Overlay Networks and the Future of the Internet

Dave CLARK, Bill LEHR, Steve BAUER, Peyman FARATIN, Rahul SAMI
Massachusetts Institute of Technology

John WROCLAWSKI
University of Southern California, Information Sciences Institute

Abstract: In recent years, we have seen the emergence of numerous types of so-called "overlay" networks in the internet. There are many diverse examples of such overlay networks including the content-delivery-caching networks, implemented by companies like Akamai, the peer-to-peer file sharing networks associated with applications such as BitTorrent, the voice-over-IP services offered via Skype, and various testbed networks such as PlanetLab. Such overlays have important technological and policy implications for the evolution of next generation internet architecture. This paper provides a first attempt to understand the implications of such overlays for internet architecture, industry structure, and policy. We introduce a taxonomy for thinking about these overlays with some examples of their scale and growing importance in the internet, and suggest some preliminary thoughts on the implications of these overlays for industry structure and policy.

Key words: internet architecture, overlay networks, ISPs, internet policy.

The internet started out as a government-funded research network running on top of the Public Switched Telecommunications Network (PSTN). The internet was a data application, mostly unregulated, that was supported on top of the public-utility regulated telephone networks. The internet was an "overlay" that complemented the underlying basic infrastructure of the PSTN by adding new functionality (packet-switched data network) to support the special needs of the research community (peer-to-peer computer communications). Most of the incremental investment in routers, servers, and access devices (PCs) was undertaken by new types of providers (Internet Service Providers or ISPs) and by end-users (Customer Premise Equipment or CPE) to complement the PSTN basic infrastructure already in place.

With the commercialization of the internet in the 1980s and its emergence as a mass market platform for broadband communications in the 1990s, the internet has evolved into the principal platform for our global public communications infrastructure. Increasingly, IP packet transport is providing the basic transport medium for telephony and other multimedia applications (voice, video, and data). What was an "overlay" application has now become basic infrastructure. Over time, the traditional PSTN providers have come to play a larger role in managing the infrastructure and investment required to support the internet. Deregulation, market growth, and innovation have resulted in a more complex and interdependent internet infrastructure ecosystem.

The success of the internet owes much to the interoperability and connectivity supported by ubiquitous adoption of the IP protocols and adherence to the "end-to-end" (e2e) design principles that have governed internet architecture for so long. However, the internet's success has also posed significant problems. Growth has brought heterogeneous services (not everyone needs or wants the same capabilities); new needs and requirements (support for real-time services or enhanced security); and complexity and size issues (arising from the sheer magnitude of today's internet measured in terms of traffic and connectivity).

To meet these challenges, the internet needs to continue to evolve. In a process that looks at times like history repeating itself, the internet is now spawning its own collection of "overlay" networks. There are many types and examples of overlays (see table 1) that arise to meet a range of purposes and needs (see further discussion in section II). The emergence of these overlays raises interesting questions for the future of internet architecture and the role of the internet as a common platform for global communications. For example, are these "overlays" precursors to the future architecture of the internet? Or, are they nasty barnacles on the internet that threaten the end to end connectivity and interoperability that have proven to be such a key aspect of the internet's value? What are the implications of overlays for industry structure and for the regulation of our public communications infrastructure?

Before it is possible to answer such questions intelligently, it is necessary to have a better understanding of what constitutes an overlay, the motivations for their deployment and use, and the potential conflicts and tensions that may arise among stakeholders. The goal of this paper is to frame such a discussion and provide further thought on the implications of overlays for Internet architecture, industry structure/business strategy, and public policy. As our analysis demonstrates, the policy questions raised by overlays are multifaceted and diverse.

Table 1: Examples of overlay networks

Type	*Function/Purpose*	*Example*
Peer-to-peer (p2p)	File sharing (mp3s)	Napster, Gnutella
Content-delivery (CDN)	Content caching to reduce access delays and transport costs	Akamai, Digital Island
Routing	Reduce routing delays, resilient routing overlays	Resilient Overlay Network (RON), Akamai SureRoute
Security	Enhance end-user security, privacy	Virtual private networks (VPNs), onion routing (Tor, I2P), anonymous content storage (Freenet, Entropy), censorship resistant overlays (Publius, Infranet, Tangler)
Experimental	Facilitate innovation, implementation of new technologies, experimentation	General purpose (PlanetLab, I3)
Other	Various	Email, VoIP (Skype), Multicast (MBone, 6Bone, TRIAD, IP-NL), Delay tolerant networks, etc.

The balance of this paper is divided into three main sections. The second section offers a preliminary taxonomy for thinking about overlays that reflects the rationale for their existence/emergence and provides further elaboration of the sorts of technical, business/economic, and policy questions that overlays raise. The following section illustrates these questions in the context of three examples of overlays (content delivery, routing, and security and privacy).

