

The SoftRouter Architecture

T.V. Lakshman T. Nandagopal R. Ramjee K. Sabnani T. Woo
Bell Laboratories, Lucent Technologies
{lakshman, thyaga, ramjee, kks, woo}@lucent.com

ABSTRACT

In current routers, implementations of the control and forwarding functions are colocated and tightly integrated by static association of control and forwarding elements. In this paper, we present the SoftRouter architecture that permits dynamic association between control and forwarding elements, and separates the implementation of control plane functions from packet forwarding functions. In this architecture, all control plane functions are implemented on general purpose servers called the control elements (CE's) that could be multiple hops away from the forwarding elements (FEs). A network element (NE) or a router is formed using dynamic binding between the CEs and the FEs. We argue that this flexibility results in several benefits including increased reliability, increased scalability, increased security, ease of adding new functionality, and decreased cost.

1. INTRODUCTION

In current routers, implementations of the control and forwarding functions are intertwined deeply in many ways. The control processors implementing control plane functions are colocated with the line cards that implement forwarding functions and often share the same router backplane. Communication between the control processors and the forwarding line cards is not based on any standards-based mechanism, making it impossible to interchange control processors and forwarding elements from different suppliers. This also leads to a static binding between forwarding elements and line cards. A router typically has at most two controllers (live and stand-by) running control plane software. Each line card is statically bound to these two controllers. The two controllers, the line-cards to which they are statically bound, and the switch fabric together constitute the router.

In this paper, we propose a new control plane architecture called the SoftRouter architecture that separates the implementation of control plane functions from packet forwarding functions. In this architecture, all control plane functions are implemented on general purpose servers called the control elements (CEs) that could be multiple hops away from the line cards or forwarding elements (FEs). Thus, there is no need for a static association between the CEs and the FEs. Each FE, when it boots up, discovers a set of CEs that can control it. The FE dynamically binds itself to a "best" CE from the discovered set of CEs. In this architecture, a collection of FEs (along with their switch fabrics), together with their associated CEs, is called a Network Element (NE) and logically constitutes a router. We envisage a standardized interface between the CEs and the FEs similar to that being standardized in the IETF ForCES working group [15]. The SoftRouter architecture and the technical challenges introduced by this ar-

chitecture are described in more detail in Sections 2 and 3, respectively.

Such an architecture has several benefits. By implementing the CE's on servers, the architecture permits easier scaling since the server capacity can be increased far more easily than increasing controller card capacities in routers. A server-based CE also facilitates use of stronger security mechanisms since such mechanisms have to be deployed at far fewer points in the network. A third aspect is the increased reliability possible both from a server-based implementation and from the architecture that permits each line card to have more than two possible controllers. A recent paper [4] argues that a server-based logically centralized implementation of BGP results in several benefits. By moving all control functionality out of the forwarding element, several other benefits, as described in Section 4, are now possible.

The proposed network evolution has similarities to the SoftSwitch based transformation of the voice network architecture that is currently taking place. The SoftSwitch architecture [14] was introduced to separate the voice transport path from the call control software. The SoftRouter architecture is aimed at providing an analogous migration in routed packet networks by separating the forwarding elements from the control elements. Similar to the SoftSwitch, the SoftRouter architecture reduces the complexity of adding new functionality into the network. We discuss other related work in Section 5 before presenting our conclusions in Section 6.

We now present an overview of the SoftRouter architecture.

2. ARCHITECTURE OVERVIEW

The SoftRouter architecture comprises of a number of different network entities and protocols between them. We describe them in separate subsections below.

2.1 Network Entities

A SoftRouter network can be described in two different views, namely, the *physical* view and the *routing* view.

In the physical view, a SoftRouter network is made up of *nodes* interconnected by *links*. There are two types of nodes: the *forwarding elements* (FEs) and the *control elements* (CEs). An FE is similar in construction to a traditional router; it may have multiple line cards (each in turn terminating multiple ports - physical or logical) and a backplane (switch fabric) that shuttles data traffic from one line card to another. Its key difference from a traditional router is the absence of sophisticated control logic (e.g., a routing process like OSPF) running locally. Instead the control logic is hosted at CEs, which are essentially general purpose server machines. A link connects any two elements (FEs/CEs). Typically, a FE has multiple in-




Figure 1: NE Possibilities

cident links (so that data traffic can be routed from one link to another) and an CE is multi-homed to more than one FE (so that it is not disconnected from the network should its only link fail). In a nutshell, the physical view of a SoftRouter network is not that different from that of a traditional routed network, except with the addition of a few multi-homed servers (CEs).

