

Sustainable Internet Architecture

PROJECT DESCRIPTION

1 Introduction

The Internet currently plays an important role in our daily lives and its impact continues to grow. But at the same time, it faces a number of challenges. Numerous research studies on a new Internet architecture (e.g., [16, 37, 48, 54, 55]) have addressed these challenges from different perspectives. Yet, researchers still do not have a consistent view on both the new Internet architecture itself and the underlying problems that motivate it. We look at the Internet architecture and its underlying problems from a higher-level perspective, considering Internet development not only within the *technical* domain, but the *economic* and *social* domains as well. From this new perspective, a clear-cut and consistent picture of the underlying problems shows up. Based on that, we propose a solution — a framework for a new, generalized Internet architecture that addresses both long and short term development goals. We will design, prototype, and evaluate core systems that serve as the basis of this framework.

Analyzing Internet development across the technical, economic and social domains, we find that the problems in the current Internet architecture stem from its lack of *sustainability* which impedes future development. This sustainability problem shows in two aspects: (i) *evolvability* issues and (ii) the Internet's pressure on a set of *cultural norms* (e.g., cooperation, trust, creativity, economic and social order) that its development relies on. This is similar to the central problem that traditional sustainable development tries to solve: economy development imposes pressure on the natural and social environments but also relies on them.

On the one hand, the current Internet does not have sufficient evolvability to keep up with new demands of services that run on it. Evolution of the Internet is driven by these demands and is expected to meet them. For example: (i) There is no effective countermeasures against *distributed denial of service* (DDoS) attacks, which prevents many good service ideas from being applied in practice. (ii) There is a long anticipated demand for *quality of service* (QoS), but it is still far from being met due to architecture restraints.

On the other hand, the Internet imposes increased pressure on the social environment, *i.e.*, a set of cultural norms upon which its development depends. We argue that the current Internet architecture does not provide leverage points, by which we can direct Internet services to evolve in a sustainable way. In particular, we should be able to make Internet services improve those cultural norms, or at least not deteriorate the norms if unable to improve them. For example: (i) Global cooperation among service providers is an essential element for Internet development, but a design dogma today is that we are better off not assuming service providers' willingness to cooperate. Can we find any leverage points to cultivate global cooperation? (ii) Copyright issue is another example of the norms. Copyright infringing objects (video, audio, software) exist in large quantities on the Internet yet there is no effective countermeasure. As a side effect, content pollution becomes justified in many cases, although it is annoying most of the time. (iii) The Internet's pressure on economic and social order as well as human mental experience keeps increasing. As an example, Internet addiction becomes a non-trivial social issue.¹

We argue that we can address both aspects by adding proper *controllability* and a *semantics aware property* to the Internet architecture. The controllability does not restrain freedoms; it instead fosters a higher extent of freedom resulting from advanced flexibility and functionality provided by the new architecture. A semantics aware property makes the entire architecture (from the highest to the lowest layer) evolvable based on service level semantics. To support the controllability and semantics aware property, we introduce four leverage points: *identity*, *standards*, *incentives*, and *auditing*. The four leverage points work in synergy, forming the base of our framework of a new, generalized Internet architecture.

Our framework adapts and applies basic principles of the traditional sustainable development [26]. We find not only that these principles are well suited in the Internet context but also they tend to work better in this new context. Most surprisingly, applying them to the Internet context can potentially create priceless products that lead to fundamental innovations for the entire technical, economic, and social development. A typical example of such products is the improvement of common human values and social trusts, which

¹According to [4], more than one out of eight Americans exhibits signs of Internet addiction. The biggest culprit of Internet addiction is not online pornography, games, and gambling as some people think, but ordinary services such as email, online chatting, and shopping.

is the key to solve many problems that are currently unsolvable by technical means.

We will deploy our research at two levels. At one level, we will perform in-depth evaluation of our framework and principles that will be introduced in Section 3. We will evaluate our theories by comparing with related work across computer science, economics, sociology, and law. The goal is to show in detail how our solutions can satisfy or reconcile with related design criteria, how our systems can potentially solve currently unsolvable problems across technical, economic, and social domains. At another level, we will design and prototype four core systems that are basic for the proposed framework:

1. A *distributed announcement database* system, which is a fundamental infrastructure that makes possible many other systems in our framework. This system is designed to announce authoritative information to public. It is a key to implement a scalable identity authentication architecture that we propose.
2. A *global platform* that facilitates *network performance auditing*. This platform provides a common infrastructure that collects network auditing information from different administrative domains, processes it, solves disputes, and disseminates results to proper targets (including the public).
3. A *semantics aware network* prototype, which explores the proposed idea of introducing a semantics aware property to the Internet architecture in depth. This prototype addresses concrete design considerations of capturing and processing application level semantics at network layer in an efficient way.
4. A *unified incentive system*, which is essential to address incentives for individual users to cooperate and apply sustainable practices. This system maintains a profile for each user based on globally agreed standards and binds the profile to the user’s permanent identity. The system also ties all Internet services together to form unified incentives such that it maximizes the effectiveness of incentive mechanisms.

How we fit into a larger overall architectural framework. While unique in its goals, methods, and, we hope, significance, our project is *not* and isolated effort detached from the rest of the networking community and FIND projects: our identity-based-capability solution relates to capabilities as well as host identity and address schemes. Unlike other capability solutions, *e.g.*, [51, 55], our solution does not rely on any routing- or link-level assumptions. It works totally independent of routing, thereby allowing maximum flexibility for routing architectures and address schemes. Our solution adopts *permanent identity* instead of using temporary ones as most host identity schemes do. Moreover, our solution ties all services together with a unified identity. These two distinct features make our identity solution not only work well for capability, but also capable of serving as an effective threat against misbehavior, due to fear of consequences.

Broader impact. This research will make significant and fundamental contributions not only within technical domain surrounding the Internet, but also on the entire technical, economic, and social development. Within the technical domain, we are bringing about innovative insights on applying principles and experience of the traditional sustainable development to solve the Internet’s fundamental problems. Moving beyond the technical domain, our research will expose the great feasibility to exploit technical means of the Internet to foster *non-technical solutions* (that demand change in human values or ideas of morality), which are keys to a number of currently unsolvable problems across technical, economic, and social domains.

2 Background

2.1 Sustainable Development

Our research learns and applies basic ideas and experiences from the *traditional* sustainable development of the last three decades [26]. We find most principles in that area can be directly applied to the Internet context. Moreover, these principles tend to work better in this new context. In this section, we introduce the background of sustainable development. To avoid ambiguities, we use the abbreviation *SusEco* to denote the traditional sustainable development (that relates to economy and ecosystem), and use the term *SusInet* to denote sustainable development in the Internet context.