■ Towards a taxonomy of overlays

In this section we provide a taxonomy for thinking about overlays that includes examining the different motivations for why they emerge. This proves relevant when thinking about the challenges overlays pose for the evolution of technology, industry structure, and communications policy. Before considering these motivations and the types of challenges, however, we need to define what constitutes an overlay.

What is an overlay?

The brief description in the introduction and list of example overlays in table 1 includes a diverse array of networks and network technologies that appear to exist "on top" of the infrastructure that supports the general purpose internet while also appearing distinct from pure end-user or "distributed applications." The services range from special purpose systems that provide advanced services like customized routing to experimental networks like PlanetLab that exists as a general platform for deploying multiple overlays [1].

As a starting point, we offer this definition of an overlay:

> An overlay is a set of servers deployed across the internet that:
>
> a) Provide infrastructure to one or more applications,
>
> b) Take responsibility for the forwarding and handling of application data in ways that are different from or in competition with what is part of the basic internet,
>
> c) Can be operated in an organized and coherent way by third parties (which may include collections of end-users)

We focus in this paper on overlays that are not thought of as part of the basic internet or provided by today’s network service providers. However, we note that our overlay definition encompasses the email infrastructure of the internet. The infrastructure of mail relays and mailbox servers is an overlay service, just one that happens today to be operated by network service providers, offering one example of how overlays may evolve.

In the balance of this section, we consider how overlays relate to the rest of the internet in a loose architecture sense (what it means to be "on top" of the general purpose internet), the functionality that may be provided by overlays; and how overlays relate to industry structure (who owns what).

Fitting overlays into the architecture of the internet

We advanced the notion of an overlay as existing "on top" of the basic internet, while being "infrastructure-like" in that the overlay is a component of or input to the applications that use the internet infrastructure. From this, it is tempting to resort to the computer science conception of protocol layers and to view overlays as a "middle layer" above the basic IP protocols, but below the application layer.

[1] Planetlab is a large set of servers, distributed all across the internet, which host programs that support applications of one sort of another. Planetlab is best thought of as a highly distributed service that makes it easy to deploy new overlays (see, for example PETERSON, ANDERSON, SHENKER & TURNER, 2005).

However, this layered view of overlays addresses only part of their architecture importance. A significant reason that overlays are interesting is that they raise challenging questions regarding the end-to-end design principles that have guided the placement of functionality and services in the internet architecture historically. Overlays change where functionality and services can be placed and which entities operate that functionality and services.

To see how overlays relate to the end-to-end design principles [2], it is necessary to consider how overlays relate to the evolving architecture of the internet. The internet, in its most simple conception, has two sorts of components, end-nodes - the computers at the edge of the net that play the role of servers and user machines and routers, which forward the packets between the end-nodes. In this simple view, one can think of the internet as a cloud of connected routers with end-nodes connected around the edge of the cloud.

End-nodes run application processes (for example games and email and web clients and servers). Routers just forward packets and do not know about application-level functions. In this simple model, an application is a program on one end-node that communicates to a related program on another end-node by sending packets that are forwarded by routers.

Many real applications are more complex than this simple model would imply. In the case of email, for example, when one user sends mail to another, it goes by way of intermediate servers that have names like "the smtp server" and "the pop server". In the case of the web, there are web caches and proxies. A network purist might make a definitional distinction and say that since these intermediate devices are not routers, they must be some form of end-node, and so the simple conception of all devices as either routers or end-nodes still applies.

However, intermediary devices such as the mail and web server have some important characteristics that distinguish them from conventional end-nodes. Firstly, they are distributed around the internet in a way that provides an infrastructure on which the application runs. From the perspective of the

[2] The "end-to-end" design principle specifies that processing of the communications protocols ought to occur in the end-nodes. This allows heterogeneous edge devices and applications running on them to share a common communications infrastructure that supports peer-to-peer communications. See SALTZER, REED & CLARK (1984) or BLUMENTHAL & CLARK (2001) for further discussion. See LEMLEY & LESSING (2000) or ISENBERG (1997) for discussion of policy rationales for preserving the principle in the internet.

router, they may be an end-node, but from the perspective of the application, they are infrastructure. Secondly, they tend to be provided and operated by third parties. If two users exchange mail, they depend on servers operated by others to forward that mail.

The introduction of these organized intermediary devices then represents the introduction of new capabilities to the internet architecture at locations where the community does not generally have clear taxonomies or ways of thinking about the implications of new functionality and services. The functionality and services is not provided by a conventional end-node, but it is certainly not provided by a router either. To distinguish these infrastructure devices from conventional end-nodes the term "overlays" is employed to describe organized sets of infrastructure devices.

