NDN, Technical Report NDN-0018, 2014. http://named-data.net/techreports.html DRAFT
Revision 1: 11 April 2014.

A World on NDN:
Affordances & Implications of the Named Data Networking Future Internet Architecture

Katie Shilton
University of Maryland, College Park

Jeff Burke
University of California, Los Angeles

kc claffy
CAIDA/UC, San Diego

Charles Duan
University of Colorado Law School

Lixia Zhang
University of California, Los Angeles

Abstract
This paper explores the potential social impacts of Named Data Networking (NDN), a proposed future Internet architecture that forwards packets based on data names rather than host addresses. We highlight four departures from today's TCP/IP architecture, which underscore the social impacts of NDN: the architecture's emphases on enabling *semantic classification*, *provenance*, *publication*, and *decentralized communication*. While all of these principles can be implemented in the current Internet's application layer, NDN enables them at the network layer, and thus encourages all applications to comport with them. We describe how these changes from TCP/IP will expand affordances for free speech, and produce positive outcomes for security, privacy and anonymity while raising new challenges regarding data retention and forgetting. We describe how these changes might alter current corporate and law enforcement content regulation mechanisms by changing the way data is identified, handled, and routed across the Web. We examine how even as NDN empowers edges with more decentralized communication options, by providing more context per packet than IP, it raises new challenges in ensuring neutrality across the public network. Finally, we introduce openings where telecommunications policy can evolve alongside NDN to ensure an open, fair Internet.

1. Introduction

The Internet has permeated the economic, political, cultural and social domains of global society and transformed the way in which we present and communicate knowledge. The infrastructure underlying these communications continues to evolve, with ramifications for not only the technical protocols that govern the way the Internet functions, but also for social, economic, and legal issues. Internet protocols affect debates about intellectual property, cyber security, and the basic performance and reliability of Internet services.

This paper discusses a proposed future Internet architecture that changes how data is delivered over the Internet: Named Data Networking (NDN). At this early point in NDN's development, this paper aims to stimulate discussion about the relationship between the architecture and the society it aims to serve. Our goal is not to present the complete story of all possible implications of NDN or even to dwell on the many technical benefits that we believe it offers. Rather, it is to expand on the current literature, e.g., (Jacobson et al., 2009), (Zhang et al., 2010), and (Jacobson et al., 2012) by outlining potential social opportunities and challenges that would arise from the architecture's widespread deployment. Most of these opportunities and challenges exist in one form or another in the current IP Internet—especially at the application layer, where data-centric semantics are often already used—but they are made more pervasive and impactful by NDN's push of that data-centricity down to the network layer used by every application and device on the Internet.

We explore NDN's four departures from TCP/IP, which underscore its potential social impacts. The first is an emphasis on *semantic classification* (by applications) through per-packet names of data used by the network layer for routing and forwarding. The second is *provenance*, the idea that all data objects are linked to their creators—in NDN, through a cryptographic signature. The third is *persistent publication*, enabled by in-network storage and abstractions oriented around data dissemination rather than virtual channels between hosts. The fourth is *decentralized communication*, the principle that devices should be able to communicate directly if they can physically reach each other, or along available device-to-device paths which may not involve any Internet Service Provider (ISP). Again, while all of these ideas can be, and often have been, implemented in the *application layer* of today's Internet, NDN enables these values at the *network layer* by design, and encourages development of applications that comport with them.

NDN's architecture represents, as many infrastructures[1] do, the architects' "shared visions of the possible and acceptable dreams of the innovative" (L. L. Bucciarelli in Star, 1999). These changes, and this representation of the possible, in turn will produce changes for the social aspects of the Internet, including privacy, intellectual property, law enforcement, governance, and political economy. By considering potential social impacts engendered by the NDN architecture while it is still in development, we hope to shape an Internet that not only works more efficiently and provides increased reliability and trustworthiness in communication, but even more fundamentally supports privacy, democracy, and equity of information access. The intentionally (by design) non-proprietary and open nature of current Internet protocols have shaped a platform for innovation in communications, demonstrated by unprecedented development of a wide range of private, public, and social goods. Technical design decisions and the framework for core algorithms and protocols of any new architecture will crucially affect its uptake and the ability to support a similar level of socially beneficial innovation.

If we take seriously the notion that running code shapes rights, behavior, and governance (DeNardis, 2012; Lessig, 2006), then analyzing how NDN would alter that code – the technical infrastructures we rely on every day – is an important task. This paper is intended to start a conversation about how NDN might alter that code. First, it lays out the fundamental architectural components of NDN and key differences between NDN and TCP/IP. It then uses these differences to reflect on implications for key societal issues, including free speech, security and privacy, law enforcement, and network neutrality.

---

[1] Here and throughout this paper we consider the Internet's function as an infrastructure, like a power grid or water system: a basic service fundamental to operating a society or enterprise. Though the architecture is abstract, its global instantiation as an infrastructure is not.

2. Fundamental Architectural Components of NDN

A team funded by the National Science Foundation – led by Principal Investigators from UCLA and incorporating Co-PIs, staff and students from other U.S. institutions as well as national and international collaborators[2] – has pursued research to design and evaluate a replacement architecture (Figure 1) for the foundational layer of the Internet: the Internet Protocol (IP). IP relies on host addresses to route packets across the network. Data publishers apply for names from domain registrars, and the names are then mapped to IP addresses assigned by their service providers. Data is retrieved according to *where* (on which host/IP address) the data is located. NDN, an entirely new architecture, operates on named data directly, without translating to the address of their containers. By naming data, NDN enables applications to retrieve their desired data by *names*, and enables data to be cached anywhere in the network(Jacobson et al., 2009). NDN focuses on the *what* rather than the *where* of IP. Each piece of data is signed by its producer together with its name, securely binding them. The signature is mandatory, and coupled with data publisher information, enables determination of data provenance, allowing the consumer's trust in data to be decoupled from how (and from where) the data is obtained.

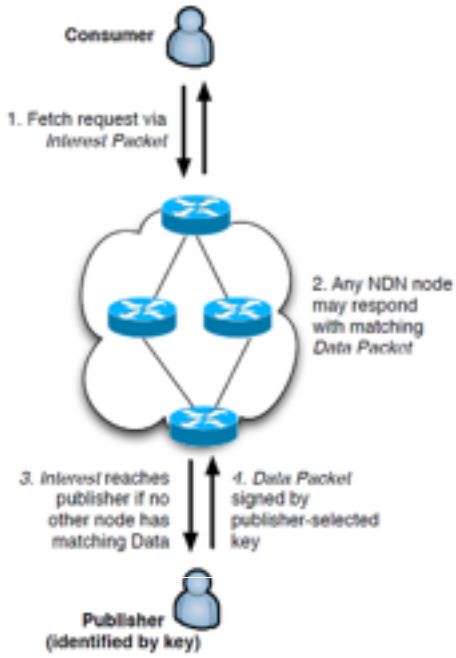

To facilitate data exchange, NDN relies on four key architectural components: names, content signatures, in-network storage, and Interest-Data exchange.