2.1.1 Basic Principles

It is increasingly recognized that we need to achieve *sustainable development* — development that not only improves economic goals, but also advances *social* and *environmental* well beings simultaneously. Until the 1980s, the overwhelming opinion was that there were inevitable and fundamental trade-offs among the three. However, progress in the *SusEco* domain in the last three decades has shown the feasibility to achieve a win-win-win goal by better balancing the short- and long-term needs and government leadership.

The following are five well known *SusEco* principles: 1. “*Business is good for sustainable development and*

sustainable development is good for business.” Business is a part of the sustainable development solution, while sustainable development is an effective long-term business growth strategy. 2. “*Good governance is needed to make business a part of the solution.*” Good governance provides solution to conflicts arising from the interaction between the short-term pressure induced by businesses’ financial goals and the emerging principles of sustainable development. 3. “*Access to markets for all supports sustainable development.*” Sustainable development is best achieved through open, transparent, and competitive global markets. 4. “*Cooperation beats confrontation.*” Sustainable development challenges are huge and require contributions from all parties — governments, businesses, civil societies, and international bodies. Confrontation puts the solutions at risk. Cooperation and creative partnerships foster sustainable development. 5. “*Thinking locally, acting globally.*” While there is much we can do locally, action is also needed at the global level. There is an inevitable need for nations to collaborate to solve common problems.

The concept of SusInet is very similar to SusEco if we treat the set of culture norms that Internet development affects (but also relies on) as the counterpart of the social and environmental well beings in SusEco. We can apply all the five SusEco principles to SusInet. Principle 1 points out *basic incentives and necessity* for ISPs and IT corporations to conduct SusInet practices. Principle 2 validates our approach to introduce *government involvement* into our solution. Principle 3 is the reason why we should emphasize the *market-based approach* for the government’s role in guiding businesses. Principle 4 corresponds to a general solution to SusInet, the *Pyramid process* (Section 3.4). Principle 5 is straightforward due to Internet’s global nature. However, in this research we will focus on *local acts* of nations.

2.2 Distinct Features of the Internet

The Internet has three distinct features: (i) *Global nature*. Internet provides global connection all over the world. (ii) *Fast service evolution*. Services on the Internet evolve much faster than the traditional business. (iii) *Fast information dissemination speed*. The Internet provides the fastest way to disseminate information — in particular, information that can attract interest of the majority. These distinct features create a specific environment which differs a lot from that of the traditional business. Such differences provide both extra challenges and opportunities to conduct sustainable development in the context of the Internet. We argue that the opportunities surpass the challenges.

Challenges. One of the largest challenges for the Internet is the difficulty to *enforce the law*, which complicates the enforcement of the social norms and order that the Internet needs. This is due to both the *complex nature* and the *fast evolution pace* of Internet services, which make cyber-law enactment unable to keep up with. Even if the law enactment were responsive, its efficacy could be limited if the majority, or at least a large percent of users, violate the law. Meanwhile, it could take a long time to popularize new laws such that the majority comply with them. Another major challenge is that Internet based businesses face more severe *short-term pressure* than traditional businesses, and such pressure could *last long*. This is due to both the more fierce competition resulting from the relaxed geographic restrictions and a higher probability that “early birds” can form a long-term global market dominance.

Opportunities. However, the Internet (in particular its fast information dissemination speed and global nature) provides good opportunities to address the above challenges. In addition, the Internet can significantly facilitate the implementation of the basic sustainable development principles. For example: (i) We can *accelerate the popularization* of laws and standards via passive and active learning mechanisms. (ii) We can enforce laws and standards globally (within each nation) by adding new features (Section 3.3) to the Internet. (iii) We can build an effective *unified incentive system* on the globally connected Internet. (iv) We can facilitate a general solution to the sustainable development (both SusEco and SusInet contexts), *i.e.*, *Pyramid process* (Section 3.4), by building an Internet platform to foster this process.

3 Framework and Principles

Here we describe our framework and design principles of a new Internet architecture towards the sustainable development goal. At a high level, we address sustainability by adding proper *controllability* and a *semantics aware property* to the Internet architecture. We define two types of controllability: (i) *High level controllability*, which is the controllability of the sustainable development, *i.e.*, the ability to direct Internet development to follow sustainable approach. (ii) *Low level controllability*, which is the controllability of the Internet access. The current Internet is an *access-by-default* network, *i.e.*, access is permitted by default. We address the low level controllability by adding a global *access-on-request* network, *i.e.*, access is denied by

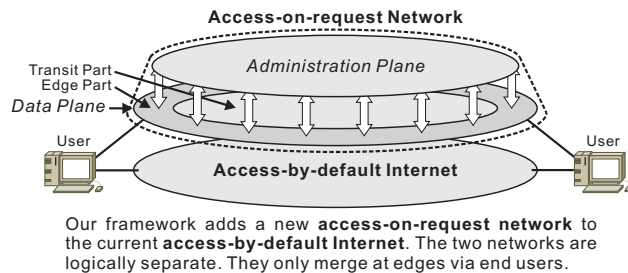


Figure 1: Abstract Structure of the New Internet

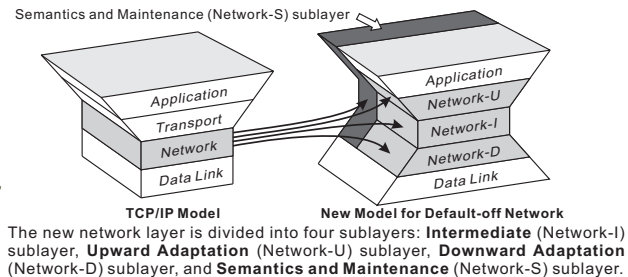


Figure 2: New Network Layer

default to the Internet. Complementing efforts of others, *e.g.*, [13,51,55], our key contribution is a generalized approach to deploy such a system at a *nationwide scale*. The semantics aware property means that all evolvable parts (that evolve as a result of new service level semantics) of the network architecture are able to capture related service level semantics. We achieve both the low level controllability and the semantics aware property by introducing a new network model (Section 3.1).