Adding functionality

A second way to understand overlays is to focus on the extra functionality offered by an overlay, beyond what is supported by the basic internet. In the context of this discussion, basic internet functionality is defined by the suite of core internet protocols (IP, TCP, UDP, DNS, BGP). These comprise the minimal set of basic protocols that any network or node must support in order to be considered part of the internet (collectively, we will refer to these as "the basic IP protocols" except when the context requires a more specific use).

The ubiquitous deployment of these basic IP protocols helps account for the internet’s great success and growth. The ability for the basic IP protocols to be supported on many different physical infrastructures (ATM, Frame Relay, SONET, wireless) and to support many different types of end-user applications (data, voice, video) helps promote connectivity and interoperability across heterogeneous infrastructure and applications and has helped the internet evolve into a global communications platform. This view of the internet is associated with the "hourglass" model (Computer Science and Telecommunications Board, 1994) where the basic IP protocols are located at the waist.

However, the basic functionality of "best effort" service offered by the standard IP protocols is not always enough. The original internet architecture was designed to support (unicast) communication between fixed locations where the source knew the address of the destination. Yet many applications have more general communication needs such as mobility, multicast, and anycast. These communication needs present new challenges that the current internet architecture does not support, such as when the source does not know the destination address (as in multicast and anycast) or the location of the receiving host is not fixed (as in mobile communications) [3]. Table 1 covers a range of functional extensions that overlays can provide, including mobility, customized routing, Quality of Service, novel addressing, enhanced security, multicast, and content distribution.

In order to address these needs and overcome the barriers inherent in the existing infrastructure, overlays blur the clean internet architecture distinction between packet forwarding and application processing. Overlays, as opposed to application-specific network solutions, are increasingly seen as the mechanism of choice for introducing functionality into the internet. This is important because overlays have become a primary means for evolving the internet architecture.

Overlays and industry structure

For many, the internet technology is a "black box". The details of technical design are not regarded as relevant in themselves, so long as the capabilities continue to evolve to meet the demands of growing markets for enhanced electronic communication capabilities. From this perspective, it might be less interesting to focus on "what the internet does or should do," and to focus instead on the industry structure (who owns what) that supports the internet.

Just as today’s overlays leverage the underlying physical infrastructure of the basic internet, the internet began as an overlay on the basic telephone network. Importantly, the phone network was subject to substantial government regulatory oversight. As part of this oversight, telephone companies were limited in the range of services they could offer and were required to provide non-discriminatory (common carriage) access to all users. This enabled new types of communication service providers, ISPs, to lease physical infrastructure from the underlying telephone companies and combine it with packet switching technology (routers and servers) to give rise to the internet.

[3] For example, see STOICA, ADKINS, ZHUANG, SHENKER & SURANA (2002).

Today, the ISP business is competitive and many of the services offered have been largely commoditized. In addition to providing the basic packet transport services supported by the basic IP protocols, ISPs also provide a host of other commoditized services like email, web hosting, and increasingly other functionality like instant messaging services. For many who take the industry perspective, basic internet infrastructure should be defined as the services that are provided by the typical ISP since that is who historically has provided basic internet services.

This suggests that overlays might be identified with offerings from non-ISP "third parties" that operate in conjunction with the basic internet services offered by the ISPs. Akamai, a provider of CDN services, is an obvious example. A less clear example is when the third party is comprised of a group of end-users as is the case with a peer-to-peer network. Many interesting and emerging overlays (peer-to-peer networks and routing overlays) are first deployed by edge users in end-nodes and may not generally be thought of as "infrastructure providers."

As traditional industry boundaries blur, however, the definition of what constitutes an ISP has become less clear. As ISPs seek to differentiate themselves (to escape the pressures of competition in a commodity market), they add services, vertically-integrating to become so-called "enhanced service providers". These include ISPs that offer data storage and back-up services as well as today’s overlay services (CDN or routing services). Additionally, other types of information technology infrastructure providers (Oracle, Microsoft) and physical infrastructure providers (Verizon, Comcast) are integrating into the ISP space.

In this changing environment, the notion that overlays are offered by some "third party" needs to be thought of as relative to the changing role of ISPs and what constitutes the "basic internet". Overlays exist between that which is provided by ISPs as part of our global communication infrastructure and the applications that ride on top of the infrastructure.

Why do overlays emerge?

To understand why overlays are important and the role they play, it is worthwhile considering the various forces that motivate their emergence. Firstly, overlays may exist to support the special requirements of a particular class of application or user community. As the internet becomes more pervasive across all segments of the global economy and becomes the general platform supporting all types of electronic communications, we would expect the need to address specialized, heterogeneous interests to increase.