2.1 Names

In NDN, applications name Data packets. NDN does not predefine name conventions. However, NDN application developers will likely develop standard naming conventions over time. For example, in the same way that hierarchical IP addresses have enabled global scaling in the current Internet, hierarchical names in NDN facilitate scalable routing, and will enable data to be found and fetched in a consistent way across the Internet. Organizations similar to those that manage IP address and domain name assignments will likely manage globally unique name prefixes (say, /ucla). Only globally-accessible data require globally distinct names; institutions, individuals, and applications may make use of local names for traffic intended for local use only.

2.2 Content signatures

An NDN network requires each piece of content to be signed by a key that binds the data content to its name. If the signature can be coupled with data producer information in the form of a key locator, it enables determination of data provenance, and serves as the basic building block of security in NDN (Jacobson et al., 2009). This signature securely binds together the tuple – <name, content, publisher's key > – authenticating that the data is what its name purports it to be. Asignature produced by a trusted key signals that a consumer can trust that the data originated from the holder of the trusted key, regardless of from where the data was retrieved. NDN does not dictate a particular trust architecture (to determine what constitutes a trusted key), but the presence of content signatures enables and facilitates a variety of trust management systems for data-centric security.

---

[2] See the Named Data Networking website, http://named-data.net/ for a full list of participants and collaborators.

2.3 In-Network Storage

Because signed data can be retrieved from any device, applications running over NDN networks can utilize in-network storage to achieve performance and scalability enhancements. Such in-network storage is similar to that provided by content distribution networks (CDNs) in today's IP networks, but is implemented at the network level. Therefore, it is available pervasively for all data, without special contracts and complex DNS configurations as needed by current CDN services. Such storage includes in-network ephemeral caches in routers, which may store data as it passes through the network. Storage also includes longer-term repositories (called repos), which are persistent stores deployed for specific namespaces, applications, or networks to replicate data in many places in the network for faster retrieval. NDN researchers are developing new primitives to interact with repos, and are researching efficient synchronization of named data collections.

2.4 Request / Response Data Exchange

NDN dictates a basic communication model based on packet-by-packet request and response. A consumer sends an Interest packet specifying the name of data it wishes to receive. Upstream routers forward Interests towards nodes that have registered data name prefixes. Each router uses a Pending Interest Table (PIT) to record the interface from which it received an Interest, creating a hop-by-hop trail (or breadcrumb) back to the data consumer for each path the Interest takes. When the Interest packet reaches a node which has the requested data, the node responds with a Data packet, which is forwarded back along the trail to the requestor, consuming (i.e. deleting) the PIT breadcrumbs along the way.

The request/response model of NDN provides inherent multicast data delivery, as requests for the same content from multiple consumers are collapsed into a single PIT entry if they flow through the same router. So if a router receives Interests for data with the same name from five different interfaces, the router only forwards the first Interest for that name while recording the incoming interfaces for the other four Interests in its PIT. When the corresponding Data packet comes back, the router forwards it to all the five interfaces.

By design, an NDN network is *loop-free* because each router keeps an entry for each outstanding Interest in its PIT, detecting and discarding duplicates. This behavior allows each router to forward one Interest to multiple upstream nodes simultaneously, and use the *feedback loop* created by the request/response to monitor packet delivery performance and losses across different interfaces. Each node may implement a *forwarding strategy* module to make Interest forwarding decisions. For example, a node may forward Interests across more than one interface simultaneously (e.g., 4G and WiFi) according to cost, measured performance, and other factors.

3. Key Differences between NDN and TCP/IP

The NDN architecture will result in application designs and network configurations that are different from that of the TCP/IP Internet. Consistent, network-visible names will encourage *semantic classification* at the network layer, as names become meaningful to finding, forwarding, and organizing data retrieval. Content signatures produce an architectural emphasis on *provenance* by providing a consistent means to verify data sources. Pervasive, persistent, and standardized in-network storage will encourage a default towards *publication* as content is cached across the web. As discussed in further detail below, NDN's mechanisms for storage and data exchange by names enable *decentralized communication* by making it more straightforward for devices to exchange data directly in a consistent and potentially large scale manner. In the sections that follow, we will illustrate each of these departures from TCP/IP using a brief application example. To best compare the affordances of NDN and TCP/IP to examine potential social impacts, we've chosen application

examples already familiar from a TCP/IP context.[3] We illustrate semantic classification through a use case from the Internet of Things. We discuss provenance and publication through a use case from video publishing. A social network use case further explores publication, while also illuminating possibilities for decentralized communication.

Table 1: Case studies that illustrate NDN's departures from TCP/IP

	Internet of Things	Video Publishing	Social Network
Semantic classification	X	X	
Provenance	X	X	X
Publication	X	X	X
Decentralized communication			X

3.1 Semantic classification of data

By using data names to make routing and forwarding decisions, NDN couples applications' data classification schemes with distribution of that data across the network. This brings the network and applications together in a new way, by explicitly enabling application-specific addressing semantics to affect operation at the network layer, something fundamentally not possible with TCP/IP.[4]

One environment that illustrates the power of such semantic awareness is the *Internet of Things* (IoT), where NDN names provide a richer and more versatile approach than IP to addressing potentially billions of devices across the world. The IoT concept envisions every device, and many objects, as network-enabled and, to varying extents, context-aware. NDN enables the Internet-connected things, and the data they create and consume, to be addressed by one or more application-specific names. For example, a manufacturer-assigned name, such as /local/appliance/kitchen/toaster/Black&Decker/<serial_number>, might be used to address a local kitchen appliance. That appliance would be configured in this namespace at the factory and listen for Interests in its prefix /local/appliance using a power-line or wireless interface. In a simple scenario, other devices in the home would issue Interests on the same broadcast media on a regular basis. Interests for /local/appliance would be used to discover the device when first plugged in; then, its more specific name could be used for direct communication. In this case, NDN enables applications to use the network layer directly to discovery nearby devices in these well-known namespaces (e.g., /local/appliance), without needing the devices to be connected to the global Internet, to have a globally-routable name, or to use middleware. More complex scenarios are a topic of current research: For example, control could be made more secure by authenticating applications via a shared secret, read from the toaster's factory packaging or otherwise transmitted out of band.

[3] Though NDN will also enable totally new applications, these are not our focus here. For discussion of this, see other NDN technical reports and publications on the project website.

[4] It is important to note that the naming semantics are opaque to networks. They come from applications, institutions, and global conventions, and these relationships are reflected in prefix forwarding rules. NDN routers view names as opaque, structured byte strings and simply use byte string matching to identify requested contents or forward Interests towards their producers.

In this case, an application communicating with the toaster for the first time would issue an authenticated Interest that, by incorporating the secret, begins the process of authorizing the application to make updates to the device and otherwise control it. Further, the user could assign a more relevant, globally-routed name, such as /ndn/ucla.edu/jburke/appliance/Melnitz/toaster, which could be used for (probably authenticated) communication.