The primary function of a FE is to “switch” data traffic between its links. The exact nature of the switching function can take different forms. We describe three possibilities here, among others: (1) Packet forwarding: this includes both layer 2 (MAC-based switching) and layer 3 (longest prefix match) forwarding. (2) Label switching: an example of this is MPLS forwarding. The data path forwarding function can include label swapping, pushing and popping. (3) Optical switching: the traffic in this case can be time-switched, wavelength-switched, or space-switched among the links. In each of these cases, the switching function is driven by a simple local table which is “computed” and “installed” by some CE on the network. In general, a FE can do more than switching. For example, a FE can perform security functions such as packet filtering and intrusion detection. The key requirement is that these functions should work off a local data structure whose management intelligence resides in some remote CE. For example, in the packet filtering case, the filtering logic consults only a local filter table, whose management (insertion and deletion of filtering rules) is performed in some remote CE.

The routing view of a network captures the topology of a network as seen by the routing control logic. To describe this view, we first need to define the concept of a *network element* (NE). At a high level, an NE is a logical grouping of network ports¹ and the respective CEs that manage those ports. A network administrator may place different restrictions on what network ports can be grouped together in the SoftRouter network, resulting in a set of allowable NEs. Specifically, there are two dimensions where restrictions can be placed (see Figure 1). The vertical dimension restricts the selection of FEs: *Single FE* means that the network ports in a NE must be from a single FE; *Connected FEs* means the ports can be selected from a set of FEs that are connected (i.e., there exists a physical path from one FE to another); and *Any FEs* means the ports can come from any set of FEs. The horizontal dimension specifies if an NE must include all ports

¹For obvious reasons, a port cannot belong to more than one NE.




Figure 2: Example SoftRouter Network: All ports in Connected FEs

of the FEs under consideration or if it can include only a subset. Clearly, any combination of the two dimensions provides a possible definition of the ports of an NE. We say that two NEs are *logically connected* if they contain ports belonging to FEs that are physically connected.

To illustrate these definitions better, we examine two specific cases, which represent the two extremes of the spectrum enabled by this definition.

Subset of ports in any FEs: any ports from any FEs can be grouped as part of a single NE. The routing view can be significantly different from the physical view. This definition can be useful for VPN applications. Essentially, the different ports in an NE represent the different sites in a wide-area VPN.

All ports in a single FE: each FE is part of one NE. The physical and the routing views become identical. With this definition, the only difference between a traditional routed network and a SoftRouter network is the offloading of the routing logic onto remote CEs.

While the SoftRouter architecture allows all possible configurations, we believe one of the most interesting and practical case is the *All ports in Connected FEs* case. This represents the clustering of neighboring FEs into a single NE and corresponds to the typical case of several routers being connected back-to-back in a central office. From a routing perspective, this can provide significant simplification: (1) the reduced number of NEs in the routing view reduces the inter-NE routing complexity; and (2) a different (possibly less complex) routing protocol can be employed for intra-NE routing (more details are discussed in Section 3). For the balance of this paper, we focus primarily on this case.

In the SoftRouter model, the routing control of the NEs is disaggregated from the FEs and the control protocol runs on the CEs. A *binding* between a FE and a CE means that the CE is performing particular control functions on behalf of the FE. Because multiple protocols (e.g., IGP and EGP, or even multiple instances of a protocol) may be required for the operation of a FE, a FE may have multiple CE bindings.

An example SoftRouter network illustrating the different concepts we have introduced is shown in Figure 2. Note that a binding can exist between a FE and a CE even if they are not directly connected (e.g., CE1 and FE1 in Figure 2). It is worthy to point out that the SoftRouter architecture thus enables the creation of a separate (log-

ically or physically) signaling network connecting all the CEs and this can make the control plane much more resilient to attacks.