Thus, the very success of the internet as an open standard leads to the need to satisfy heterogeneous requirements that, in itself, provides one justification for an overlay. It is worth noting, however, that even in this case, drawing the boundary between what functions are really general and hence should be implemented as "basic infrastructure" and which are specialized and so may be appropriate as "overlays" is neither static (i.e., yesterday's niche might be tomorrow's mass market need) nor clear (i.e., are mail servers basic internet infrastructure or part of an overlay?).

Secondly, and related to the above, overlays may play a role in the dynamic evolution of Internet technology. One of the great advantages of the internet's end-to-end architecture is the ability to incrementally deploy/adopt edge-based innovations. Applications can be deployed virally by a growing number of edge-nodes without requiring modifications to the basic internet. The ubiquity with which TCP/IP are implemented provides a stable platform to support communications among and across heterogeneous edge-nodes. However, the very benefit of ubiquitous availability becomes a challenge when it comes to upgrading the internet's own basic infrastructure.

Overlay networks can provide a way to first experiment with new routing and architecture designs (PlanetLab) and then as a way to incrementally deploy new solutions. Functionality that is missing in the current internet may be deployed first in an overlay for those users/uses that most require (and are willing to pay for) enhancements that may not be available yet in the general internet. This can include such things as enhanced quality of service or security/privacy. Over time, successful innovations will become ubiquitously adopted and, as such, de facto components of basic internet infrastructure. Whether the functionality offered by an overlay is viewed as an enduring specialized need (static) or simply the early version of functionality that is generally needed and will be deployed in basic infrastructure over time (dynamic) may be a question of perspective.

Thirdly, overlays may arise because of conflicts in stakeholder interests, reflecting a tussle between and among customers, service providers, and policy-makers. For example, the missing functionality (privacy, caching support for content delivery, or delay-minimizing routing) may be intentional. Routing overlays that seek to improve on the basic internet route selection process (BGP), may be in conflict with policy-based routing implemented by

ISPs in response to other, non-delay-related considerations (long-term interconnection agreements or regulatory-jurisdiction issues). Or, overlays that implement privacy (obscure the source, content, or type of traffic) might be in conflict with public policy rules that seek to make traffic auditable by the police (CALEA) or to support carrier efforts to price discriminate (price voice higher than data).

The forces that lead to each of these rationales (heterogeneous interests, dynamic evolution, and tussle) are fundamental and enduring and so we should expect that overlays will remain an important and growing feature of the internet landscape.

Implications of overlays for architecture, industry structure, and policy

Overlays pose an interesting challenge for policy in multiple dimensions. They raise important questions for the evolution of internet technology (architecture), business/industry structure (how do overlays impact infrastructure costs and who owns/controls what investment?), and regulatory policy (how do overlays interact with open access, interconnection, or other basic infrastructure policies?). Before turning in section III to a more detailed discussion of these challenges in the context of specific types of overlay networks, we highlight some of the generic issues that overlays pose for technical, economic, and policy analysis.

Implications for internet architecture

To understand how overlays pose a challenge for the evolution of internet architecture, it is worth considering overlays in the context of the internet's "end-to-end" design principle, that has played such an important part in the internet's architecture to date [4].

The end-to-end principle can be thought of as operating at two levels in the internet, the packet level and the application level. At the packet level, the end-to-end principle leads to the original design of the internet where the routers know nothing about applications, but just forward packets, and knowledge of the applications is confined to the end-nodes. This view is consistent with the simple model of the internet we offered in our definition of

[4] See note 2 *supra*.

overlays. However, at the application level, the end-to-end principle could be interpreted as leading to an application architecture in which data is transferred directly and without intervention from the original source to the ultimate destination. By this definition, many real applications and overlays today are not consistent with end-to-end.

Consider the email example we discussed above. Normally, mail goes from sender to receiver via two servers, an smtp server and a pop server. There is no end-to-end confirmation that the mail actually gets to the receiver - the sender and receiver depend on the servers to be reliable. The example of careful file transfer from the original end-to-end paper (SALTZER, REED & CLARK, 1984), which involved confirmation from the destination to the source that the data was delivered correctly, is totally missing from email. From the application perspective, email is not consistent with end-to-end [5].

Seen from the packet level, overlays that support one or a class of applications do not in themselves erode the end-to-end principle. Overlays that provide a new form of the basic internet service (routing overlays) also do not erode the end to end principle. They simply signal that the user wants to purchase a service that the ISPs do not want to offer them.