The IoT example illustrates that semantic classification can facilitate discovery of new devices on a network—from a new light bulb to a digital television—using network names directly. Given appropriate conventions, the approach above could be used for all IoT devices, using NDN's hierarchical naming to make scaling manageable. For example, a television might respond to Interests issued over WiFi in the namespace, /local/appliance/media with a Data packet /local/appliance/media/television/Sony/00-AA-31-49-BC. Like the toaster, it would be discovered by applications issuing Interests in /local, /local/appliance, etc., but not /local/appliance/kitchen, allowing applications to focus on only the types of devices they are interested in.

Because naming conventions will delineate publishing authority (i.e., how a publisher will know what prefix in which it can publish), an open question is who will assign and manage top-level names for both globally routable prefixes (e.g., /ndn/ucla.edu) and local conventions used for standardized mechanisms like discovery (/local/appliance). Further, how language will impact these name-based mechanisms must be considered—e.g., in Japan, the same Sony television above might be addressed using a standardized Japanese namespace rather than an English one. The example of local discovery illustrates that there are many new naming contexts created by the architecture. New authorities may emerge to assign names according to topological, geographical, institutional, device-related, or other schemes.

3.2 Provenance

NDN emphasizes data provenance through one of its basic building blocks, content signatures. These signatures enable Data Packets to identify a producer. For example, in video publishing, consumers identify and verify a video's publisher for each packet by checking the key used to sign every data segment of the video. This mechanism is quite different than trusting the connection over which the video data traveled, as in a TLS/SSL tunnel over TCP/IP. In NDN, data is signed when it is produced, independently of its transmission over the network; the *data* is secured, not the *channel*. Applications can receive signed data packets in NDN from any node that has them—one or more original producers, a network cache, a local repository, etc.—and verify them independently of how they are obtained. This contrasts toTLS/SSL and other channel-based mechanisms that secure only one end-to-end communication at a time.

A viewer or viewing application might trust a content producer's key for a variety of reasons. They may have encountered the key before—perhaps in other videos, or just in previous frames of the current video. Or, the viewer may recognize the key as matching a recognizable and trusted service identity. For example, NDN-Vimeo.com might re-publish content under its own key or a Vimeo-controlled key for the user. Or, such a trusted service might vouch for a pre-existing key. The NDN-Vimeo application could issue a daily key for a user, sign it with their master key, and allow the user to publish his/her video signed by that daily key. In each of these cases, a viewing application would walk the trust chain from the video's key up to the well-known key for Vimeo. Similarly, viewers might choose to trust content by a key connected to a social network identity (e.g., Facebook ID), workplace credentials, or some other form of online identity that is strongly mapped to real world identity.

3.3 Persistent publication

By encouraging caching of such verifiable content in both routers and persistent storage in standardized repositories, NDN encourages persistent publication by default. Consider a video captured by a user on a mobile phone and disseminated to a large number of Internet viewers; a content producer might produce and wish to disseminate a video of a cat riding a robotic vacuum cleaner or a live stream of a protest. On the current TCP/IP-based Internet, the end user uploads a pre-recorded video (the cat) or forwards the live stream (the protest) to a cloud-based service such as YouTube or Vimeo. The service stores the video and responds to web requests with the help of a content-distribution network (CDN) that dynamically distributes popular videos to servers that are topologically close to requestors.

NDN incorporates content distribution support directly into the network layer. The mobile user might maintain data under a home namespace prefix provided by a connectivity provider (e.g., /verizon.net/jburke) or a service provider (e.g., /ndn-vimeo.com/kshilton). The namespace provider ensures that content published by the mobile user and signed by an appropriate key is accessible over the global Internet via the name prefix.

In NDN, the cat or protest video could be published on the mobile phone in the user's home namespace. NDN forwards Interests issued by (potentially many) viewers directly to the phone, which responds with video frames. NDN nodes along the paths to this video source reduce traffic in two ways: 1) by collapsing many requests for the same data into a single forwarded Interest and 2) by caching the returned Data packet. In this way, NDN helps distribute popular content even when it is published from low-capability devices. Caching enables intermediate nodes to respond to subsequent requests from their caches rather than forwarding Interests all the way back to the mobile phone. Collapsing outstanding Interests ensures that only one request is forwarded towards the source, even if the data has drawn multiple Interests.

The same published video could persist in a number of nodes. The original publisher's device might choose to store the content and respond to Interests until it runs out of storage or needs to conserve battery life. Services providing storage for the user could republish it—content producers might pay services (like Vimeo or YouTube) to express Interests for their content and then store and serve it over the long term. Caches would also hold the video, though they might not store the Data for long unless it was very popular, given their primary purpose to reduce network load and provide good performance to viewers. Finally, any node could choose to capture and republish the video without the content producer's intervention, in much the same way that the Internet Archive operates.[5]

3.4 Decentralized Communication

An affordance of NDN related to, but distinct from, publication is decentralized communication. Application-assigned data names, content signatures, persistent storage, and NDN's intrinsic multicast data delivery model together enable peer-to-peer and other decentralized communications models that are impractical or unwieldy under TCP/IP.

Imagining a social networking service, such as Facebook or Twitter, on NDN further reveals how the architecture promotes publication and decentralized communication. To operationalize a basic social network on NDN, one could establish a branded namespace (e.g., /twitter or /facebook) and then procure distributed storage (aka repos) throughout the world for it. Then, the service

[5] Given the ease and potential breadth of republishing in NDN, rather than the conceptually more ephemeral channel of IP-based streaming, encryption-based access control becomes the primary method for protecting content in the NDN Internet. Producers would need to encrypt any sensitive videos on the phone to protect them from unauthorized access.

could provide its users with keys that are authorized to sign content in some child namespace (/facebook/lixia.zhang) and facilitate the creation and replication of that content. Thus, a simple version of NDN-Facebook or NDN-Twitter would function similarly to a content distribution network in a branded namespace, achieving fast content delivery via basic NDN building blocks. This distributed approach enables social network users to publish their content directly in NDN-Facebook's namespace, making their own decision about where to first store it. However, trust is still centralized by having a root of trust originating at a single entity's key.

To control and transform content in ways similar to that of an existing social network, the implementation of such a platform would be more involved. A trust management approach for one of these social networks could work as follows: The end-user application (e.g. the browser or mobile application) would sign new content using either the user's NDN-Facebook key or a key provided by their ISP and signed (authorized) by NDN-Facebook. Then, it would publish the content and associated metadata to a namespace provided by the user's ISP. To provide similar control as in IP-Facebook, NDN-Facebook would retrieve, process/transform, and re-publish the content, signed by a user-specific key controlled by NDN-Facebook or with NDN-Facebook's key and metadata identifying the user.[6]

On NDN, it requires only a shift in the trust model to turn the distributed NDN-Facebook (with a centrally controlled key) as described above into a decentralized one: Content could be published directly by each users, and a web-of-trust model employed rather than a centralized hierarchy with the company's root key at the top. In this case, the social network could be created by a group of people reserving a namespace and then enabling any number of publishers to directly publish content in this namespace. A decentralized social network built in such a way becomes a virtual rendezvous point and set of publishing conventions, with processing and transformation implemented in local applications instead of a centralized service. Such a model would remove the control by a central service provider, and would improve privacy from data surveillance by distributing data ownership.

