

How to Lease the Internet in Your Spare Time*

Nick Feamster
Georgia Tech
feamster@cc.gatech.edu

Lixin Gao
University of Massachusetts
lgao@ecs.umass.edu

Jennifer Rexford
Princeton University
jrex@cs.princeton.edu

ABSTRACT

Today’s Internet Service Providers (ISPs) serve two roles: managing their network infrastructure and providing (arguably limited) services to end users. We argue that coupling these roles impedes the deployment of new protocols and architectures, and that the future Internet should support two separate entities: infrastructure providers (who manage the physical infrastructure) and service providers (who deploy network protocols and offer end-to-end services). We present a high-level design for Cabo, an architecture that enables this separation; we also describe challenges associated with realizing this architecture.

Categories and Subject Descriptors

C.2.1 [Computer Communication Networks]: Network Architecture and Design

General Terms

Design, Management

1. Introduction

The Internet is relatively resistant to fundamental change. The last fifteen years have offered countless “false starts” in the deployment of new services. For example, differentiated services, IP multicast, and secure routing protocols have not seen wide-scale deployment, despite offering tangible value and making significant headway through the protocol standardization process. A major impediment to deploying these services is the need for *coordination*: an Internet service provider (ISP) that deploys the service garners little benefit until other domains follow suit. For example, an ISP that deploys a secure routing protocol like S-BGP incurs substantial cost but still is not protected from bogus route announcements unless *other* ISPs also deploy S-BGP.

ISPs are under immense pressure to offer “value added” services, in response to both customer demands and the increasing commoditization of Internet connectivity. Building a network with global reach requires either “building it yourself” or relying on other ISPs for connectivity; ISPs naturally adopt the latter approach to contain cost. Unfortunately, because a single ISP rarely controls the entire path, new services either have been deployed only in small islands

or have languished entirely. Some ISPs, hard-pressed to offer profitable services, are driven to extortionary measures such as degrading service for some, while providing “better service” (though not better *end-to-end* service) for others, as evidenced by the ongoing “net neutrality” debate.

Researchers are also under pressure to justify their work in the context of a federated network by explaining how new protocols could be deployed one network at a time, but emphasizing incremental deployability does not necessarily lead to the best architecture. In fact, focusing on incremental deployment may lead to solutions where each step along the path makes sense, but the end state is wrong. Rather, we argue that substantive improvements to the Internet architecture may require fundamental change that is *not* incrementally deployable. Unfortunately, ideas that are not incrementally deployable are typically relegated to the library of paper designs that are either never seen again, or, in rare cases, dusted off as “band aid” fixes only when crisis is imminent (as with IPv6 in the face of address depletion in IPv4).

We argue that decoupling *infrastructure providers* (who deploy and maintain network equipment) from *service providers* (who deploy network protocols and offer end-to-end services)¹ is the key to breaking this stalemate. We propose Cabo (“Concurrent Architectures are Better than One”), which exploits virtualization to allow a service provider to simultaneously run multiple end-to-end services over equipment owned by different infrastructure providers. Cabo extends network virtualization beyond its current use for supporting shared experimental facilities, such as PlanetLab [2] and GENI [6]. Rather than simply serving as an evaluation platform for selecting a single “winning” architecture, *support for virtual networks itself should be the architecture*. Cabo’s design adopts the *pluralist* philosophy [1] that advocates a flexible and extensible system that supports multiple simultaneous network architectures. We take the pluralist approach further by refactoring the economics of the Internet to have service providers form business relationships with one or more infrastructure providers to construct virtual networks that offer end-to-end services.

Decoupling service providers from infrastructure providers is consistent with the business models that have resulted from the commercialization of the Internet, and some ISPs are already pushing the trend toward decoupling service from infrastructure. The early days of the commercial Internet saw the rise of “carrier hotels”, which reduce

*Apologies to Staniford *et al.* [11]. **Note:** This article is an editorial note submitted to CCR. It has *not* been peer reviewed. Authors take full responsibility for this article’s technical content. Comments can be posted through CCR Online.

¹Throughout the paper, we use the term “service provider” as an organization that composes network services and protocols on top of physical infrastructure, and “Internet service provider” to refer to a status quo ISP.




Figure 1: Cabo architecture.

the cost of interconnection between ISPs by locating the physical equipment of many different ISPs in the same building. Co-location amortizes the high fixed cost of maintaining a physical footprint (*e.g.*, racks, power supplies, backup generators, switches, fiber, “hands and eyes” support, etc.) by sharing capital and operational expenditure across ISPs. In the same way that connectivity providers share infrastructure like backup generators, Cabo allows service providers to share physical network infrastructure. Now, several companies are starting to separate service from infrastructure to provide better service or lower costs. FON, a Spanish ISP, acts as third-party broker for existing wireless access points deployed by private households [4]. FON brokers Internet access using physical infrastructure deployed by other parties. Packet Fabric amortizes the cost of inter-exchange connectivity in a single city by allowing ISPs at these exchanges to share access to the same switch [5]. Cabo pushes this philosophy further, by allowing service providers to offer a wide range of *end-to-end* services and network architectures, not just Internet access or single switch ports.