2.2 Protocols

A number of different protocols are needed in the operation of a SoftRouter network. We describe the three most important ones here.

Discovery Protocol: In order for a FE to establish a binding with a CE, it must first know about the existence of the CE and be able to reach it using some route. A discovery protocol finds out what CEs are available and lays out paths to them for the FEs.

FE/CE Control Protocol: Once a binding is established, the FEs and the CEs communicate using a control protocol. On the uplink (FE to CE) direction, this control protocol provides link state information (e.g. link up/down signal) to the CE. On the downlink direction, the protocol carries configuration and control information (e.g., enable/disable a link, forwarding information base (FIB), etc.). The key issues in the design of this protocol relate to the frequency, bandwidth and delay requirements for the FE/CE communication. In this paper, we assume that the ForCES [15] protocol would be sufficient for our purposes.

CE/CE Protocol: In a pure SoftRouter network, especially one with a separate signaling network, a CE to CE protocol is necessary for the CEs to discover each other and determine the routes between them. In this case, the CEs would also exchange control packets (e.g., OSPF database exchange packets) directly between themselves. On the other hand, in a heterogeneous network, a CE to CE protocol is not strictly necessary. In this case, interoperability with existing routers can be preserved by using the FE/CE protocol to tunnel any control packets received (sent) from (to) any of the ports in the FE (e.g., all OSPF protocol packets if CE is running OSPF for the FE).

3. TECHNICAL CHALLENGES

In any given NE, the FEs are strongly connected (with the same underlying topology), and the CE can be many hops away from the corresponding forwarding set. This separation leads to many new scenarios and technical challenges that do not normally occur in existing networks.

In this section, we highlight some of these issues and present potential means of addressing them. We will use the term *integrated NE* to denote a router in an existing network, where the control element and the corresponding forwarding elements are part of a single physical device.

3.1 Bootstrapping

In an integrated NE, the configuration pertaining to that NE is obtained upon bootup, since the control card on the box is in direct contact with the forwarding engine through a *bus/interconnect*. As soon as the router (NE) comes up, the forwarding engine knows how to obtain the configuration information immediately from the control processor.

In the SoftRouter model, upon bootup, the forwarding engine has to obtain its configuration information, which includes the IP addresses of its interfaces, from a remote

control element that resides on a server. This poses a potential paradox: in order to discover a CE and send packets to it, the FE requires routing information; however, the routing information is supposed to come from the CE. This paradox can be resolved using a *discovery protocol* that lets FEs and CEs discover each other and maintains a routable path between these entities.

If all the FEs in the SoftRouter network support ethernet services, a spanning tree protocol can route the packets of the discovery protocol. If the FEs are heterogeneous (see Section 2.1) and do not support spanning tree, a hop-by-hop or source-routed routing layer over IP would suffice.

CEs and FEs would then advertise their existence over this routing layer. FEs can then dynamically bind to the “best” CE using either bindings pre-configured by the network administrator or obtained using distributed clustering algorithms. In the latter case, it may be possible to devise zero-configuration self-managing networks.

Note that the routing provided by the discovery protocol, connecting the FEs and the CEs, would be used only for routing control messages between the CE and the FE and thus, optimal routing is not essential in this case. On the other hand, the routing of data traffic would be optimal, governed by forwarding tables on the FEs that are computed and installed by the CEs.

Further, the discovery protocol can be enhanced by additional features such as CE load balancing and failover (discussed later), constrained clustering algorithms that increase CE-FE binding resiliency etc., increasing the availability of the network.

3.2 Routing and Forwarding

Since the FEs are assumed to be simple with minimal control on-board, the CE is responsible for maintaining the knowledge of the links’ status between FEs within the same NE. In addition, the CE must also be able to integrate topology changes within the NE with external (inter-NE) route changes, and update the FIBs of individual FEs accordingly. This can be done by using a protocol similar to ForCES [15].

Due to the difference between the physical view and the routing view of the network as a result of this architecture, IP TTL and IP options behavior might deviate from traditional router design, which assumes geographical closeness of control and forwarding plane. In the SoftRouter architecture, the FEs of a given NE might be distributed over a large geographical area, which makes it expedient to decrement TTL and process IP options on a per FE (rather than a per NE) basis. For an NE with co-located FEs, one can similarly realize this behavior or imitate traditional router behavior.