Note that both of the above examples describe mechanisms for verifiable, but not access-controlled, content. An open research challenge is how to best implement group access control for a particular application without resorting to IP-like channel or session semantics. The same key infrastructure used for verification in the example above could be used as the basis for encryption-based access control, though creating and managing group access is challenging, given the cryptographic tools currently available in practice. These challenges exist in other communications architectures as well, but NDN provides intrinsic mechanisms for naming, distributing, and verifying keys, which simplify the basic operations needed to deploy such access control.

4. Social Implications of NDN's Components and Departures

> The English language … becomes ugly and inaccurate because our thoughts are foolish, but the slovenliness of our language makes it easier to have foolish thoughts.
> - George Orwell, *Politics and the English Language*

Transitioning the Internet to NDN would produce a number of social changes by shifting the language used to envision and create networked applications. Some of these changes are difficult to predict; both protocol architects and deployers of Internet infrastructure purposefully provide

[6] The service's key is used for publishing in this case in order to enable applications to authenticate the content as coming from Facebook, and because the original content may have been transformed or altered, thus invalidating the user's original signature.

adaptable mechanisms, reconfigurable properties, and interpretive flexibility. However, changing the fundamental architectural components will change the nature of Internet interactions.

The following five sections expand on the use cases outlined above to illustrate how four important affordances of NDN could impact free speech (4.1), trust and security (4.2), privacy (4.3),regulation of content ownership or legality (4.4), and network neutrality (4.5). NDN's affordances for *semantic classification* have the potential to expose more information about the data that each packet contains. NDN's emphasis on *provenance* will change social models for identity, security and trust by requiring a publisher be identified for each Data packet. *Persistent publication* will shift the economics of caching and content distribution and, therefore, network neutrality, and will increase affordances for free speech while complicating information privacy. And *decentralized communication* emphasizes the ability to quickly and easily communicate without centralized servers or even wireline infrastructure, improving opportunities for free speech and privacy while frustrating law enforcement techniques that, according to recent accounts, have relied heavily on surveillance of traffic and storage within a few key application providers (Gellman & Soltani, 2013).

4.1 Free Speech

How a decentralized publishing model, like that in the social networking example, can support free speech becomes clearer when we consider a regime with authoritarian tendencies, which allows Internet access but constrains what is published. NDN makes it easier than IP to use alternative communications paths and opportunistic communication—NDN applications will be more likely to be able to communicate without global infrastructure. Users moving in cars or planes or people with ad-hoc wireless on their mobile devices can exchange data by leveraging in-network storage and persistent publishing namespaces even when they have intermittent connectivity. Any NDN node that has access to multiple networks – say, wireless and wired connections – can act as a bridge between those networks simply by forwarding Interests and/or answering them from its content store, broadening the scope of routes that data can take to a consumer (Zhang et al., 2010). NDN's data exchange mechanisms and affordances for decentralized communication make it more tolerant of disruptions by authority (such as those seen in Syria and Iran) than the client-server model of IP.

Because blocking a small number of well-known websites is currently an effective censorship scheme (Best & Wade, 2007), and in most cases, a website's content is centrally controlled, enabling decentralized communication can allow users to route around censorship, creating positive impacts for free speech. For example, NDN would enable a group of phones at a protest to use data muling — a combination of data storage and direct phone-to-phone communication in which phones carry video data from place to place rather than relying on infrastructure that might be subject to global surveillance. With NDN, the individually-signed packets of one video, carried by any number of devices to others, can be reassembled based on common naming conventions, and verified as being from the same publisher using data signatures. Such peer-to-peer muling can be done in IP networks, but is both more complicated at the network level (e.g., requiring IP address assignment even for local communication and providing only limited support for broadcast and multicast requests to local peers) and at the application layer (e.g., a higher-level data exchange protocol must include a mechanism to sign the data chunks and track authenticity of packets for reassembly).

4.2 Trust and Security

NDN requires all content be signed, whether produced by videographers or toasters; NDN applications can then verify the publisher of the data packets they receive. Though NDN cannot stop the use of false but similar-sounding names (NDN typosquatting), NDN's content signatures will increase consumers' recognition of, and reliance on, data provenance. For example, a user of NDN-Vimeo's video player will identify and verify each packet of a downloaded video by checking

the key used to sign every data segment. Including provenance information directly in packets will help to address consumer concerns of whether data has been compromised by either hosts or the communication channel over which the data was transmitted. Certificates, signatures, and provenance are more likely to be brought to the end-user's attention because of their intrinsic role in most NDN applications. Increased consumer recognition of, and reliance on, intrinsic provenance information would be a cultural change, not an architectural one, but we believe the architecture encourages the cultural change. Through this change, architecture's widespread deployment should improve data security and thus consumer trust in content, as well as mitigating some current Internet problems such as spoofing data and phishing.

Signatures alone are insufficient, of course, to determine whether the key used to sign data can be trusted. As in the NDN-Facebook example given above, trust models applied to Data packet verification must be generated for various classes of applications, which will require application developers, providers, and/or standards bodies to develop and communicate those models. Additionally, key/certificate distribution must be implemented as needed for various application classes and deployments.

While data encryption is not part of the basic NDN architecture, the inclusion of per-packet cryptographic signatures for provenance, and the resulting requirement for trust management and key/certification distribution mechanisms, will also provide most of the necessary prerequisites to support data packet encryption. The NDN architecture (philosophically, if not directly in the basic protocol) encourages applications to *secure the data* (and thus control access) by encrypting it, rather than attempting to secure the path over which the data flows as is currently done in the IP Internet using SSL, VPNs, and other similar schemes. Content producers will have the capability to easily encrypt sensitive packets and distribute keys via NDN. This reliance on encryption-based access control further emphasizes publication. After encrypted data is published, it can be duplicated many times and hosted in many (potentially hostile) locations, but only those with access to the right keys can decrypt the information.

However, NDN's reliance on widespread encryption to ensure access control raises usability challenges. NDN may make it easier for anyone to request a chunk of anyone else's encrypted data, and that data chunk is more likely to be cached with an unencrypted name in an NDN architecture than in the TCP/IP Internet. Ensuring privacy will require careful design of encryption mechanisms and incorporation of techniques, such as forwarding secrecy, that are now becoming increasingly common for security-conscious IP Internet applications. For example, encrypted NDN data may be widely available for long periods of time, increasing the long-term potential for attack and requiring the application of research in long-term encrypted storage. Further, encrypting data comes with tradeoffs, such as the computational burden of pervasive cryptography[7] and the challenges of key distribution and revocation.

To provide robust application support for these trust and security mechanisms, NDN requires new work on, for example, establishing, exchanging, and revoking keys within data-centric networks. Because application verification of NDN data is dependent on the trust models associated with those keys, standardized support for application-specific trust models is also critical. We hope such needs will stimulate development and improvement of practical encryption systems and their application to NDN. A particular opportunity for innovation is the need to develop flexible and easily-deployed mechanisms for encrypted group communications that fit the NDN communication

[7] Using unencrypted data internally is easier on resource-constrained devices, such as those found in IoT scenarios, and this approach is often used on IP networks. However, it is difficult to keep IP network perimeters secure, especially when there is a growing emphasis on Internet integration of diverse devices and systems, such as in the IoT vision.

paradigm. While basic examples exist, continued research is needed to provide usable, secure implementations of more complex multi-participant encryption schemes.