The high level controllability is the core of our framework. We address it by introducing four critical leverage points: *identity* (Section 3.2), *standards*, *incentives*, and *auditing* (Section 3.3). The four leverage points work in synergy, forming the base of our framework. *Identity* is a key to implementing our access-on-request network and is also the premise to build a *unified incentive system*, in which we bind extensive information of a user to her identity. *Standards* exploit the semantics aware property and are the key to directing Internet evolution. We recognize that both enforcement and enactment of new standards are non-trivial tasks. We exploit *incentive* mechanisms and *auditing* to help enforce standards. And we use *identity* as an ultimate weapon of enforcement. For standards enactment, we take advantage of the *Pyramid process*.

3.1 Network Model

Abstract Structure of the New Internet. Figure 1 depicts the structure of the new Internet architecture that we propose. As shown in the figure, we are adding an access-on-request network to complement today’s access-by-default Internet. The access-on-request and access-by-default networks are logically separate networks which only merge at edges via end users (end users can choose to use a service on the new access-on-request network, or to use a counterpart service on the old access-by-default network). The access-on-request network consists of two separate planes: the *data plane* and the *administration plane*. The data plane is where user traffic is actually transferred. The administration plane provides *centralized administration* (backed by a single trusted third party — government) for the data plane. The data plane is further divided into the *edge part* and the *transit part*. The edge part performs access control and processes semantics by contacting the administration plane. The transit part is responsible for routing. Although it can contact the administration plane as well, our access-on-request solution does not rely on such communication. This allows the maximum flexibility for routing architecture deployment (any implementation would be acceptable, *e.g.*, any routing mechanisms or multiple routing architectures in parallel).

Administration Plane. The administration plane is controlled by a single trusted third party, which is supposed to be a government body. The administration plane provides three basic services: (i) *Identity authentication*, which is the fundamental service that brings about the access-on-request network and is subject to be used massively. Therefore, its scalability is a major design concern. In Section 4.1.2, we will elaborate our solution to providing high scalability of this service. (ii) *Unified incentive system*, which supports aforementioned critical leverage points. It gives individual users *incentive* to comply with *standards*. It is also essential for motivating the public to help in *auditing*. (iii) *Distributed announcement database*, which is designed to announce authoritative information to public.

Identity Based Capability and Keeping State. Our framework exploits *capability* [10,35,51,55] to implement the access-on-request network. To access the network, a user must first acquire a capability, which is issued based on her permanent user identity (Section 3.2). Such *identity based capability* solution makes *keeping state* (at network devices) become readily feasible.² The network design therefore benefits from ex-

²Currently, denial of service (DoS) attacks pose a significant threat to network devices that keep state. In our framework, we can apply identity based rate-limiting to effectively counter DoS attacks and exploit appropriate solution (*e.g.*, [35]) to solve the bootstrap issue

traordinary flexibility. One essential case is the design flexibility of control information. For example: (i) We can easily support *variable length or long control information*, which is problematic in the current Internet. (ii) We can safely use *out-of-band* channels (which are separate from, but associated with, the data channels) to transfer control information. (iii) We can significantly improve the processing efficiency of control information by *amortizing its overhead* across entire service sessions.

New Network Layer. To support sustainability, we introduce a new layered model for our access-on-request network. As shown in Figure 2, this new model is adapted from the TCP/IP model by revising the network layer. The new network layer has three distinct features: (i) *Semantics aware*. It is aware of service level semantics (*i.e.*, *abstract* application layer semantics) and handles traffic based on it. (ii) *Bidirectional extensibility*. It provides good extensibility in both upward and downward directions. The TCP/IP model instead only supports good extensibility in the upward direction. A typical example of the downward extensibility in our model is that it can easily upgrade the network routing architecture and its underlying data link infrastructure; it can even support multiple routing architectures working in parallel. (iii) *Dual Layers*. It integrates both the transport layer and the network layer of the TCP/IP model. This necessarily results from the semantics aware property. The new network layer is divided into four sublayers:

- The *intermediate* (network-I) sublayer defines *abstract* and *immutable* network layer primitives (network-I primitives) that are generally applicable to all higher level semantics and lower level infrastructures.
- The *upward adaptation* (network-U) sublayer adapts network-I primitives to service specific protocols. It also defines additional protocols to meet new demands of evolving application layer services.
- The *downward adaptation* (network-D) sublayer adapts network-I primitives to protocols specific to different lower level infrastructures (*i.e.*, routing architectures) based on service level semantics.
- The *semantics and maintenance* (network-S) sublayer provides a *common and compact* interface (i) to convey application layer semantics to lower layers in an efficient way and (ii) to support network-wide maintenance functionalities (*management, fault diagnosis, measurement, etc.*)³

Such division decouples the network layer’s *variant part* (subject to evolve) from its *invariant part* (evolves very slowly, nearly immutable). The variant part includes the network-U, network-D, and network-S sublayers while the invariant part is the network-I sublayer. The decoupling helps to ensure good evolvability. A central criteria to implement the variant part is making it *globally upgradable*, *e.g.*, all routers can upgrade specific software components consistently. In our framework, we leverage *standards, incentives, and auditing* to enforce the global upgradability. However, we may also try certain innovative network solution (*e.g.*, the “active network” proposed in [48]) to help achieve the global upgradability.

Network Layer Metadata. The network layer *metadata* are control information (out-of-band typically) associated with each user service session. They play a similar role as the TCP and IP headers in the TCP/IP model, but in a much more comprehensive way. One distinct and fundamental functionality of the metadata is to carry semantics and maintenance information being processed at network devices. The metadata are crucial for the semantics aware property and the high level controllability of our proposed architecture.

3.2 Identity

Identity is a key to our access-on-request solution. We use *identity based capability* to control network access. We introduce two forms of identity: *user identity* and *data identity*. Capabilities are issued via identity authentication which mandatorily inspects user identity and optionally checks data identity.

User Identity. A capability is issued based on a user’s permanent identity, which is the same for all services on the access-on-request network. This permanent user identity is issued and managed by a single trusted third party — *identity authority* (a government body). The user identity is also the premise for us to build user profiles such that we can apply effective incentive mechanisms.

A practical user identity solution is crucial for our framework. To this end, we propose a comprehensive solution of user identity, as we will introduce in Section 4.1.2. This solution can provide a scalable identity authentication service capable of being accessed massively across the network (although authentication

of capability. In addition, with fear of consequences enhanced by traceable user identity, DoS attacks are fundamentally deterred.
³As pointed out in [12, 14, 29]: (i) the difficulty of network management today is partially attributed to too many details exposed by heterogeneous infrastructures and diversified protocols at different layers; (ii) fault diagnosis can benefit a lot from correlating diversified elements at different layers. The network-S sublayer provides a common and compact interface that can significantly reduce complexities resulting from the ever-evolving data plane.

data are under centralized administration). Meanwhile, it implements the permanent user identity without compromising user privacy. In addition, this solution bears strong resilience to identity theft.