The FE must also be able to support a minimal set of *slow path* functions such as support for ICMP. Further, many of the IP options such as Timestamp, Strict/Loose Source Route, etc. can also be handled by the FE. Other IP options such as Router Alert [11] will be tunneled to the CE for processing.

3.3 Protocol Optimization

In standard internet routing protocols, there are various messages that fulfil different functions. For example, OSPF has *hello* packets to probe link status and various *LSA* messages to advertise link status to the rest of the net-

work. In a SoftRouter architecture, these messages have to be differentiated according to their purpose. A FE can choose to either forward all received routing protocol packets to its CE, or it could handle a subset of the protocol packets by itself without forwarding it to the CE, thereby reducing response time and control traffic in the network.

Let us consider the OSPF routing protocol as an example. One of the main purposes of the *HELLO* message is to check that links are operational. We could move this functionality to the FEs which would then notify their respective CEs only when there is a change in the status of the links². On the other hand, LSA messages serve to advertise connectivity information to the rest of the network, and hence these will originate from the CE, and will be exchanged between CEs.

Within the SoftRouter architecture, routing protocols can be optimized in terms of message overhead. In a traditional router network running OSPF, for example, LSA messages from each FE are flooded over the entire network. In a SoftRouter network, LSAs are flooded only over the network of CEs. Given the potential difference in the magnitude of CEs and FEs, this reduces the number of OSPF messages sent over the network. We are investigating further optimizations enabled by the SoftRouter architecture.

3.4 Failures and Load Balancing

By separating the control elements, we provide more choices for CEs to control a given FE. This choice comes at the risk of vulnerability to loss of CE connectivity due to both intermediate link and node failures, in addition to failure of the CE itself. This failure event is compounded by the fact that, in a SoftRouter network, a single CE may be responsible for managing a large number of FEs, that are now unmanaged.

However, note that, CE connectivity does not result in a critical failure immediately, since the FEs can continue to forward packets using its existing FIB. The discovery protocol will discover an alternate route to the CE when intermediate links and/or nodes fail, thereby recovering from the loss of control.

We address CE node failures by using a standard server concept, that of backup CEs. Each FE will have a *primary CE*, and a *secondary CE set*. When the discovery protocol or the control protocol signals a loss of association with the primary CE, the FE can choose a candidate from the secondary CE set to become its primary CE. The re-association is done using the discovery protocol.

Apart from CE failures, a CE can voluntarily relinquish control of a FE for load balancing purposes. In that scenario, the CE can ask the FE to find another CE to manage it, or the CE can guide the FE to a new CE. In the former, the new CE will have to restart the routing instance for the FE, while in the latter, the old CE can transfer the FE's state to the new CE, resulting in a hot fail-over. The ability to load balance and recover quickly from control failures is a salient feature of this architecture.

4. BENEFITS

There are five significant benefits to the proposed architecture:

²We might still retain other functionalities of the *HELLO* message such as designated router election on the CE.

1) Increased reliability: The reduced software in a forwarding element makes that element more robust. As regards the control element, protocol specific and independent mechanisms can be incorporated to enhance reliability;

2) Increased scalability: control elements can be implemented on general-purpose servers, and thus can be easily scaled up using well-established server scaling techniques;

3) Increased control plane security: fewer management points makes it easier to manage and provide a strong defense around the control elements, thus making the overall network more secure;

4) Ease of adding new functionality: adding new functionality is easier on a separate control server, executing on general purpose processors and operating systems;

5) Lower costs: SoftRouter decouples the innovation curve of the control and forwarding elements, allowing economies of scale to reduce cost.

We now discuss each of these benefits in more detail.

4.1 Reliability

The separation of control and forwarding elements in the SoftRouter architecture provides several reliability benefits. First, the reduced software in a forwarding element implies that it is easier to make that element robust. On the control plane server side, sophisticated reliability enhancing mechanisms such as automatic fail-over and the use of hot or cold standby's, as discussed in Section 3, can be incorporated. Further, a server can have much higher redundancy capabilities such as 1:N failover capability rather than a very limited 1:1 control blade failover capability of today's routers. Finally, the ability to host different protocols on different CEs imply that the failure of one CE does not render the corresponding FE useless (e.g. the failure of the CE hosting BGP would still allow the FE to process OSPF protocol messages). Apart from these generic benefits that result in higher availability, we detail protocol-specific optimizations with respect to two-key protocols, BGP and OSPF, below.