On the other hand, the end-to-end principle is traditionally important because it helps preserve the openness of the internet. This traditional openness has made it, in general, possible for anyone to communicate with anyone else on the internet. (This feature is also a drawback when one of the communicants is malicious or disruptive.) Overlays, however, may be a means to build "gated communities in cyberspace".

Commercial implications

To understand the commercial implications of overlays, it is worthwhile considering how they impact industry structure. We have already mentioned how overlays shift the ownership of infrastructure functionality between basic (ISPs) and enhanced service providers (CDNs like Akamai), which in turn influences the allocation of costs/profits across the infrastructure value chain. It remains an open question who will or should own and operate overlays in the future.

[5] In contrast, voice-over-IP is an example of a design with a very high-level end-to-end error check - if the message is garbled the human listener can say, “What?”

In addition, overlays have important implications for industry costs and the realization of scale and scope economies. The global market for equipment and support services for basic internet technology helps keep markets competitive, which keeps costs low and expands consumer choice. The global network also realizes large positive network externalities.

Yet overlay networks pose a threat of fragmenting markets if they lead to islands of incompatible and different overlay technologies and bundles of functionality. If this occurs, it may lead to reduced scale/scope economies and reduced network externalities. In the most extreme case, the proliferation of heterogeneous overlays may threaten internet connectivity. Furthermore, overlays may result in redundant investments in overhead/functionality. For example, using an overlay to address routing deficiencies in the basic internet may be more costly than upgrading the basic routing capabilities. However, the need to upgrade implies adjustment costs and, perforce, invokes dynamic considerations.

Implications for public policy

The tussle over internet technology and industry structure leads directly to a tussle over public policy. In recent decades, industry convergence has been challenging the traditional silo-based model of communications regulation whereby cable companies are treated differently from telephone companies, mobile providers from fixed line providers, and basic service providers from enhanced service providers. These regulatory distinctions are increasingly problematic in the emerging world of converged, broadband platform services.

One solution that has been proposed to this dilemma is to migrate toward a layered model of regulation (SICKER, 2002; WERBACH, 2002; FREIDEN, 2002). While such a strategy may prove appealing, overlays challenge possible definitions of what constitute appropriate layers (how to distinquish between basic and enhanced telecommunication services?). Moreover, even when the regulatory distinction can be made, enforcement is difficult when the technical architecture blurs boundaries. For example, there are many ways to deploy VoIP or DRM services that makes it challenging to craft a simple policy to enforce the phone tapping rules (CALEA) or copyright protection.

Alternatively, overlays may be used to circumvent communication (access charges or CALEA)[6] or other policies (copyright protection or censorship rules)[7]. Finally, because overlays challenge industry structure, they impact competition policies (open access rules, antitrust)[8]. For example, asymmetric access to the services of a content-delivery network may result in preferential access to some content at the expense of others. Whether this is a problem or not may depend on ones perspective (is it a specialized service to serve a specialized community that is paying for the incremental services provided, or is it a violation of common carriage non-discrimination rules?).

■ Challenges for different types of overlays

In this section we provide an overview of the technical, commercial, and policy challenges posed by three different types of overlays: Content Delivery Networks (CDNs), routing overlays, and security overlays. For each example, we provide a description of how the overlays operate and then identify some of the interesting questions raised by the growth of such overlays.

Content Delivery Networks (CDNs)

CDN overlays address a fundamental challenge on the internet – how to cost effectively distribute and acquire content while simultaneously lowering latencies experienced by end-hosts.

At a technical level CDNs consist of caches of content and services distributed across the internet. These caches contain copies of content and

[6] VoIP has been used to bypass regulatory-mandated intercarrier payments both domestically and internationally.

[7] A key driver for the growth of peer-to-peer file-sharing programs like Gnutella or Napster was the opportunity to trade MP3 files in violation of copyright rules.

[8] For example, a key policy for implementing open access in telephone services was structural separation of long distance and local telephone services. Local telephone providers were required to provide non-discriminatory interconnection to all long distance providers, thereby enabling the growth of competition in long distance. In the USA, this was supported by restrictions that prohibited the largest local telephone companies from offering long distance services. Once local companies are allowed into long distance, they have an obvious incentive to discriminate against non-affiliated carriers.

services retrieved either on-request or proactively from publishers and providers. The heart of a CDN is the method by which requests and content are routed and redirected in the overlays to accomplish the load balancing. Example ways of accomplishing this routing and redirection in distributed CDNs include DNS and URL rewriting and http redirection. In balancing the request and content load, CDNs optimize different criteria including technical measures such as response time and server loads and economic measures such as bandwidth costs.

CDNs fall into three distinct categories 1) commercial 2) cooperative and 3) peer-to-peer based overlays. Commercial CDNs distribute dynamic services such as on-line airline reservation applications and static content such as patches to Microsoft software products. Currently the largest commercial provider of CDN services is Akamai, which claims to serve approximately 10-15% of web content and collocates with around 10,000 ISPs globally [9].