Data Identity. In addition to user identity, we propose *data identity*, in which each data object is assigned a unique identity. This allows services to perform data-object-based authentication in addition to user-based authentication. We can assign each newly created data object a new data identity. This data identity encodes (i) the identity of the user who creates the object, (ii) the length, and (iii) the digital fingerprint of the data object. The length and the digital fingerprint encode unique properties of the object such that we can verify the consistency between the data identity and the object. The data identity also includes a digital signature granted by the identity authority which helps to verify the integrity of the above three components. Prior to creating the signature, the identity authority verifies (via identity authentication) that the user who is requesting the signature is the same person as the user encoded in the data identity.

3.3 Standards, Incentives, and Auditing

In order to guide Internet users (both individuals and businesses) to behave well, we need methods that can evaluate and respond to users' behavior in a qualitative and quantitative way, as a cyber law does. However, as pointed out in Section 2.2, a cyber law can not keep up with the fast evolution pace of Internet services. Therefore, we need a more responsive solution. In addition, at a high level, we need a mechanism to guide Internet evolution in the right direction.

Standards. We introduce *standards* to address both the above issues. (i) At the low level, standards work as "responsive cyber laws." Standards enforce users to behave well in a *more relaxed* way than a cyber law does. Rather than addressing how to punish violations, standards emphasize on how to encourage and reward good behavior. Although standards also define penalties for violations, they are usually simply the deduction of benefits. Therefore, standards can be enacted *more tentatively*. We can introduce more tentative clauses to standards. If undesired side effects show up, the standard can be modified, suspended, or even canceled. (ii) At the high level, standards bring about new protocols that *reach global agreements* among service providers. In this way, standards facilitate controlled Internet evolution. Standards can be enacted through the Pyramid process (Section 3.4), which ensures both the comprehensiveness (to be globally agreed upon) and the responsiveness requirement of new standards.

Incentives. Incentives are a central topic of sustainable development (for both SusEco and SusInet). Neither businesses nor individual users spontaneously follow the sustainable practices. (i) To address incentives for businesses, we directly apply a central guideline of SusEco — introducing government involvement and having government focus on the *market-based approach* to motivate businesses to follow sustainable practices. In SusEco, such practices are to address economic development's pressure on natural resources and environment; in SusInet, such practices are to address Internet development's pressure on a set of culture norms that Internet development itself also relies on. Rather than direct intervention, government provides *incentives and market signals*, e.g., via subsidies and levies, for businesses to conduct sustainable practices. (ii) To address incentives for individual users, we adapt the market-based approach for businesses to create a version that can be applied to individual users, i.e., the *unified incentive system*. The unified incentive system motivates users *to follow standards* and *to contribute*.

Auditing. Auditing is the method to verify whether a user's or a network device's behavior complies with the standards. The ideal way of auditing is that we can both detect and provide irrefutable evidences for compliance with and violation of the standards. However, this could be extremely hard to implement in practice due to various difficulties: (i) auditing without infringing on users' privacy and data confidentiality; (ii) auditing without introducing too much overhead and cost; (iii) auditing without losing generality (e.g., make an audit method compatible with different routing mechanisms and path properties). In our framework, we relax the concept of auditing and focus more on auditing compliance than auditing violations. We exploit *incentives* to encourage the majority of users to comply with and to cooperate with auditing. We can use *identity* as an ultimate weapon to threat and to counter the minority that violate the standards. In addition, we enable the public to audit government such that we can quickly detect and respond to any government (intentional or unintentional) wrongdoings and prevent the abuse of power.

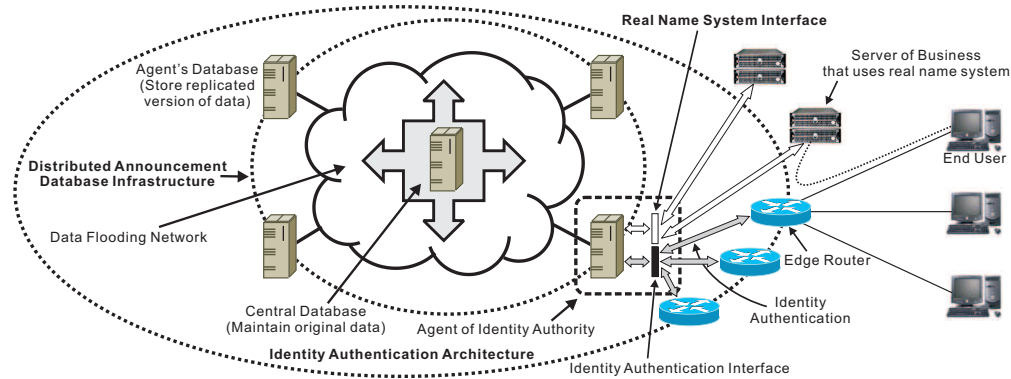


Figure 3: Distributed Announcement Database Infrastructure (in the Scenario of Identity Authentication Architecture)

3.4 The Pyramid Process

The Pyramid process (Chapter 23 of [26]) is a multi-stakeholder engagement process that is especially suitable to address the *huge challenges* of sustainable development: (i) Sustainable development requires contributions from all parties — governments, business, civil societies and international bodies. (ii) It creates unprecedented demands for learning, thinking, planning and decision-making. (iii) Initiatives seeking to promote sustainability are often doing so under a sense of time urgency, with limited resources; we do not have time or money to waste on suboptimal solutions or difficult-to-achieve agreements.

The Pyramid process supports group learning, planning, and decision processes. In addition, it can produce, as a purposeful by-product, improved levels of interdisciplinary understanding and innovative thinking. We apply the Pyramid process to the SusInet context. On one side, the Pyramid process helps with the enactment of standards. It ensures that new standards reach global agreements among service providers and other entities. Meanwhile, it guarantees that the process meets the responsiveness requirement of standards. On the other side, Internet can provide efficient platforms to help the Pyramid process; and progress in SusInet can improve the effectiveness of such platforms (by adding proper controllability).

4 Research Agenda

We will design and prototype four core systems crucial for the proposed framework: (i) the distributed announcement database infrastructure of the administration plane and systems derived from it; (ii) a common platform that collects and disseminates network auditing information; (iii) a prototype of the semantics-aware network; and (iv) a basic version of the unified incentive system. These systems (or some of their sub-systems) can even be directly applied to the current Internet to benefit *services*, e.g., a p2p or a social network. The only difference from applying them to the new access-on-request network is that they are applied at the service level for specific businesses instead of working at the network level globally to benefit everyone and every service.