Whether a “horizontally stratified” Internet would ever arise hinges on many important economic questions.² We focus, however, on the many *technical* hurdles that must be surmounted to realize Cabo. (As we describe in Section 2.2, many of Cabo’s benefits can be realized even when a single ISP deploys Cabo for the purposes of running multiple architectures in parallel on the same network infrastructure.) These technical challenges, which we describe in Section 2.3 include discovering the available physical infrastructure, mapping virtual networks into the underlying network topology, and accounting for and provisioning resources.

2. Concurrent Architectures Better Than One

In this section, we describe a high-level overview of Cabo, along with its associated benefits and challenges.

2.1 Cabo Architecture

Cabo separates the notion of conventional ISPs into two distinct entities: infrastructure providers and service providers. An *infrastructure provider* owns and maintains

²Separating infrastructure providers from service providers has precedence in other industries. For example, the airline industry has airports (infrastructure providers) that allocate certain gates or terminals to particular airlines (service providers). The leasing of commercial planes is another example. In the U.S., airlines (service providers) rent planes from financial-services firms (infrastructure providers). Similarly, shopping malls are typically owned by different companies (infrastructure providers) than the commercial chains that lease space for their stores (service providers).

the network equipment (*e.g.*, routers and links) that forms an *infrastructure network*. A *service provider* establishes agreements with one or more infrastructure providers for access to a share of these router and link resources. Cabo facilitates sharing of physical resources by subdividing a physical node (*i.e.*, router) or link into many virtual nodes and virtual links. A *virtual node* controls a subset of the underlying node resources, with guarantees of isolation from other virtual nodes running on the same machine. Similarly, a *virtual link* is formed from a path through the infrastructure network and includes a portion of the resources along the path. Cabo can guarantee bandwidth or delay properties on these links using schedulers that arbitrate access to shared resources, such as CPU, memory, and bandwidth.

A *virtual network* consists of virtual nodes and links that belong to the same service provider. For example, in Figure 1, service provider 1 has a virtual network using physical resources belonging to infrastructure providers 1 and 3 to provide end-to-end services between end hosts A and B. An end host may run virtual machines that connect to different virtual networks, possibly run by different service providers, over a physical connection to one infrastructure provider. Service providers may install software (*e.g.*, a customized routing protocol) on their virtual components and may even program the hardware (*e.g.*, a customized packet-forwarding algorithm implemented on a network processor). A single service provider may have multiple virtual networks tailored to specific services or topologies. For example, one virtual network may run an Interior Gateway Protocol (IGP) like OSPF and conventional longest-prefix match packet forwarding, while another virtual network may support source routing based on flat addresses.

One service provider may offer service to another via “nesting” of virtual components. That is, a virtual node might even be subdivided into multiple virtual nodes, and a virtual link itself comprises multiple virtual links. For example, one service provider might provide end-to-end connectivity (akin to an ISP today) and sell that connectivity to another service provider that offers some other end-to-end service. Also, an infrastructure provider might offer some services beyond the basic support for virtual components. For example, to reduce the number of nodes that other service providers would need to manage, an infrastructure provider might run a virtual network of its own, with virtual links between pairs of its edge routers.

2.2 The Benefits of Cabo

In this subsection, we present examples that illustrate the benefits of Cabo, including easy deployment end-to-end network services, the ability to run custom routing protocols, and better accountability.

2.2.1 Better network services

End-to-end network services. Some players in the “net neutrality” debate have advocated a *tiered Internet*, where Internet service providers provide “better” service to edge networks and content providers (*e.g.*, Google) who pay more money directly to those ISPs. This “enhanced service” is disingenuous: a tiered Internet cannot inherently provide

better service, since no single ISP controls any given end-to-end path (e.g., between a home user and Google). In Cabo, service providers can add real value by exposing control over *end-to-end* paths. In Cabo, infrastructure providers can achieve a competitive advantage by running more efficient and robust networks, and service providers differentiate themselves by running different end-to-end services on a common physical infrastructure.

Customized protocols. Cabo allows service providers to build virtual networks with dramatically different characteristics on top of the same physical infrastructure. For example, one service provider might deploy a virtual network that provides strong security guarantees (e.g., by using a secure routing protocol and self-certified addresses) at the expense of complete reachability, while another offers global reachability and greater anonymity (e.g., by using a conventional routing protocol with ephemeral IP addresses that depend on a host's current location). Similarly, one service provider might perform conventional IP routing and forwarding, while another permits end hosts to perform source routing on a relatively small virtual network, consisting of virtual links that span multiple hops in the infrastructure. Deploying source routing today is immensely difficult, since most ISPs disable the feature; in Cabo, a service provider could decide to offer source routing on its virtual network without having to coordinate with other ISPs.