4.1.1 Increased BGP reliability

Figure 3(a) shows a typical deployment of BGP with Route Reflectors [2] in a large Autonomous System (AS) today. This deployment has two main drawbacks. Specifically, under certain conditions, the network can go into persistent route oscillation where a subset of routers may exchange routing information without ever reaching a stable routing state [1]. Another issue with the Route Reflector architecture is IBGP reliability. While the failure of one IBGP session will affect only two routers in the case of a full mesh IBGP architecture, the same failure of a session between two Route Reflectors could partition the network, resulting in significantly lower reliability.

In the case of BGP deployment in the SoftRouter architecture (Figure 3(b)), the number of control elements that run BGP will typically be at least an order of magnitude smaller than the number of routers. Thus, a full IBGP mesh can easily be maintained among the control elements. The control elements would then download the appropriate forwarding tables to all the forwarding elements using the ForCES protocol [15]. *Thus, the persistent route oscillation problem is trivially solved in the SoftRouter architecture (since there are no Route Reflectors), thereby*

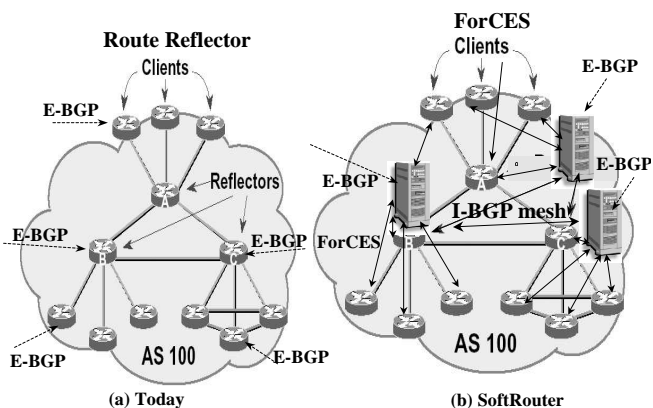


Figure 3: BGP deployment

increasing the availability of the network.

Further, in the SoftRouter architecture, the IBGP mesh is between servers that can employ a higher degree of redundancy such as 1:N ($N > 1$) as compared to 1:1 redundancy in the control processors of the routers. Thus, BGP reliability in the SoftRouter architecture can be significantly improved over the current BGP deployments³.

Note that the implementation of BGP in the SoftRouter architecture is logically similar to the router control platform (RCP) [4] proposal. The key differences are that a) unlike RCP, the IGP protocol is also executing in the server, b) we use the ForCES protocol to communicate with the forwarding elements rather than using IGP and IBGP as is the case in RCP (drastically simplifying the software on the forwarding element), and c) we use an IBGP mesh between the control element servers inside a single AS.

4.1.2 Faster OSPF convergence

OSPF convergence in the presence of failures is known to take tens of seconds in large networks today [6]. This delay can have a significant impact on the availability of the network, especially for critical services such as voice-over-IP. Compared to the 5 9's availability of the telephone network (translates to 5.25 minutes of downtime a year), today's data network can ill afford even a few link failures before its availability falls below the 99.999% availability target. The SoftRouter architecture allows for faster OSPF convergence through several complementary techniques:

- (a) Order of magnitude fewer control elements for the same number of forwarding elements, resulting in a smaller OSPF control network and faster convergence
- (b) Faster processors (see section 4.2) result in faster computations (shortest-path calculations)

In the extreme case, there can be a single CE (with appropriate backup CEs) in the SoftRouter network, executing the all-pairs shortest path algorithm and computing the forwarding tables of all the FEs in the network in a centralized manner. Thus, OSPF convergence time can be minimized to about 50ms or less, resulting in a highly available data network.