4.1 Distributed Announcement Database Centric Infrastructures

The *distributed announcement database* (DAD) is designed to announce authoritative information to the public. It is a core service of the administration plane, from which many other basic services can effectively be derived. Figure 3 depicts a use of the DAD in the case of the identity authentication architecture, which is fundamental for our access-by-request solution. In this case, the DAD is the key to the scalability of identity authentication. It replicates authentication data (e.g., of all clients subscribed to the access-on-request network nationwide) from a central database of the identity authority (a government body) to a large number of identity authority's agents. Edge routers therefore can perform identity authentication distributedly by contacting these agents. One subtlety of data replication in this case is that for those confidential authentication data, we replicate a hashed version to each agent's database instead of the original data. In this way, we achieve scalability through large scale replication without compromising security. Based on this identity authentication architecture, we can further develop a very useful system — a common interface for real name systems of businesses. This system builds upon the identity authentication architecture by adding a real name system service interface at each identity authority's agent. Authorized businesses then

can import proper information from this interface to their servers such that they can manage customers and build user profiles efficiently. Below, we first introduce our research plan for a more generalized version of the DAD infrastructure. Then, we describe the aforementioned two systems derived from the DAD, the identity authentication architecture and the real name system interface, in more detail.

4.1.1 A Basic Version of DAD

A Common DAD. We will first design a framework for a general purpose DAD infrastructure by addressing its basic functionality — to implement large-scale replication of authoritative data across agents (*i.e.*, authorized entities that hold replicated versions of data). This includes the following design considerations: How to effectively flood data across agents by exploiting the single-trusted-third-party nature and what is the appropriate flooding structure? What is the proper format of each data unit to be flooded? What is the common format of control information? What are the basic data manipulation operations (for both individual data entry and bulk data) of the DAD and how to implement them?

We will study these questions by considering all possible uses of the DAD. Indeed, many important systems and services can be built on the DAD. For example, (*i*) The DAD is essential to our proposed identity authentication architecture (Section 4.1.2), which can perform user-based and object-based authentication; (*ii*) Based on the DAD, we can build an efficient platform for announcing network audit information to public (Section 4.2); (*iii*) The DAD can help to build a unified incentive system (Section 4.4), allowing the public to access a user’s qualification information; (*iv*) We can use the DAD to implement a global seed distribution service, which is the key to enable an approach proposed in [35] that can solve the bootstrap issue⁴ of capability.

Our goal here is to design a basic framework of the DAD that is compatible for all possible uses of the DAD. Actually, the above examples exhibit very different subtleties of possible uses. For instance, example (*i*) requires that the data replication process is able to additionally ensure data confidentiality; example (*ii*) and (*iv*) require the support of more dynamic data; example (*ii*) and (*iii*) additionally call for the support of data “writing” from potentially a large number of authorized entities.

Announcing Agents’ Public Keys. The DAD is designed to announce authoritative data. To this end, each agent of the DAD is assigned a public and private key pair by authority (the single trusted third party) such that the public can verify the authenticity of each agent through digital signature. Therefore, it is essential to have a mechanism that can announce agents’ public keys in an authoritative way.

We will design, implement, and evaluate a protocol for such a key announcement mechanism. The protocol will be build on the DAD infrastructure itself. Its design will learn from digital certificate solutions that already exist on the Internet. The main design consideration is how to efficiently revoke previously announced public keys. The goal of this protocol is to achieve both high efficiency and strong robustness for key announcement and revocation.

4.1.2 Scalable Identity Authentication Architecture

Based on the basic version of DAD, we will derive an identity authentication architecture that can serve identity authentication requests massively. This research includes two parts: (*i*) an identity solution that satisfies several basic design criteria of identity and (*ii*) a large-scale replication mechanism for confidential authentication data, which is built upon the DAD infrastructure.

Identity Solution. We propose an identity solution that satisfies the following basic design criteria:

- *Unified Identity.* Each user is assigned a permanent identity that is used for all services. This unified permanent identity can serve as an effective threat against misbehavior due to fear of consequences (access to all services could be affected). The unified identity also makes possible a unified incentive system that encourages users to comply and to contribute.
- *Use Single Trusted Third Party.* Our unified identity solution uses a single trusted third party, an *identity authority* (a government body), which manages identity and authentication related data.
- *Routing Independent Capability.* Network access capabilities are issued based on a user’s identity. Our capability scheme allows the maximum freedom in the design of routing architecture by making capabilities totally independent of network routing.

⁴An axiom [10] for capability-based solutions is that they rely on non-capability-based solutions for their bootstrap, *i.e.*, to establish the capability. Recent research [35,51] has given good solutions for this bootstrap issue.

- *Retain Privacy.* Our identity solution provides traceability of user identity by the identity authority. Meanwhile, it also ensures that user identity is *untraceable* by unprivileged users. Here “untraceable” means: (i) *unresolvable*, i.e., a user’s real identity is not resolvable; (ii) *undistinguishable*, i.e., unable to distinguish whether two observed identities belong to a same user.
- *Address Identity Theft.* Our solution addresses identity theft in the following way: (i) A stolen identity can be easily reclaimed by its owner. (ii) We reduce the impact of identity theft by restricting the negative impact on its owner if it happens. (iii) We provide proper mechanisms to allow users to quickly detect identity theft and respond to it. The quick response is desirable because it can effectively help to mitigate many kinds of misbehavior on the Internet that rely on identity masquerading.

Large-Scale Replication of Confidential Authentication Data. Using single trusted third party significantly simplifies the management of identity and authentication related data. But it also raises scalability challenges for identity authentication which is supposed to support massive requests in our framework. To solve this, we can perform large-scale replication of authentication data. Such replication mechanism can be built on the DAD infrastructure. Its main design consideration is to retain privacy of identity and authentication related data. We can exploit cryptographic hash functions and the pair of public-private keys assigned to each agent to address this issue. We replicate a hashed copy to an agent instead of the original data. Each agent receives a different hashed version seeded by its assigned keys. In this way, each agent can independently perform identity authentication without knowing the original data and replicated data on one agent is totally useless elsewhere.

We will design and implement both the large-scale replication mechanism and a corresponding identity authentication protocol between agents and end users. We will evaluate their scalability and security resilience quantitatively via emulation and Internet based experiments.