Co-location for expanded network presence. In today's Internet, an organization that needs a global footprint must deploy physical infrastructure in a wide variety of locations; each router deployed in a new remote facility incurs a relatively high fixed cost. Today, these organizations can contract with an ISP that offers a Virtual Private Network (VPN) service, though finding a single ISP with facilities at every location may be difficult. In contrast, Cabo allows that enterprise (or its service provider) to instantiate virtual nodes and links on equipment managed by an infrastructure provider in diverse regions, allowing the organization to run its own virtual network without incurring the costs of deploying and managing its own physical equipment.

2.2.2 *More robust management and operations*

Testing and deploying new protocols. Today's router software is typically evaluated in a test lab before deployment. Large lab configurations that mimic a production network are expensive, and limiting tests to simple topologies and traffic patterns may not give operators an accurate view of how the new software would perform "in the wild". In Cabo, new router software (including new experimental services) could be evaluated on a separate virtual network on the same underlying infrastructure; this virtual network could initially carry only test traffic or support users willing to serve as early adopters. Also, migrating a network from one protocol to another can be painstaking. In Cabo, a new protocol could be deployed in its own virtual network, followed later by a cut-over of the data traffic from the old virtual network to the new one.

Protection against misconfiguration. Cabo provides isolation between different network components and services,

which can provide protection against misconfigurations and bugs. Network protocols are commonly misconfigured and are subject to implementation bugs. Adding a new service, provisioning a new customer, or rebalancing traffic each requires an operator either to invoke certain configuration commands or to install new software; these actions may cause instability or temporary service disruptions. Cabo allows services that might interact to be compartmentalized into different virtual networks, thereby preventing configuration errors or software bugs related to one network service from interfering with others.

Accountability at every layer. In the current Internet, a single ISP manages its network from the physical infrastructure, all the way up to applications, but that ISP typically does not have purview over an entire end-to-end path. When performance or security problems arise, the ISP must initiate the arduous process of locating the source of the fault (often a different ISP) and coordinating to diagnose and fix the problem. This process is inherently difficult because both monitoring and mitigation require coordination across one or more administrative boundaries. Cabo, on the other hand, allows each entity to have complete, end-to-end control over the layer it is managing. For example, when the virtual components do not behave as expected, the service provider has direct recourse (and a direct business relationship) with the infrastructure provider managing the equipment.

2.3 Challenges in Building Cabo

Realizing Cabo introduces many challenges that we are exploring in our ongoing work. First, although Cabo allows virtual networks to run customized protocols, we must demonstrate that the underlying equipment can provide this flexibility at high enough speed. In recent years, the major router vendors have started supporting virtual routers to simplify network designs, reduce capital expenditure, and lower the barriers to co-location [9]. To better support new protocols and forwarding algorithms, Cabo could also make use of programmable routers [7, 13], which can be simultaneously used by multiple parties. Conventional techniques for packet and CPU scheduling can ensure appropriate isolation between virtual networks that run on the same infrastructure.

We must explore whether managing multiple specialized virtual networks is less complex than managing one large general purpose network. In particular, Cabo must support provisioning and embedding, so that a service provider that specifies a virtual topology, traffic demand matrix, availability requirements, or some combination of these criteria can be allocated an appropriate virtual network. Cabo also introduces many management subtleties when physical components fail. First, the Cabo substrate must notify virtual network components when a failure or other disruption occurs. Second, it must account for performance on each physical and virtual component to enable a service provider to identify the cause of a performance or reliability problem in the virtual network in the underlying physical infrastructure. This approach to accountability is simpler than today's situation, where an ISP many hops away may disrupt end-to-end performance, without having any direct accountability to the senders or receivers of traffic.

In Cabo, a service provider coordinates with one or more infrastructure providers to create virtual networks, which should be easier than coordinating across ISPs to deploy new protocols and services today. Such coordination might be achieved with a signaling protocol that allows service providers to request virtual components; infrastructure providers could apply admission control and embedding algorithms to satisfy these requests. Requesting new virtual networks requires service providers to communicate with the infrastructure providers. To resolve this circularity, we believe that Cabo would ultimately have some virtual networks that provide global reachability, much like today's Internet, which provides complete connectivity even though the network does not intrinsically provide any such guarantee.

3. Related Work

Today's Internet offers several examples of refactoring connectivity to create new services. Equinix and Internap allow edge networks to change upstream providers on relatively short timescales, but these services can only control the *first* ISP along the path to the destination; in contrast, Cabo provides control over the entire end-to-end path. Other systems allow hosts to request overlay paths with certain properties [8], without providing programmability in the forwarding nodes. Content distribution networks and bandwidth brokers also extend basic connectivity by creating paths from source to destination (or content). Cabo provides this functionality by making the construction of virtual links a first-order primitive.