Cooperative CDNs such as CoralCDN and OpenCDN seek to offer similar benefits to non-commercial users. Unsurprisingly, the performance of these networks is often not as good since they rely on voluntarily contributed infrastructure. They do however offer the potential for content publishers to reach a larger audience and sustain service during larger flash crowds than would be possible from a resource limited server. Any web publisher for instance can "Coralize" their web URLs and cause their content to be cached in the CoralCDN by appending nyud.net:8090 to their URLs (rewriting http://www.x.com into http://www.x.com.nyud.net:8090).

Finally, many of the peer-to-peer overlays function as content distribution networks. Peer-to-peer content distribution networks differ from cooperative caches in terms of functionality (often including for instance search capabilities), content (larger percentage of files likely to raise issues of copyright infringement), and overlay structure (majority of the nodes are both servers as well as clients whereas in the cooperative caching networks many nodes are contributed purely altruistically and serve only as content caches).

Currently a significant amount of content is being served from the decentralized peer-to-peer caching overlays. BitTorrent, in particular, has been a popular distribution network in the past few years. Notably, BitTorrent

[9] Estimate provided in private conversations with Akamai personnel.

first gained popularity as a distribution channel for Linux software distributions. BitTorrent remains vital to the distribution of content from a number of popular publishers that rely on the peer-to-peer CDN to lower their distribution costs. Some (much of?) this content might not be available in the absence of these cost savings. Thus, peer-to-peer content distribution may lower the costs of accessing diverse content (a public good!) while at the same time providing a platform for copyright infringement (a policy challenge!).

Questions:

- As CDNs evolve how will they effect the industrial organization of the internet? Who should own/control the content caches?
- If CDNs provide superior access selectively to some content, would this give rise to a two-tiered internet: one that is high quality for commercial content and one that is lower quality for non-commercial content? If so, would this raise concerns about equal, non-discriminatory access or free speech? Should CDNs have net neutrality obligations?
- Suppose an ISP sought to vertically integrate with a major CDN provider like Akamai. Would that raise antitrust concerns?

Routing overlay networks

A routing overlay is an overlay that exists for the purpose of controlling or modifying the path of data through the network. In a routing overlay the endpoints of the information exchange are unchanged from what they would have been in the absence of the overlay, but the route through the network that the packets traverse may be different [10].

The routing overlay is unique among classes of overlays we discuss in this paper because the overlay network performs a function that is already implemented by the existing internet infrastructure. In contrast with other classes of overlays, which exist to provide new functionality, routing overlays in their purest form exist to change the way an existing function is performed. It is this overlap, between routing as a base internet function and routing as an overlay network function that leads to the most interesting properties of routing overlays.

[10] For example, this contrasts with a CDN, wherein the source/destination addresses of the communicating pairs of nodes may be changed.

Routing - the determination of a path between the source and destination of transmitted data - is a basic function of all computer networks. In any reasonably large network, there will be several possible paths between any given source and destination. When more than one path is available, the object of routing is to choose the "best" path. This choice is not necessarily obvious even in a single network.

The problem, however, is further complicated in a network of interconnected ISPs, each of which may be making routing choices based on divergent criteria and partial information about the overall status of the Internet. These decisions are driven by a number of factors. Chief among these are the internal structure of the ISP, which determines the cost to carry a packet across that ISPs network and the business arrangements between ISPs which determine the cost of exchanging traffic across ISP boundaries.

These individual decisions are coordinated by a network protocol known as the Border Gateway Protocol (BGP). Broadly speaking, BGP allows each ISP to express its policies for accepting, forwarding, and passing off packets using a variety of control knobs. BGP then performs a distributed computation to determine the "best" path along which packets from each source to each destination should be forwarded.

This formulation raises two difficulties, one fundamental and one pragmatic. The first of these is that the notion of "best" is in fact insufficient to fully express the routing task. "Best" is a single dimensional concept, but routing is a multi-dimensional problem. Individual ISPs, in making their routing decisions, may choose to optimize a wide variety of properties. Among these might be 1) the cost of passing on a packet; 2) the distribution of traffic among different physical links within their infrastructure to maximize utilization and minimize congestion - so-called traffic engineering; and 3) performance in some dimension, such as bandwidth available to the traffic or transmission delay across the ISP. Furthermore, because the management of each ISP chooses its own objectives, different ISPs may choose to optimize different quantities, leading to an overall path that captures no simple notion of "best", and rarely if ever is best for the user.