4.3 Privacy

Beyond security, NDN's defaults towards publication, provenance, and semantic classification of data create both problems and opportunities for information privacy. In particular, NDN departures from TCP/IP impact three fundamental information privacy issues: anonymity, data retention, and reasonable expectations of privacy. Privacy scholarship about the current Internet and digital technology more broadly expresses these concerns in terms of *information privacy*, generally understood as control over personal information (Waldo, Lin, & Millett, 2007); *contextual privacy*, which refers to limiting information flow to appropriate social contexts (Nissenbaum, 2009); and *individual privacy*, which can be protection from harms of exposure or invasion of personal space (Solove, 2010). NDN's changes to network communication impact each of these dimensions. In particular, NDN's request/response data exchange improves anonymous information-seeking (there is no source address in an Interest), but not anonymous publication (all Data packets have signatures). The architectural emphasis on publication and in-network storage presents new challenges for limiting *data retention*, and thus control over personal information, and will also likely change prevailing *expectations of privacy*, bringing them more in line with current US law.

4.3.1 Anonymity and obscurity

Anonymity - one traditional facet of information privacy – has both positive and negative social consequences. Anonymous communication can encourage free speech, help individuals evade censorship, and promote civic dialogue (Solove, 2010). Anonymity can also promote intellectual creativity and discovery (Cohen, 1996; Richards, 2013). Anonymity can also be used to evade prosecution for criminal behavior. And some scholars worry that there is a strong link between anonymity and mob behavior online, in particular hate crimes (Citron, 2010).

There are two issues for anonymity in NDN: obscuring who is seeking data (data consumers), and obscuring who is creating data (data producers). As described above, NDN strengthens the anonymity of data consumers. Though Interest packets create a trail as they are routed towards a Data packet, the entries in the PIT are erased as soon as a Data packet satisfies the Interest and each router's table only indicates the next hop. Though this trail of breadcrumbs could be logged, individuals requesting information are not likely to have their Interests traced back to them, unless an authoritarian regime can access and correlate state across all routers in the (possibly many) paths that individual Data packets have taken. ISPs might log Interests and forward them to governments, but decreased reliance on ISPs for connection due to NDN's possibilities for decentralized communication might enable users to circumvent such logging. Providing routes for anonymous data retrieval could strengthen privacy, allowing individuals to consume controversial political material or socially-stigmatized content without fear of embarrassment or harm.

NDN's impact on content producers is more complicated, since they may be identifiable in multiple ways, including by the key used to sign the data, by the namespace in which it is published, or by the content itself. While NDN data must be signed, it may be signed with either ephemeral keys or persistent keys unlinked to real-world identities. But the pervasive use of signatures makes it easier for infrastructure providers and content consumers alike to demand persistent or even verified real-world identities. An indicator of such a trend in the current Internet is Google's shift to requiring Google+ social network logins, which are tied to real-world identities, for participation in many of its platforms. In an NDN Internet, online forums might not accept comments without verified signatures. Data producers might also use multiple namespaces to enable pseudonymous or

anonymous communication, but the legal ownership of the namespace prefix may give away publishers' identities. Thus, because of NDN's emphases on semantic classification and provenance, more advanced technical measures will be required to secure anonymity for content producers in NDN. Researchers have explored Tor-like routing to preserve content producer anonymity (DiBenedetto, Gasti, Tsudik, & Uzun, 2012).

Similar to but distinct from anonymity is information *obscurity* (Hartzog & Stutzman, 2013), which individuals sometimes rely on (even for identified, non-anonymous data) as a form of privacy. Recent work has defined obscurity as a continuum based on a combination of four attributes: search visibility, unprotected access, identification, and clarity (Hartzog & Stutzman, 2013). Empirical research has demonstrated that individuals believe that obscurity can help protect their personal information on the Internet (boyd & Marwick, 2011). NDN may have impacts for the third attribute of obscurity in particular: identification. Common use of unencrypted and semantically meaningful names, as well as the need to authenticate data with a persistent identity, will tend to reduce the obscurity of information by default.

However, NDN may eventually strengthen information obscurity in much the same way it encourages free speech – by encouraging decentralized, rather than cloud-based, communications. IoT applications can use data within local or private networks, or encrypted namespaces, rather than globally-available and/or plaintext namespaces. Private namespaces will limit the likelihood of centralized data collection and surveillance, a privacy concern as cars and household devices gain the ability to monitor and report our activities (Schneier, 2013). Many social impacts of the NDN architecture will depend upon how NDN applications are designed and implemented .

4.3.2 Data retention and forgetting

Recently, international privacy scholars as well as policymakers in Europe have paid increased attention to data retention and disposal, or what has been termed the "right to be forgotten" (Blanchette & Johnson, 2002; Mayer-Schoenberger, 2007; Rosen, 2012). More recently, California adopted Senate Bill 568, which requires websites to enable minors to easily remove their own posts from websites. NDN's defaults towards persistent publication will further complicate how society addresses the social role of forgetting and data retention more generally, especially as it facilitates applications that adopt data distribution models over connection-oriented models. As personal data proliferates on the web, policymakers are increasingly concerned that such data cannot be erased or forgotten. The specter of total accountability for our past actions is considered unpleasant at best and potentially limiting to social interaction and democracy at worst (Blanchette & Johnson, 2002; Mayer-Schoenberger, 2007).

Today's routers purge data from buffers as soon as it is passed on towards the requesting party, or they never store it. In contrast, NDN nodes cache incoming data for future requests and may use repos for more persistent storage of data. NDN routers default towards remembering (caches); IP routers default towards forgetting (buffers). With an IP Internet, parties can request that publishers remove data from the hosting site. Although copies may proliferate in caches and on end hosts on the network, new requests for the original hosted data will go unsatisfied. In NDN, copies will proliferate on routers, repositories, and other application-specific stores and remain accessible in response to Interests, due to the caching model at the heart of the architecture. NDN creates an even more memory-intensive model of the Internet, which will require protocols that include information to define and respect expectations for data deletion.

4.3.3 Network surveillance and reasonable expectations of privacy

In the long term, our expectations of privacy may also evolve more fundamentally based on our communication capabilities. Today's IP internet suffers from a critical disconnect between user expectations for private communication (Urban, Hoofnagle, & Li, 2012) and legal protections. For

example, some consumers may expect email conversations to be private because they are typically shared with only a few individuals, but most email, especially if hosted in the cloud, is not secret in any significant technical sense. U.S. courts have ruled that the reasonable expectation of privacy usually applied to telephone conversations (*Katz vs. United States*, 1967) does not apply to the Internet, because communications on the Internet are shared with many third parties (Glancy, 2000). The Supreme Court has ruled that a reasonable expectation of privacy disappears once a communication is handed off to a third party. These rulings have resulted in a situation where courts consider various sorts of communications to be essentially public, while users envision an Internet that shares some regulatory protections afforded telephone communications. Will NDN networks bring social expectations of the network in alignment with the legal interpretation? NDN makes no guarantee of privacy for published information—applications must encrypt their data.