Cabo must allow many virtual networks to operate on the same physical infrastructure. Cabo is similar to "switchlets", which allow constructing virtual networks according to some set of specified properties using subdivisions of physical ATM switches and a standardized switch control interface [14]. This architecture relied on the design of a common interface to the switches but did not allow custom routing software or forwarding engines run directly on the switches themselves. Today's layer-3 VPNs [10] also provide virtualization functionality that is similar to Cabo. However, these VPNs do not (in and of themselves) provide resource isolation, they do not span multiple ISPs, and they offer service providers neither access to the physical routers nor the ability to run custom routing software on these routers.

Some research infrastructures use virtualization to support multiple simultaneous experiments. PlanetLab [2] supports virtualization of network servers, but not complete networks. GENI [6] and VINI [3] focus on network virtualization and programmability, with the goal of supporting experiments rather than acting as an architecture itself. In addition, Cabo must bootstrap its own communications to network elements, rather than relying on the legacy Internet.

In supporting programmable routers, Cabo resembles active networks. Previous research on active networking focused on issues with mobile code (and the resulting language and security issues) and providing control to end users [12]. In contrast, Cabo provides service providers (rather than users) with their own virtual networks, with a fairly general programming environment on the virtual nodes. In fact, a

service provider could run an active-network architecture in a virtual network on Cabo.

4. Conclusion

Although we are focusing on solving the technical problems with Cabo, we recognize that Cabo requires fundamental changes to the Internet's operation that might hinder deployment. For example, infrastructure providers would have to provide service providers access to their infrastructure. We note that Cabo does not prevent a single entity from acting as both an infrastructure provider and a service provider, and Cabo could provide benefits even to single ISPs. Limited cooperation is already compatible with incentives today, such as establishing geographical footprints by leasing a virtual router in other ISPs.

In searching for a single "right" future Internet architecture, Cabo suggests that the right future network architecture is not an end state comprised of a collection of addressing, routing, and forwarding paradigms, but rather a platform that allows these functions to evolve as demands on communication networks change. Indeed, the designers of IP aimed for generality, recognizing that they could not predict what networked applications would ultimately run on top of the network. Continual rapid advances in communication technologies and the sheer difficulty of predicting future requirements suggest that the architecture itself should be sufficiently general to enable support for network protocols, services, and architectures that we cannot imagine today.

REFERENCES

- [1] T. Anderson, L. Peterson, S. Shenker, and J. Turner. Overcoming the Internet impasse through virtualization. *IEEE Computer*, 38(4):34–41, Apr. 2005.
- [2] A. Bavier, M. Bowman, D. Culler, B. Chun, S. Karlin, S. Muir, L. Peterson, T. Roscoe, T. Spalink, and M. Wawrzoniak. Operating System Support for Planetary-Scale Network Services. In *Proc. Networked Systems Design and Implementation*, Mar. 2004.
- [3] A. Bavier, N. Feamster, M. Huang, L. Peterson, and J. Rexford. In VINI Veritas: Realistic and controlled network experimentation. In *Proc. ACM SIGCOMM*, Sept. 2006.
- [4] FON: WiFi everywhere! <http://en.fon.com/>, 2006.
- [5] Private communication with Avi Freedman, Oct. 2006.
- [6] GENI: Global Environment for Network Innovations. <http://www.geni.net/>.
- [7] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek. The Click modular router. *ACM Transactions on Computer Systems*, 18(3):263–297, Aug. 2000.
- [8] K. Lakshminarayanan, I. Stoica, and S. Shenker. Routing as a Service. Technical Report UCB-CS-04-1327, UC Berkeley, 2004.
- [9] D. McPherson et al. Core Network Design and Vendor Prophecies. In *NANOG 25*, June 2003.
- [10] E. Rosen and Y. Rekhter. *BGP/MPLS VPNs*. Internet Engineering Task Force, Mar. 1999. RFC 2547.
- [11] S. Staniford, V. Paxson, and N. Weaver. How to Own the Internet in your spare time. In *Proc. 11th USENIX Security Symposium*, Aug. 2002.
- [12] D. L. Tennenhouse and D. J. Wetherall. Towards an active network architecture. *ACM Computer Communications Review*, Apr. 1996.
- [13] J. Turner. A proposed architecture for the GENI backbone platform. In *Proc. Architectures for Networking and Communications Systems*, Dec. 2006.
- [14] J. van der Merwe, S. Rooney, I. Leslie, and S. Crosby. The Tempest - A Practical Framework for Network Programmability. *IEEE Network*, 12(3):20–28, May 1998.