by the MSF vision, shared by a broad range of service providers and equipment vendors. This vision will continue to gain support as the industry realizes the benefits of an integrated network and a service-oriented architecture that includes interoperable and interchangeable elements. The MSF, along with other industry fora and standards bodies, is dedicated to fostering global acceptance of this vision, one that will lead to a rapid proliferation of network services and lasting improvements in global communication.

ACKNOWLEDGMENT

The authors are especially indebted to Dave McDysan, Technical Chair of the MSF, and to the members of the MSF Technical Committee for their groundbreaking contributions. We thank Satoshi Yoshisawa for updates on OpenSig and P1520.

REFERENCES

[1] The Multiservice Switching Forum: http://www.msforum.org
[2] The Parlay Group, "Introduction and Objectives," White paper: http://www.parlay.org
[3] Int'l. Softswitch Consortium: http://www.softswitch.org
[4] IEEE P1520, "Proposed IEEE Standard for Application Programming Interfaces for Networks": http://www.ieee-pin.org

[5] A. Campell *et al.*, "Open Signaling for ATM, Internet and Mobile Networks," *OPENSIG '98*, Oct. 1998: http://comet.columbia.edu/opensig/activities

[6] N. Bjorkman *et al.*, "MSF System Architecture Implementation Agreement," Apr. 2000: http://www.msforum.org

[7] S. Rooney *et al.*, "The Tempest a Framework for Safe, Resource Assured, Programmable Networks," *IEEE Trans. Commun.*.

[8] P. Newman *et al.*, "Ipsilon's General Switch Management Protocol Specification v. 2.0," IETF RFC 2297, Mar. 1998.

[9] IETF GSMPv3: http://www.ietf.org/html.charters/gsmp-charter.html

[10] H.248/MEGACO: http://www.ietf.org/html.charters/megaco-charter.html

BIOGRAPHIES

NILS BJORKMAN (Nils.P.Bjorkman@telia.se) received his M.Sc. degree in engineering physics from the Royal Institute of Technology, Stockholm, Sweden, in 1985. The same year he joined Swedish Telecom and later Telia Research. Original research on fiber optics and IP and ATM networks has been demonstrated in dozens of courses, lectures, and publications. He was editor of the MSF Architecture group in 1999-2000. Currently he develops applications and markets the Essential Bandwidth estimator for the Irish start-up Measure Technology Ireland Ltd.

YONG JIANG received his Bachelor's degree from Beijing University of Science and Technology, China, in 1997. After working for one year as an IT consultant, he moved to Sweden and pursued his Master's degree at Linköping University in 1999. Currently he is a researcher at Telia Research. His research area is policy-based networking, switch partitioning in an open network, and future network access nodes.

TORBJÖRN LUNDBERG received an M.Sc. 1970 at the University of Stockholm. He has worked since 1992 for Telia Research in different areas, such as human factors, network architectures, and IP switching.

ALEXANDER LATOUR-HENNER holds an M.S.E.E. degree from the Royal Institute of Technology, Stockholm, Sweden. He has, over the last 25 years, worked in different areas from electrochemical power sources to tele- and data communications. While at Telia AB, he managed development of advanced service control architectures for future Internet services. Currently, he devotes his time to product development and marketing with two startup companies.

AVRI DORIA is chief architect in the Routing Architecture Laboratory for Nortel Networks. She serves on the Board of Directors of the MSF and is active in the IETF as co-chair of the GSMP working group. Her primary technical interests include component network architectures and improvements to the Internet routing architecture.

> *This vision will continue to gain support as the industry realizes the benefits of an integrated network and of a service-oriented architecture that includes interoperable and interchangeable elements.*

INTELLIGENCE IN OPTICAL NETWORKS

Intelligent Networks were aimed to add intelligence in telecom networks so that the end user could get a service delivered to him without knowing how it has been delivered and what it takes to deliver that service. Service logic was built in the central office, which was the core element responsible to provide the service, by reserving resources for each connection/call.

A central office (CO) switch has always been the core component in Telecom Networks enabling the communication between two end points. It reserves resources per connection/call and delivers the service to the end user, based on the service logic built in it. A CO switch has evolved from its basic form to today's form, loaded with lots of intelligence. This evolution has improved service delivery and insulated the user from knowing how the service has been delivered and what it takes to deliver that service.

The other important components of a telecom network, i.e., access (local loops) and transmission circuits, have so far played the role of circuit termination and the information carriers. With enormous growth in transport technology in the optical domain and the increasing information carrying capacity of the optical media, different approaches have been proposed to realize this potential to the end user in terms of services.

Many services have been explored as to how they can be delivered intelligently by the optical transport (i.e., how different optical mechanisms can apply and provide bearer services to the user).

With full control on wavelengths, there is opportunity to add intelligence in DWDM-based optical networks, and emerging optical systems (OXCs and OADM) can have knowledge of:

- The wavelengths in the network
- Traffic carrying capacity of each wavelength
- Their status

Such intelligence could create self-connecting and self-regulating networks as envisioned for next-generation transport networks (i.e., optical networks).

The Feature Topic "Network Intelligence in Optical Networks" scheduled to appear in the September 2001 issue of *IEEE Communications Magazine* is well timed considering the planning and R&D efforts going on in the optical market. It would provide a collection of papers aiming at different aspects of building intelligence in optical networks. A few of those identified are:

- Service-based next-generation networks
- Optical services considered for being delivered in these networks
- Optical service creation, delivery using network intelligence
- Network management and control in current and next-generation networks
- Alternatives to network intelligence like IN concepts, softswitch concepts, active networks, and programmable networks
- Intelligent access networks (intelligent local loops)

The acceptance of papers is subject to reviews by the editorial board and its members and other experts whom they identify.

Schedule (Planned for Sep. 2001):

Call for Submissions: 15th Nov. 2000
Manuscripts Due: 15th Jan. 2001
Acceptance Notification: 1st April 2001
Final Revised Manuscripts due: 1st June 2001
Manuscripts to Publisher: 1st July 2001

Submission Guidelines:
Please visit the following Web page for submission guidelines:
http://www.comsoc.org/~ci/sub_guidelines.html