4.4 Content Regulation

Because the Internet is widely used for commerce across international borders, it must contend with diverse national and international policies regulating publication and use of content. Some content types may be illegal in some countries (for example, sale of Nazi memorabilia in France); other forms of content may have use restrictions designed to guarantee a profit to content creators (for example, movies produced by major studios). Enforcing publication and use regulations on content across the global Internet is an intractable problem with today's IP Internet. Corporate interests often use the loose geography of IP addresses to enforce market-based restrictions on content access. Law enforcement uses a range of tactics – ranging from IP address tracing to deep packet inspection – to track and prosecute both producers and consumers of illegal or pirated content. A transition to NDN will impact each of these mechanisms of law enforcement by changing the tools needed for tracking individuals and monitoring and restricting communications.

4.4.1 Law Enforcement

NDN's emphasis on semantic names and required content signatures may make certain types of law enforcement easier, as the source of much Internet data will be readily traceable. For example, if NDN conventions evolve so that data names reflect data types, application-specific names may make application-level packet-sniffing (and therefore, evidence-seeking) more efficient and less processing-intensive. Criminals are not likely to give identifying names to illegal material, regardless of common practice in NDN. But those involved in illegal activity will have to have access to a namespace, and follow conventions that enable their files to be routable on the network. Law enforcement will likely be able to trace criminal activity to namespaces much as they would to IP addresses in today's Internet.

NDN's emphasis on publication may trigger a social shift towards encrypting more data. Police and regulatory regimes have long been wary of cryptography, as developers have resisted providing back doors for law enforcement to inspect or wiretap communications. NDN's reliance on cryptography with decentralized trust schemes could face similar resistance from law enforcement as well as operators; encrypted traffic makes wiretapping, deep packet inspection, and traffic management more difficult (Bendrath & Mueller, 2011).

NDN will also necessitate a change in how governments currently assert regional jurisdiction on the Internet. IP addresses are often used to determine who to target in a law enforcement action (Cooke, 2007); IP source address spoofing reduces the effectiveness of surveillance techniques that require source identification. NDN further disassociates addressing from location, which might dissuade law enforcement from identifying the subjects of actions based on network data. NDN makes local geography more difficult to track, which means that law enforcement must rely on other evidence such as credit card and financial trails for enforcement. While it may temporarily

complicate policing regional jurisdictions, NDN might provide a lever to encourage law enforcement methods that are more accurate and effective.

Finally, NDN's strengthening of anonymity for content consumers may bring changes to how crimes are prosecuted. It may become easier to pursue the producers of illegal or infringing information, rather than the consumers. Although some may argue that consumers of pirated or illegal material should be punished as well, we argue that eliminating the source of infringing material is a more fair and effective mechanism than punishing consumers (Cohen, 1996; Gillespie, 2009).

4.4.2 Copyright enforcement

Because NDN retrieves data by name, rather than by host, and encourages widespread storage of content chunks that are not authored by the storage owner, the concept of hosting content is weaker than in IP. This evolution implies changes for how copyright law is enforced, including both prevention of infringement and allowed fair uses. One open question is how content producers can facilitate digital rights management (DRM) in an NDN world. DRM attempts to provide strict enforcement for copyright holders and, in some cases, content identification as in YouTube's Content ID system, while limiting rights to consumers or libraries (Cohen, 2003). DRM typically controls distribution of content, including whether consumers may redistribute content. NDN supports the first kind of DRM well, but the second kind poorly.

NDN supports control over distribution of copyrighted content well. Just as in TCP/IP, copyright holders can easily distribute verified, encrypted media, and consumers would access the content with the proper key. (Setting aside the challenges of group communication described earlier, this approach could follow current DRM strategies in the IP Internet.) Producers might allow fair use by giving copies of keys to libraries, or by providing portions of the content in the clear for scholarship, critique, parody, or other protected fair uses.

But once consumers have received and decrypted verified content, they may distribute unauthorized versions in clear text. Content industries may object to NDN's default caching, because so many copies of both licensed (presumably encrypted) and pirated (presumably decrypted) media can reside on countless routers and repos. A world where countless copies proliferate across the Internet would challenge a major US mechanism of copyright enforcement, the Digital Millennium Copyright Act (DMCA) takedown notice. On the TCP/IP Internet, for scalability, a video either must be hosted by a major company such as YouTube or Hulu, which attempt to accommodate organizations wielding takedown notices. Who would take down an infringing video that is duplicated on routers across the world? Automatic caching of content is less of an issue, since routers will likely cache data only briefly. A second type of duplication — curated replication to repos operated by legally responsible entities — is more likely the recipient of future takedown notices. The political economy of repos – who owns them, and in what legal jurisdictions – will likely impact the future efficacy of DMCA takedown notices.

4.4.3 Geographic content controls

Law enforcement personnel are not the only stakeholders that rely on the loose geography provided by IP addresses for content control. Major sports franchises restrict subscribers in local markets from watching games online. Gambling operations restrict participation from countries in which such operations are illegal. Search results might be tailored to a searcher's location. All of these industries will need to look for alternative ways to enforce location-based content restrictions. Because Interests can come from anywhere, a system of encryption and key distribution based on location-verified subscribers will likely be the result. Encrypting content for a single subscriber loses much of the economy of scale provided by in-network caching and storage, creating tradeoffs for NDN application developers that wish to restrict access to content.

4.5 Network neutrality

The network neutrality debate focuses on what actors pay for Internet resources such as routers and bandwidth, and whether those actors (in IP, traditionally ISPs) can throttle or privilege traffic to increase revenue. In NDN, the network of actors controlling traffic routing decisions is likely to be much broader, expanding the scope of network neutrality. In NDN routers, strategy algorithms controlling the operation of three routing tables – the content store, the PIT, and the FIB – may impact network neutrality by enabling the router owner to express particular traffic shaping choices in terms of NDN namespaces. These routers may be owned by ISPs, but they may also be owned by individuals, small groups, other corporations, or governments.

An NDN router's FIB (forwarding information base) is roughly similar to the FIB in an IP router, except that it contains name prefixes instead of IP address prefixes, and it may show multiple interfaces for a given name prefix. Routing protocols and/or manual setup of static routes are used to configure the FIB; the resulting configuration expresses the policies of the router's administrators. For example, routing administrators may choose to discriminate based on data types (indicated within clear-text data names) or based upon data's namespace of publication. While similar possibilities exist in IP, they are at higher layers; NDN routers will be capable of such choices at the network layer and at wire speed.