Implementation Highlights. Here we introduce our implementation scheme for both the identity solution and the large-scale replication mechanism. Due to space limits, we only highlight the critical points.

PID and TID. The identity authority issues each user a *permanent identity (PID)* through registration. Along with the *PID*, the identity authority also issues a number of secret codes (*SEC*) to the user, each *SEC* is corresponding to a day. Both the *PID* and the *SEC* are stored on a small hardware device — *network passport (N-pass)*. To access the Internet, a user plugs her N-pass into the computer. The N-pass helps her computer to generate *temporary identities (TID)*, which are used for authentication. The relationship between *TID* and *PID* can be represented by Equation (1). Such an implementation makes the observed identity, *TID*, be *untraceable* by unprivileged users while resolvable by the identity authority.

$\begin{cases} TID = f(PID, time, seeds, PubKey) \\ PID = g(TID, time, seeds, PriKey) \end{cases} \quad (1)$	$VC = v(TID, SEC) \quad (2)$
<i>PubKey, PriKey:</i> the public, private key pair of the identity authority. <i>f:</i> a function to generate <i>TID</i> from <i>PID</i> . <i>g:</i> a function to resolve <i>PID</i> from <i>TID</i> . <i>seeds:</i> any commonly known parameters (e.g., service type).	<i>v:</i> a commonly known hash function.

Authentication. To perform authentication, we require the user to supply a verification code (*VC*) along with the *TID*. The *VC* is computed based on the user’s *SEC*, which is only known by the user herself and the identity authority. The method to generate *VC* is represented by Equation (2). Upon receiving the *TID* and *VC*, the identity authority first resolves the user’s *PID*, which in turn helps to retrieve the user’s *SEC*. Then it verifies the *VC* by regenerating a *VC* the same way as the owner does.

Address Identity Theft. To address identity theft, we can embed a clock into the N-pass. Using the clock, we can enforce that the stored *SEC* for a specific day can be accessed *only* on that day (i.e., impossible to access in advance). This effectively reduces the risk of leaking *SECs*. In addition, if the N-pass itself is stolen, the identity owner can easily reclaim her N-pass via the identity authority by changing her *SECs*.

Large-scale Replication. We transfer to each agent the following data: *PID* and *SEC* of all users, the private key (*PriKey*) of the identity authority. Meanwhile, we address security issue by exploiting cryptographic hash functions. First, we have the identity authority generates a public and private key pair (*PubKey_i*, *PriKey_i*) for each agent, *i*. Second, we modify Equations (1) and (2) to Equations (3) and (4) respectively. Third, we disseminate the following data to each agent *i*: *PID_i* and *V_i(SEC)* of all users, *G(PriKey_i)*, and *V_i*.

$\begin{cases} TID = f(PID, time, seeds, PubKey_i) \\ PID_i = g(TID, time, seeds, G(PriKey_i)) \end{cases} \quad (3)$	$\begin{cases} VC = v(TID, SEC) \\ V_i(VC) = v(TID, V_i(SEC)) \end{cases} \quad (4)$
<i>G, V_i:</i> cryptographic hash functions.	

When a user generates a *TID* for authentication, she uses the public key of an agent ($PubKey_i$) instead of that of the identity authority ($PubKey$). Each agent stores a hashed copy of user identity (PID_i) and it is resolvable from the *TID* by using a hashed version of private key ($G(PriKey_i)$). The way a user generates a verification code (*VC*) remains the same (first line of Equation (4)). Each agent stores a hashed copy of secret codes ($V_i(SEC)$), based on which it can check the validity of verification codes (using the second line of Equation (4)), thereby authenticating the user. In this way, we achieve to replicate authentication data without compromising security. Authentication data stored by one agent is useless for other agents (can neither resolve nor distinguish user identities).

4.1.3 A Common Interface for Real Name Systems

For purposes such as management and personalized services, businesses often prefer to use real name systems — one account per user, user identity traceable by authority, partial user information (e.g., age, gender) resolvable. To this end, we propose to design a common interface provided by the identity authentication architecture that serves authorized businesses to build their real name systems.

We will design, implement, and evaluate a protocol for this common interface. Our work will focus on the implementation subtleties of this interface as a practical business service. A business can import a hashed copy of its clients' authentication data through this interface. The business can also import partial user information and use this information in an authorized way. The authorization of a business's real name system can be easily suspended, revoked (if it violates the authorization agreement), and resumed. Impacts of a business's misbehavior should be limited. In addition, the identity authentication architecture should effectively protect a business's privacy for its customer information.

4.2 Network Auditing Information Platform

A basic design guideline of auditing in our framework is that methods to audit network transit part, *i.e.*, routing architectures, should focus on network performance auditing. To facilitate such auditing methods, we propose to design a common platform that collects and disseminates network auditing information, *i.e.*, authoritative performance measurement results. Consequently, the design of each specific auditing method can simply focus on how to measure and audit each specific performance metric (e.g., packet loss, delay, outage, capacity, or available bandwidth). While the common platform will take care of most of the rest parts: (i) Where and how to collect raw measurement and auditing results? (ii) Where and how to generate authoritative measurement results from raw ones? (iii) How to reach arbitration for disputes among results from different administrative domains? (iv) How to archive evidences for both resolved and unresolved disputes? (v) What kind of auditing information, and how, to disseminate to the public?

4.2.1 Quasi Real Time DAD for Network Auditing Information

One essential part of this platform is a specific DAD system that can announce authoritative network auditing information in a quasi real time manner. End users therefore can acquire recent network performance information from this system. The system only keeps data of a recent period of time such that its data amount is relatively bounded. Storage space for outdated data will be reclaimed for fresh data.

We can derive this system from the basic version of DAD by adding enhanced features that support the very dynamic auditing information. The main design consideration includes two aspects: (i) How to design an effective data scheme for the auditing information such that we can minimize dissemination overhead? (ii) How to design an efficient dissemination scheme to deliver auditing information across this quasi real time DAD infrastructure to endpoints that really (or potentially) want it, avoiding unnecessary delivery?

4.2.2 Information Feed-in System

Figure 4 describes the structure of this auditing information platform. As it shows, another essential part of this platform is an information feed-in system that *collects* raw measurement results, *processes* them, and *prepares* authoritative results to be disseminated. Raw measurement results are collected by *auditors*, which are machines administered by a single trusted third party. Each administrative domain reports measurement results to a (or several) designated auditor(s). The results could be further relayed to other auditors. Each auditor processes measurement results reported from related administrative domains and generates a signed auditing summary (information to be disseminated to public). If a dispute among administrative domains happens, the auditor signs measurement results reported from each administrative domain and sends them to designated *arbitrators* — authorized entities that archive and resolve disputes.