4.2 Scalability

Some of the fundamental limitations in scaling routing protocols in existing architectures are control processor ca-

³This assumes that the CEs are multi-homed and that there are multiple CEs in the network for fail-over.

capacity and on-board memory. The decoupling provided by the SoftRouter architecture makes it easier to upgrade control hardware which is based on general-purpose servers; contrast this to the difficulty in obtaining an upgraded control processor card from a router manufacturer who needs to accommodate upgrades with other constraints such as power availability, slot availability, etc.

We now highlight this advantage through a specific example of requirements of a highly scalable Mobile IP home agent [9]. One common approach today for deploying highly scalable home agents is to use routers that support several hundred thousand home agents. However, signaling scalability is limited due to limitations of the control processor to about hundred bindings/sec (or less than two updates per hour per user). This is a significant limitation as updates generated through both mobility as well as the periodic refresh mechanism can easily exceed two per hour per user. Another approach used today is to implement Mobile IP on a cluster of general purpose processors. Signaling scalability will not be an issue here. However, scaling the number of home agents becomes difficult since IPsec processing (for each agent) is CPU intensive and will not scale efficiently to several hundred thousand home agents without specialized hardware.

SoftRouter architecture admits a complementary combination of both of these approaches, allowing server-based signaling scalability while retaining hardware-based transport scalability. Thus, transport will still be handled by FEs with hardware support for IPsec using regular router blades while signaling capacity can be easily scaled up using multiple server blades.

4.3 Security

The SoftRouter architecture enables the adoption of a multi-fence approach to security with each fence adding a layer of security. These include:

- (a) Off-the-shelf versus special-purpose operating system: specialized operating systems are not as widely-tested for security holes;
- (b) Multi-blade server platform versus one or two control blades in the router: overload due to malicious traffic can be distributed over a large number of processors and sophisticated compute-intensive intrusive detection mechanisms can be deployed;
- (c) Fewer control elements: managing fewer elements is easier (e.g. changing security keys frequently) and it may be possible to place these few elements in a more secure environment (physically or logically firewalled) compared to the numerous forwarding elements;
- (d) Separate signaling network: using a physically or logically separate signaling network for the Internet can limit attacks on control plane protocols.

4.4 New Functionality

The separation of control and forwarding elements and the use of general purpose servers to host the control processes enable easier introduction of new network-based functionalities such as traffic engineering, network-based virtual private network (VPN) support, active network services [13] etc. *Since each FE can have multiple CE bindings in the SoftRouter network (see Section 2.1), each of these services can be hosted on separate CEs, thereby isolating the effects of new applications on critical functions*

such as routing. We now illustrate the benefits of the SoftRouter architecture using VPNs as an example.

There has been significant recent activity in defining network-based VPN services using BGP/MPLS [10]. In this application, a VPN server dynamically creates MPLS or IPSEC tunnels among the provider edge routers. While the VPN server would execute on the router control board in today's architecture, migrating the VPN server functionality into a control element in the SoftRouter architecture has several benefits: 1) VPN server upgrades can now be independently performed without impacting basic network operations such as forwarding; 2) Network-wide failover of VPN control servers can be performed without impacting existing or new VPN sessions; 3) Configuring BGP policies for the provider edge routers connected to the VPN customer sites can be done in a central location at the VPN server rather than at multiple routers (e.g. edge routers and route reflectors involved in the VPN); 4) MPLS tunnels can be engineered in a centralized manner to meet customer requirements; 5) Scalability for support of large number of VPNs can be easily handled using generic server scaling techniques.

4.5 Cost

In the SoftRouter model, the forwarding element is mostly hardware-based and requires little management. This allows economies of scale to help reduce the cost of each forwarding element. The individual router control element is consolidated into a few dedicated control plane servers. These control elements will run on generic computing servers rather than in specialized control processor cards in routers as is the case today. This allows "complete sharing" of control processor resources resulting in better efficiency than the "complete partitioning" approach adopted today. Further, since the control plane servers are just another application for generic computing servers, the SoftRouter architecture can leverage the CPU price-performance curve of these server platforms. Finally, the reduced number of control plane elements means fewer boxes to manage, thus reducing operating expenses.

5. RELATED WORK

The Open Signaling approach [8] advocated the separation of control plane from the forwarding plane in ATM networks for increased network programmability. This separation enabled the Tempest framework [12] to create architectures where multiple control planes could simultaneously control a single network of ATM switches.