Data that passes through an NDN router is stored in the content store. Though initially conceived of as a cache, the content store can be extended into persistent storage for expanded caching based on the business logic of the router owner. NDN's encouragement of content stores on each router will disrupt the current market for content distribution networks (CDNs[8]) and hosting services. NDN could introduce competition for current CDNs by spreading out caching and its costs. In a similar fashion, NDN will impact the economics of content hosting. Content producers on an NDN Internet can use a cheap server and low-bandwidth connection, and their viral videos can still be reachable on their own server, with the network scalable serving content requests. NDN will reduce dependence on third-party services, while allowing such services to continue to provide added value. Users can host content on their own terms (or those of their ISP), rather than being subject to a third party CDN provider or hosting service's terms. But hosting services and CDNs won't disappear, as there will still be a need for persistent storage in NDN. Very little content is sufficiently popular to be constantly serviced through content store caches. So today's CDN companies may go into the business of running repos to provide this storage, for a fee. Content producers might pay for longer-term storage. Or content store owners might cache more popular content over less popular content. ISPs may take on greater responsibility for providing caching resources in NDN, giving ISPs access to a market currently dominated by CDNs.

The PIT is a fundamentally new entity that does not exist in an IP router. PIT entries in an NDN node record the Interest packets that have been forwarded and await for Data packets to return. An entry records the requested data name, the incoming interface(s) of the Interest(s), and the outgoing interface(s) to which the Interest has been forwarded. Policies that set how long Interest information is retained could impact retrieval performance. Whether consumers or namespace providers are able to pay for better quality of service through longer Interest storage in the PIT or, for example, more aggressive re-issuing of Interests across multiple outgoing interfaces is a strategy question that may impact the neutrality of the node.

Prior to being stored in the PIT, arriving Interests are immediately forwarded, also according to a strategy module. As currently envisioned, policies for the strategy module are relatively neutral,

[8] CDNs are companies that replicate data across a geographically distributed network connected to the IP Internet, moving content close to urban centers and thus providing faster data access over a broader area (often globally) than a traditional web hosting model, for a fee.

dependent upon network conditions or link costs and not data type or producer. However, we can imagine ISPs that would author their own strategy modules to prioritize certain types of data or data namespace origins. NDN pushes away from the concept of sessions, which renders the utility functions currently used for congestion management obsolete. This obsolescence gives us an opportunity to rethink what fairness might mean for traffic management in a world running on NDN. What will fair congestion management look like if semantically-rich names enable data to easily be segmented by type or namespace of origin? Data names may reveal types of content (such as video, VOIP, scientific data, or emergency response data); keys may reveal even more information about origin. Evolving standards might require quality of service requests to be included in names – for example, inelastic VOIP data could be named as such for quick delivery. Perhaps routing algorithms could help prioritize less-popular information (such as emergency response traffic) to avoid tyranny of the majority. However, such algorithms would require a system to delineate and charge for trusted quality of service. For scalability, NDN network administrators could use well-known name prefixes or components to configure resources.

So, a more likely (but perhaps less equitable) marker for quality of service information would be namespace of origin, as most data names will express an originating namespace. Namespace of origin would be a poor indicator of data elasticity, but a good indicator of the power and status of the originating institution. For instance, the network could route emergency response data quickly using algorithms sensitive to namespace of origin. Networks could similarly prioritize today's *New York Times* or Netflix videos. But a system of traffic management based on institutional providers raises many of the political concerns addressed in the network neutrality literature (Peha, 2007) by increasing the power of institutions (and their data) relative to individuals. Semantically-meaningful names may facilitate such data priority distinctions.

Beyond router policy, NDN's support for mobility and disruption-tolerant networking will impact network neutrality. Even if prioritized networking evolves using semantically meaningful names or pay-for-retention policies on routers, NDN's ability to route around ISPs will give consumers more options for data transmission, empowering users.

5. Similarities with TCP/IP

Though there are many technical differences between TCP/IP and NDN, there are also fundamental ways that NDN does not depart from TCP/IP. This section discusses social and political issues that are unlikely to change in a World on NDN.

5.1 Top-Level Name Allocation

NDN does not change the need for top-level, globally-unique names for globally-routable information. Mechanisms for global namespace governance will be necessary, and top-level names will likely be associated with real people and legal entities. If a central namespace allocator authoritatively knows who owns what name prefixes, then law enforcement can use that namespace allocator's records to determine the subjects of their actions in NDN just as in IP. If names are chosen by publishers without a central allocator, then it is possible that names will be useless for identifying particular subjects. But such a network would raise name squatting and trademark issues, which governments would most likely seek to prevent, most likely through establishment of a central namespace allocator.

5.2 Censorship

Though NDN supports geographically local, decentralized content (such as with a group of phones at a protest), broadcasting censored content more widely over public networks may not be any easier than with IP. Someone on the network must ask for content for it to traverse an NDN-based

Internet. A WikiLeaks-style organization, an opposition movement, or any number of private citizens might express Interests for content published by unpopular or even unknown authors. But in what namespace would content authors publish controversial or censored content? Well-known namespaces for controversial material and content publishing keys associated with organizations like WikiLeaks could be blocked, as could anonymizers. So long as the regime allows arbitrarily named, encrypted content to be published, consumers and producers could share content with impunity using methods similar to those on the TCP/IP Internet (e.g., passing keys out-of-band).

In addition, NDN's data exchange mechanisms do not allow for regimes to block certain types of traffic via port number. However, regimes or network authorities could still block or slow traffic with certain kinds of names. If data names evolve to include application type (likely for both prioritizing inelastic data for network management, as well to enable servers to direct packets to the proper application face), such application-specific names may also enable blocking types of traffic.

If a regime blocks specific types of traffic or bans encrypted content outright, NDN could support a variety of countermeasures. Steganography might be used to embed messages in normal content. The rich header fields of both Interests and ContentObjects might be used by either Interest or content producers to relay secret content. (For example, an application could specify the hash of the content in messages.)

5.3 Markets

Whether NDN will fundamentally reshape the role of either ISPs or content owners is difficult to predict, though many existing market forces would likely persist. Though NDN enables more networking around the edges—decentralized networking support is part of NDN's architectural vision—the political economy of the industry makes it likely that most routing and storage infrastructure will remain owned by ISPs and paid for by consumer subscription. NDN is also unlikely to change models of content ownership. Though data producers will find their content to be much more dispersed across the web, they could still control access through encryption.

6. Openings for Policy

As the ways in which NDN is similar to TCP/IP suggest, NDN cannot solve all of today's Internet challenges on its own. To flourish, NDN will likely need a set of policy regulations that evolve alongside the architecture.

If content is to be widely distributed and cached over a variety of personal devices, we will need to define ownership and legal jurisdiction for pervasive in-network storage. For example, we must resolve whether individuals should be accountable for the content on a given device. If illegal material is cached by your phone as it makes its way to another consumer, should you be responsible for that content? Current prosecution is based on whether that material is found on your machines. New legislation will need to protect device owners from unrequested content if a distributed, peer-to-peer model of distribution is to thrive.

Policy will also need to define fair congestion management policies when semantically-rich names are widely used. If we wish to restrict the ways that ISPs can discriminate based upon names, legislation to support this most likely will have to be created.

Finally, legislation will need to define next-generation DRM and intellectual property in an NDN world. How will we enforce fair use in a system where content must be encrypted in order to be controlled? How can we prevent unencrypted copies of intellectual property from circulating broadly? Such questions must be addressed by policy rather than network architecture.