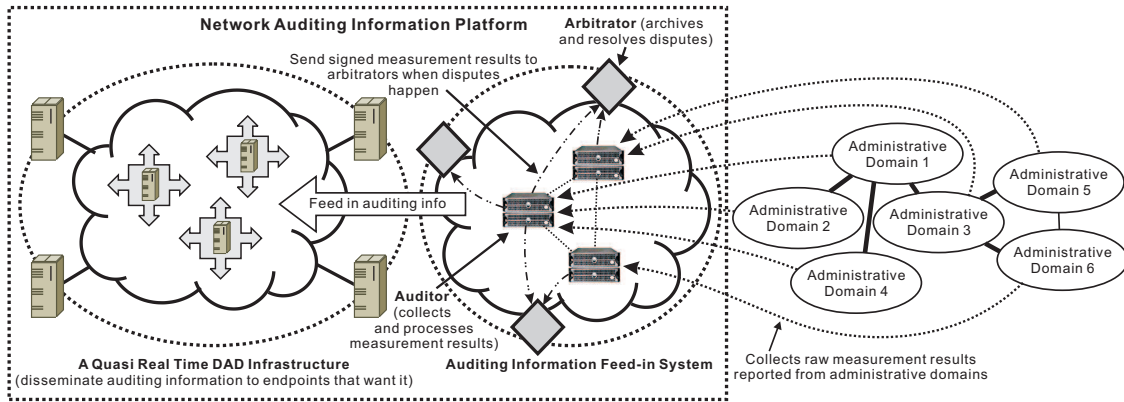


Figure 4: Structure of the Network Auditing Information Platform

The main design considerations of this information feed-in system include: What is the appropriate scheme to assign auditors to administrative domains? How do auditors collaborate with each other and how do they inject data to the quasi real time DAD system? How to generate auditing summary sufficiently useful for public but at the same time retain service providers' privacy? What kind of information (or events) is mandatory to report to the auditor and how to enforce this? To answer these questions, we will focus on the auditing information of administrative-domain-level packet loss and delay (assuming they are defined and measured in a way described in [11]) as a research starting point.

4.3 Semantics Aware Network Prototype

Adding a semantics aware property to the Internet architecture is a fundamental principle of our framework. We will follow our proposed new network layer model (Section 3.1) to perform in depth research on implementing this semantics aware property.

4.3.1 Invariant Primitives And Bidirectional Adaptation

Our research here will focus on an in-depth exploration for the invariant, network-I primitives. The main questions we will study include: (i) What are these generally applicable network-I primitives which can cover most network layer functionalities needed? (ii) Are these primitives abstract enough such that they are immutable in the long run? (iii) How can we adapt these primitives upward to concrete protocols that serve diversified application level services and semantics? (iv) How can we make these primitives compatible to all possible underlying routing architectures? (Regardless of whether a routing architecture is in the traditional end-to-end mode or in a cache and forward style [56], forwarding via single- or multi-path, via reliable or lossy channels, requiring timely transfer or not, *etc.*) (v) How can we adapt these primitives downward to each specific routing architecture?

For the upward and downward adaptation, we will focus on two kinds of redirection methods at network edges, as described in Figure 5(a). For the upward adaptation, we will explore how to redirect the network layer metadata to appropriate semantics processing infrastructures (*e.g.*, a data-object-based authentication system). For the downward adaptation, we will study how to redirect user traffic to an appropriate underlying routing architecture according to semantics encoded in the metadata. We will evaluate both kinds of redirection methods in terms of efficiency and scalability.

4.3.2 Network Layer Metadata Prototype

The network layer metadata is a crucial element for the semantics aware property in our framework. They are typically out-of-band control information for user service sessions and they carry semantics information being processed at network devices, especially at network edges. Our research here will focus on implementing a prototype of the network layer metadata and the network layer functionalities surrounding the metadata. As shown in Figure 5(a), the metadata encode semantics information from applications at endpoints, being processed at edge routers. Edge routers redirect the metadata to appropriate semantics processing infrastructures if necessary. After the metadata gets processed, states are established for associated sessions. Edge routers store the states and redirect user traffic of associated sessions to appropriate

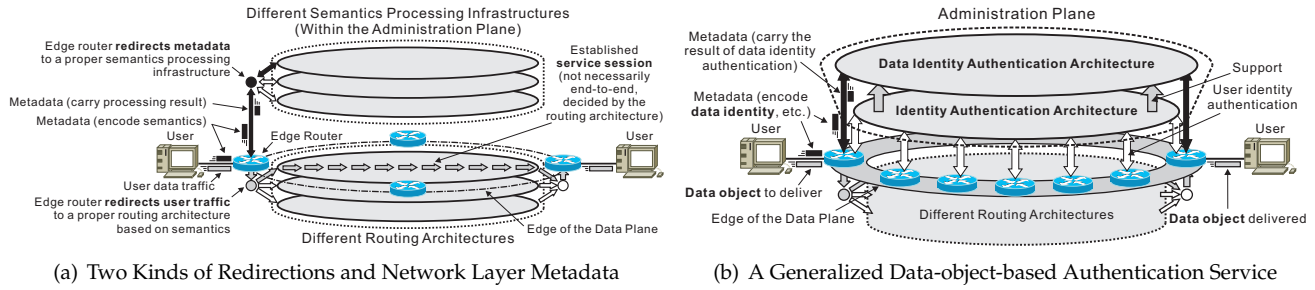


Figure 5: Semantics Aware Network Prototype

underlying routing architectures according to the states. Edge routers and semantics processing infrastructures can also exploit the metadata to pass control information to routing architectures.

The main design consideration include: (i) What is the general format of the metadata exchanged between endpoints and edge routers? What is the general format of the metadata sent from edge routers and semantics processing infrastructures to routing architectures? (ii) How to make the above formats common and compact (*i.e.*, globally understandable by network devices, capable of being efficiently processed)? (iii) How to process the metadata at edge routers and semantics processing infrastructures? (iv) How to amortize the metadata’s processing overhead across an entire service session? (v) How to evolve the metadata (both the format and processing methods) without compromising the common and compact property?