The Internet Engineering Task Force (IETF) is working on standardizing a protocol between the control element and the forwarding element in the ForCES [15] working group. However, unlike the SoftRouter architecture, the focus is on an architecture where the control element is directly connected to the forwarding element.

The case for separating some of the routing protocols (specifically, BGP) multiple hops away from the routers have been made by several researchers [4, 5]. While it is possible to migrate a few selected protocols out of the forwarding element, such an approach does not deliver the full benefits of the SoftRouter architecture.

One of the key benefits of the SoftRouter architecture is that it makes it easier to add new functionality into the network. Researchers have proposed other techniques such

as Active Routers [13] or open programmable routers [7] to increase flexibility in deploying new protocols in the Internet. By separating the forwarding and control elements and hosting the control protocols on general purpose servers, a lot more resources are available for adding new software services in the SoftRouter architecture.

Finally, as the tussle between the service providers and the end users occurs in the deployment of various functions in the Internet [3], an architecture like the SoftRouter may be necessary for allowing flexible and resource-intensive solutions to be deployable.

6. CONCLUSIONS

In this paper we presented the SoftRouter architecture that permits establishing dynamic association between control and forwarding elements, allows flexible definition of network elements, and separates the implementation of control plane functions from packet forwarding functions. The forwarding elements were simple hardware devices with little intelligence and were controlled by control elements that might be multiple hops away. We then highlighted both the flexibility accommodated by the partitioned view as well as technical challenges unique to this architecture. Further, we argued that the SoftRouter architecture provides significant benefits including increased reliability, increased scalability, increased security, ease of adding new functionality and decreased cost. As data networks become integral to everyday life and as voice-over-IP services become increasingly deployed on data networks, these issues will become critical necessities. The SoftRouter architecture is well placed to meet these critical requirements.

7. REFERENCES

- [1] A. Basu et. al., "Route Oscillations in I-BGP with Route Reflection," In Proc. of ACM SIGCOMM, Aug. 2002.
- [2] T. Bates, R. Chandra and E. Chen, "BGP Route Reflection - An Alternative to Full Mesh I-BGP," RFC 2796, April 2000.
- [3] D. Clark, J. Wroclawski, K. Sollins, and R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet," In Proc. ACM SIGCOMM Conference, August 2002.
- [4] N. Feamster, H. Balakrishnan, J. Rexford, A. Shaikh, and J. van der Merwe, "The Case for Separating Routing from Routers," In Proc. of FDNA workshop, August 2004.
- [5] R. Govindan, C. Alaettinoglu, K. Varadhan and D. Estrin, "Route servers for inter-domain routing," *Computer Networks and ISDN systems*, Vol. 30, 1998, pp 1157-1174.
- [6] M. Goyal, K. Ramakrishnan and W. Feng, "Achieving Faster Failure Detection in OSPF Networks," In ICC'03.
- [7] M. Handley, O. Hudson, and E. Kohler, "XORP: An open platform for network research," In *HotNets*, 2002.
- [8] A. Lazar, "Programming Telecommunication Networks," *IEEE Network*, vol 11., pp. 8-18, Sept/Oct. 1997.
- [9] C. Perkins, "IP Mobility Support for IPv4," RFC3344, August 2002.
- [10] E. Rosen and Y. Rekhter, "BGP/MPLS VPNs," RFC2547, March 1999.
- [11] D. Katz, "IP Router Alert Option," RFC2113, Feb. 1997.
- [12] Sean Rooney, Jacobus E. van der Merwe, Simon Crosby and Ian Leslie "The Tempest, a Framework for Safe, Resource Assured, Programmable Networks," *IEEE Communications*, Vol 36, Number 10, October 1998, pp.42-53.
- [13] D. Wetherall, U. Legedza and J. Guttag, "Introducing New Internet Services: Why and How," *IEEE Network Magazine*, July/August 1998.
- [14] S. Williams, "The softswitch advantage," *IEE Review* Volume: 48 , Issue: 4 , July 2002, Pages:25 - 29
- [15] L. Yang et al., "Forwarding and Control Element Separation (ForCES) Framework," RFC 3746, 2004.