7. Conclusion

NDN brings the semantics of the current Internet's data-centric application layers to the network layer. In doing so, it may provide significant benefits to applications and network operators alike. We have explored the possible social and cultural impacts of both layers, including some of today's most pressing challenges: free speech, security and privacy, control of content, and network neutrality. NDN departs from TCP/IP in its emphasis on data publication, data provenance, decentralized communication, and semantic classification. Analyzing these departures has illustrated that NDN is likely to improve conditions for free speech and security, while complicating both privacy and content regulation. NDN's overall impact on network neutrality remains an open question, dependent upon choices still to be made in naming and routing.

The practical impact of NDN will also depend on how a number of open research areas are addressed, specifically how to: (1) balance application-driven semantically meaningful, consistent names that simplify application development and opaque names that better protect privacy; (2) develop practices for key assignment, distribution and revocation, given NDN's reliance upon content signatures for identity and security; (3) provide usable, secure implementations of more complex multi-participant encryption schemes; (4) standardize mechanisms for establishing trust relationships; (5) mitigate information leakage in names; and (6) create fair congestion management when semantically-rich names are widely used.

Indeed, most of NDN's potential changes for free speech, security and privacy, content regulation and law enforcement, and network neutrality are speculative, as the NDN architecture continues to evolve as these questions are explored. But imagining the social changes NDN might encourage is a useful exercise in relating infrastructure, social challenges, and impacts. We hope this work will spark continuing discussion of the current Internet's impact on society. Thinking creatively about how network usage has changed helps us reimagine the relationship between infrastructure and our world.

Acknowledgements
The authors wish to thank colleagues Van Jacobson, David D. Clark, and Steven Bellovin for feedback and ideas that shaped this work.

8. References

Bendrath, R., & Mueller, M. (2011). The end of the net as we know it? Deep packet inspection and internet governance. *New Media & Society*, *13*(7), 1142–1160.

Best, M. L., & Wade, K. W. (2007). Democratic and Anti-Democratic Regulators of the Internet: A Framework. *The Information Society*, *23*(5), 405–411. doi:10.1080/01972240701575684

Blanchette, J.-F., & Johnson, D. G. (2002). Data retention and the panoptic society: the social benefits of forgetfulness. *The Information Society*, *18*(33-45).

boyd, danah, & Marwick, A. E. (2011). Social Privacy in Networked Publics: Teens' Attitudes, Practices, and Strategies. In *A Decade in Internet Time: Symposium on the Dynamics of the Internet and Society*. Oxford, UK: SSRN. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1925128

Citron, D. K. (2010). Civil rights in our information age. In *The offensive internet: privacy, speech, and reputation* (pp. 31–49). Cambridge, MA and London: Harvard University Press.

Cohen, J. E. (1996). A Right to Read Anonymously: A Closer Look at "Copyright Management" in Cyberspace. *Connecticut Law Review*, *28*, 981–1039.

Cohen, J. E. (2003). DRM and Privacy. *Berkeley Technology Law Journal*, *18*(2), 575–617.

Cooke, L. (2007). Controlling the net: European approaches to content and access regulation. *Journal of Information Science*, *33*(3), 360–376. doi:10.1177/0165551506072163

DeNardis, L. (2012). Hidden levers of internet control. *Information, Communication & Society*, *15*(5), 720–738. doi:10.1080/1369118X.2012.659199

DiBenedetto, S., Gasti, P., Tsudik, G., & Uzun, E. (2012). ANDaNA: Anonymous Named Data Networking Application. In *19th Annual Network & Distributed System Security Symposium*. San Diego, CA: Internet Society. Retrieved from http://arxiv.org/abs/1112.2205

Gellman, B., & Soltani, A. (2013, November 1). NSA infiltrates links to Yahoo, Google data centers worldwide, Snowden documents say. *The Washington Post*. Retrieved from http://www.washingtonpost.com/world/national-security/nsa-infiltrates-links-to-yahoo-google-data-centers-worldwide-snowden-documents-say/2013/10/30/e51d661e-4166-11e3-8b74-d89d714ca4dd_story.html

Gillespie, T. (2009). *Wired Shut: Copyright and the Shape of Digital Culture* (Reprint.). The MIT Press.

Glancy, D. (2000). At the Intersection of Visible and Invisible Worlds: United States Privacy Law and the Internet. *Santa Clara Computer and High-Technology Law Journal*, *16*, 357.

Hartzog, W., & Stutzman, F. (2013). The Case for Online Obscurity. *California Law Review*, *101*(1), 1–49.

Jacobson, V., Smetters, D. K., Thornton, J. D., Plass, M. F., Briggs, N. H., & Braynard, R. L. (2009). Networking named content. *Proceedings of the 5th international conference on Emerging networking experiments and technologies*, 1–12.

Jacobson, V., Braynard, R. L., Diebert, T., Mahadevan, P., Mosko, M., Briggs, N. H., Barber, S., Plass, M. F., Solis, I., Uzun, E., Lee, B.-J., Jang, M.-W., Byun, D., Smetters, D. K., and Thornton, J. D. (2012) "Custodian-Based Information Sharing," *IEEE Communications Magazine,* July, 2012.

Katz vs. United States. , No. 389 U.S. 347 (1967).

Lessig, L. (2006). *Code: version 2.0*. New York: Basic Books.

Mayer-Schoenberger, V. (2007). *Useful void: the art of forgetting in the age of ubiquitous computing* (Working Paper No. RWP07-022). Cambridge, MA: Harvard University.

Nissenbaum, H. (2009). *Privacy in context: technology, policy, and the integrity of social life*. Stanford, CA: Stanford Law Books.

Peha, J. M. (2007). The benefits and risks of mandating network neutrality, and the quest for a balanced policy. *International Journal of Communication*, *1*, 644–668.

Richards, N. M. (2013). The Perils of Social Reading. *Georgetown Law Journal*, *101*(3). Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2031307

Rosen, J. (2012). The Right to Be Forgotten. *Stanford Law Review Online*, *64*, 88.

Schneier, B. (2013, May 21). Schneier on Security: Surveillance and the Internet of Things. Retrieved June 5, 2013, from https://www.schneier.com/blog/archives/2013/05/the_eyes_and_ea.html

Solove, D. J. (2010). *Understanding Privacy*. Harvard University Press.

Star, S. L. (1999). The ethnography of infrastructure. *American Behavioral Scientist*, *43*(3), 377–391.

Urban, J. M., Hoofnagle, C. J., & Li, S. (2012). *Mobile Phones and Privacy*. Berkeley, CA: University of California at Berkeley - Center for the Study of Law and Society. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2103405

Waldo, J., Lin, H. S., & Millett, L. I. (2007). *Engaging privacy and information technology in a digital age*. Washington, D.C.: The National Academies Press.

Zhang, L., Estrin, D., Burke, J., Jacobson, V., Thornton, J. D., Smetters, D. K., … Yeh, E. (2010). *Named Data Networking (NDN) project* (PARC Technical Report No. NDN-0001). Palo Alto, CA: PARC.