A second, pragmatic problem with the current internet routing infrastructure is that it has evolved over time from one in which simple technical objectives dominated to one in which ISPs often wish to express complex policy requirements. For this reason the knobs - the methods available within BGP to control routing choices - have also evolved over time, and are presently somewhat haphazard and baroque. This compounds the fundamental problem by making it harder for ISPs to express precisely the policies they desire, even after those policies are known.

The objective of a routing overlay then is to override the routes determined by the ISPs. While ISPs route traffic based upon cost and operational efficiency, an overlay can route traffic based upon metrics directly related to application performance.

Questions:

- How do routing overlays impact infrastructure costs and who owns/controls what investment?
- How do routing overlays interact with open access, interconnection, or other basic infrastructure policies?
- What is the sustainable business model? Who will run the routing overlay service when it grows up, and who will pay for it. Who is the provider, and who is the customer?
- Would the market accept a global service owned and offered by a single ISP, or would there be concerns about fair and equitable treatment? Would this be an issue the marketplace could sort out on its own?

Security and privacy overlays

The final class of overlay networks we discuss are ones that we broadly characterize as "security and privacy overlays". These overlay networks provide different forms of communication protection (HERSCOVITZ, 1999), user or server anonymity (DINGLEDINE, MATHEWSON & SYVERSON, 2004; CLARKE, SANDBERG, WILEY & HONG, 2000), censorship resistance for online content (WALDMAN & MAZIERES, 2001; DINGLEDINE, MATHEWSON & SYVERSON, 2004; Entropynet, 1984; FEAMSTER, BALAZINSKA, HARFST, BALAKRISHNAN & KARGER, 2002), or deniability of the knowledge of traffic (CLARKE, SANDBERG, WILEY & HONG, 2000) or content (WALDMAN, RUBIN & CRANOR, 2000). This is a particularly interesting class of overlays because even if the volume of traffic on these overlays is not large, the policy and social implications can be significant.

In many ways these overlay networks mirror the content and routing overlays. Security overlays change the routing and caching behavior of communications and content on the internet. The difference is that instead of changing the behavior to optimize performance or money flows, these overlays enhance some aspect of end-user security or privacy. Some provide for secret communications or anonymity for end users; others make content robust against attempts of powerful adversaries to remove it from the internet and enable users to establish legal deniability of traffic or content ownership.

This class of overlays tends to make the internet opaque to regulation, easily frustrating policy makers' objectives. Through clever use of cryptographic techniques and system engineering these networks provide provable properties about how hard they are to break or the legal deniability afforded to network participants. In many ways, this class of overlay networks re-raises questions from the encryption debates in the 1990's (HOFFMAN, ALI, HECKLER & HUYBRECHTS, 1994).

While encryption hides the content of communications in a network, some of these overlays hide the entire network.

Most security and privacy overlays are limited in terms of current deployment, but are significant for the policy questions that they raise. Table 3 below provides a descriptive summary and examples of several types of security overlays. Many of the security and privacy overlays have overlapping goals. We classify the networks by the functionality they are most commonly associated with.

Table 2 - Description and examples of prominent security overlays

Onion routing overlays	Onion routing networks, or mix nets, are overlay networks that enable pseudo-anonymous communication over the internet. Current examples include the Tor (DINGLEDINE, MATHEWSON & SYVERSON, 2004) and I2P (I2P Anonymous Network) networks.
Anonymous content storage and retrieval overlays	These overlays protect the identity of authors, publishers, and content providers when they store, query, and download content from the internet. Current examples include the Freenet (CLARKE, SANDBERG, WILEY & HONG, 2000) and Entropy (Entropynet, 2006) networks.
Censorship resistance overlays	These overlay networks attempt to make it very difficult for powerful adversaries to remove content or pollute the overlay network with distracting materials. Current examples include Publius (WALDMAN, RUBIN & CRANOR, 2000), Infranet (FEAMSTER, BALAZINSKA, HARFST, BALAKRISHNAN & KARGER, 2002) and Tangler (WALDMAN & MAZIERES, 2001) networks.

Questions:

- Where will the balance between freedom of speech, civil liberties and law enforcement needs be found for users of security overlays?
- Many security overlays are designed to diminish the threat from powerful adversaries. But this also makes enforcing things like copyright enforcement or identifying criminal perpetrators next to impossible. How will this fundamental philosophical clash be resolved?
- How do ISP obligations like CALEA apply to operators of security overlays?

■ Conclusion

The internet emerged as an overlay on the telephone system, and triggered a massive shift in the structure of the telecommunications industry, with economic, policy and social implications. We believe that overlay systems on top of the internet may signal yet another shift and while overlays may not be as dramatic or a far-reaching as the internet itself, they again have important economic, policy and social implications.

Given that the internet has matured and the landscape of functions and applications is somewhat set, overlays will be the source of disruptive innovation. They represent a way to innovate at a higher level, creating a new order, with a new set of players, a new economic landscape, and a new set of rules. As the discussion here makes clear, however, there are many kinds of overlays and they arise to serve different purposes. Understanding them will pose a challenge for policy-makers and industry participants, but they are where a lot of the future action is likely to be.

References

BLUMENTHAL Marjory & David CLARK (2001): "Rethinking the design of the Internet: the end-to-end arguments vs. the Brave New World", *ACM Transactions on Internet Technology (TOIT)*, vol. 1, issue 1, August, pp. 70-109.

CLARK David, William LEHR, Steve BAUER, Peyman FARATIN, Rahul SAMI & John WROCLAWSKI (2005): "The Growth of Internet Overlay Networks: Implications for Architecture, Industry Structure and Policy", 33rd Telecommunications Policy Research Conference, September, available at:
http://web.si.umich.edu/tprc/papers/2005/466/TPRC_Overlays_9_8_05.pdf.

CLARKE I., O. SANDBERG, B. WILEY & T. HONG (2000): "Freenet: A Distributed Anonymous Information Storage and Retrieval System", Proceedings of Designing Privacy Enhancing Technologies: Workshop on Design Issues in Anonymity and Unobservability, July, pp. 46-66.

CSTB (1994): Computer Science and Telecommunications Board, National Research Council, Realizing the Information Future: The Internet and Beyond, National Academy Press, Washington DC.

DINGLEDINE Roger, Nick MATHEWSON & Paul SYVERSON (2004): "Tor: The Second-Generation Onion Router", Proceedings of the 13th USENIX Security Symposium, August.

Entropynet (2006): "Entropy (Emerging Network To Reduce Orwellian Potency Yield)", Project website available at: http://entropy.stop1984.com/

FEAMSTER N., M. BALAZINSKA, G. HARFST, H. BALAKRISHNAN & D. KARGER (2002): "Infranet: Circumventing Web Censorship and Surveillance", Proceedings of the 11th USENIX Security Symposium, August.

FREIDEN Robert (2002): "Adjusting the Horizontal and Vertical in Telecommunications Regulation: A Comparison of the Traditional and a New Layered Approach", mimeo, Pennsylvania State University.

HERSCOVITZ (1999): "Secure virtual private networks: the future of data communications", *Int. J. Netw. Manag*. 9, 4 August.

HOFFMAN L.J., ALI F.A. HECKLER S.L. & HUYBRECHTS A. (1994): "Cryptography policy", *Commun. ACM* 37, 9 September.

I2PNET: "I2P Anonymous Network": Project website available at http://www.i2p.net/

ISENBERG David (1997): "The Rise of the Stupid Network", AT&T Research Labs, June 4, See: http://www.isen.com/stupid.html

LANZI Diego & Massimiliano MARZO (2005): "Content Delivery and Vertical Integration in On-line Content Markets", *Review of Network Economics*, vol. 4, no. 1, March, pp. 63-74.

LEMLEY Mark A. & LESSIG Lawrence (2000): "The End of End-to-End: Preserving the Architecture of the Internet in the Broadband Era", October, UC Berkeley Law & Econ Research Paper, no. 2000-19.

PETERSON Larry, Thomas ANDERSON, Scott SHENKER & Jonathan TURNER (2005): "Overcoming the Internet Impasse Through Virtualization", *IEEE Computer Magazine*, April, pp. 62-69.

SALTZER Jerome, David REED & David CLARK (1984): "End-to-end arguments in system design", ACM Transactions on Computer Systems (TOCS), vol 2, no. 4 pp. 195-206.

SICKER Douglas (2002): "Further Defining a Layered Model for Telecommunications Policy", draft mimeo, University of Colorado, Boulder, September 1.

STOICA Ion, Daniel ADKINS, Shelley ZHUANG, Scott SHENKER & Sonesh SURANA (2002): "Internet Indirection Infrastructure", Proceedings of ACM SIGCOMM, August.

WALDMAN M. & D. MAZIÈRES 2001) "Tangler: a censorship-resistant publishing system based on document entanglements", Proceedings of the 8th ACM Conference on Computer and Communications Security (CCS 2001), November, pp. 126-135.

WALDMAN M., A. RUBIN & Lorrie CRANOR (2000): "Publius: A robust, tamper-evident, censorship-resistant and source-anonymous web publishing system", Proceedings of the 9th USENIX Security Symposium, August, pp. 59-72.

WERBACH Kevin (2002): "A layered model for Internet policy", *Journal of Telecommunications and High Technology Law*, vol. 1, pp. 58-64.