4.3.3 Case Study

We will exploit a specific service context to help us understand and evaluate subtleties of the prototype design. This service context is a *generalized data-object-based authentication service for data delivery*, as shown in Figure 5(b). The data delivery can be either in client-server mode, in P2P file sharing mode, or in form of an email attachment, *etc.* The distinct semantics information carried in the metadata here is the data identity (Section 3.2) for the object to be delivered. Edge routers redirect the metadata to a special semantics processing infrastructure — a data identity authentication architecture. This architecture authenticates each data identity in terms of copyright and immunity information. If the authentication is successful, the edge routers will redirect associated data delivery traffic to an appropriate routing architecture according to other semantics information. We will describe this specific service in more detail in Section 5.3.

4.4 Unified Incentive System

In our proposed framework, the unified incentive system (UIS), designed to bring all Internet services together, is a fundamental administration plane’s service that addresses incentives for individual users.

4.4.1 Introducing Two Incentive Elements

We will design and prototype a basic version of the UIS by introducing two incentive elements: (i) *Reward point*, which is modeled off of diversified forms of credit points used by businesses that can be redeemed for gifts or money. (ii) *Compliance & contribution score*, which is modeled off of the credit score in the credit card system. This score quantifies the extent to which a user complies with standards and contributes to public. A user can benefit from a high score and can suffer from a low score. For example, if her compliance & contribution score is too low, she might be blocked from accessing the Internet.

Our research will focus on developing a common interface provided by a single trusted third party for manipulating these two incentive elements. The main design considerations include: (i) What are the proper data formats and components of the reward point and the compliance & contribution score? (ii) For each of these incentive elements, what are the basic data manipulation operations we should provide to support possible uses? (iii) How to exploit the DAD infrastructure to build this interface? (iv) How to ensure the fairness and integrity for updates of the two incentive elements? (v) How to adapt ideas of financial transactions to implement manipulations for reward points? (vi) How to make the “writing” of compliance & contribution score scalable given that a large number of authorized entities can update it?

4.4.2 Updating and Auditing Compliance & Contribution Score

An authorized entity (*e.g.*, a business) can update a user’s compliance & contribution score according to related standards. The standards provide concrete definitions for “compliance” and “contribution” behavior of users in the related context. The standards also quantitatively define how to update a user’s compliance & contribution score according to her behavior. To facilitate this, the common interface of the UIS will not let an authorized entity to directly modify a user’s compliance & contribution score. Instead, it lets the entity describes the nature (compliance, violation, or contribution) of a user’s behavior by indicating the corresponding standard item. The interface then updates the user’s compliance & contribution score according to the description and the standard. Each update event will be logged for auditing purpose. In addition, users will be informed of such events if necessary to help guarantee the integrity of updates. Main design considerations for the above mechanism are: (i) How to easily and smoothly add new standards to this common interface of the UIS? (ii) How to support the tests of new standards?

5 Evaluation

5.1 Theoretic Evaluation

An important part of our work is to evaluate our design theories in depth. We will compare our theories and actual design with related work across computer science, economics, sociology, and law. The goal is to show in detail how our solutions can satisfy or reconcile with related design criteria, how our systems can potentially solve currently unsolvable problems across technical, economic, and social domains. Here is a list of the main design criteria and questions we will evaluate: How to (i) enable a global access control to service traffic without compromising user privacy, data confidentiality, diversity of routing choices, and scalability; (ii) facilitate network management, measurement, and fault diagnosis by correlating different layers and diversified devices across the network in an efficient way; (iii) evolve the Internet architecture through incremental deployment and end-user-chosen service evolution (a.k.a. bottom-up control); (iv) promote multi-stakeholder engagement processes of decision making and learning to help the Internet development; (v) introduce government involvement and have government play the correct role in the Internet development; (vi) reach a balance between controllability and neutrality, such that no specific interest group can control the Internet in a strategic way; (vii) foster a solution to currently unsolvable conflicts between creativity and copyright laws?

5.2 Performance Evaluation

Using a single trusted third party significantly facilitates system design for controllability. In addition, it serves as an effective basis to foster global trusts, which is essential for sustainable solutions. However, it brings about performance concerns on system scalability. We therefore will evaluate performance of our systems by focusing on their scalability. Meanwhile, we will also evaluate major performance benefits resulting from the single autonomous system nature (*e.g.*, multicast-like forwarding can be effectively exploited for the administration plane). To this end, we will build prototype implementations of proposed systems in the *Emulab* [1] and on the PlanetLab [3]. For those systems and protocols that can be directly applied to the current Internet, we will particularly focus on evaluating their prototype implementations, with a goal of developing ready-to-use solutions that can benefit the current Internet. Such systems and protocols include: the DAD infrastructure, solution of user identity, identity authentication protocols, the real name system interface, the compliance & contribution score and its manipulation interface, solution of data identity, and data-object-based authentication protocols.

5.3 Service-Driven Evaluation

Integral Evaluation. We will evaluate the proposed systems *as a whole* by considering a specific service scenario — *a generalized data-object-based authentication service for data delivery* — as depicted in Figure 5(b). We select this service because it puts together most of our proposed systems. The key part of this system is a specific semantics aware network which can provide comprehensive control for data delivery service based on data identity and other semantics. Meanwhile, since this network assumes the existence of a access-by-request mechanism, it necessarily relies on the identity authentication architecture that we propose. In addition, it also closely relates to the unified incentive system, with which users can be effectively motivated to comply with standards of this service and cooperate with auditing.

Copyright Protection Service. We will design the data-object-based authentication service to provide effective solutions to (i) copyright protection and (ii) immunity (*i.e.*, to counter virus, malware, spam, *etc.*). It accomplishes this by enforcing data identity authentication globally at the network level for all data delivery services (*e.g.*, in the client-server mode, the P2P file sharing mode, or in form of an email attachment, *etc.*). We will in particular focus on the solution to copyright protection in this service. This is because: (i) A practical solution to copyright protection is especially meaningful for the Internet development. (ii) Our solution to copyright protection is a typical case that is required to leverage *identity, standards, incentives, and auditing* in synergy, hence a representative implementation example of our entire framework. Therefore, showing the effectiveness of this solution to a large extent will prove the effectiveness of our framework. Here we highlight our solution to copyright protection in this data-object-based authentication service:

- We set a standard for data delivery services: The source endpoint must provide the data identity in the metadata for the data object that it wants to send during the session setup phase; both endpoints are required to verify the consistency between the data identity and the actual data object; if inconsistent, the source endpoint should not deliver the object, and the destination endpoint should drop it.
- The source edge router redirects metadata (which encodes the data identity and user information) to the data identity authentication architecture, which performs authentication by checking: (i) whether the session associates with a valid *copyright authorization*, *e.g.*, whether the destination user owns an authorization for downloading or whether the source user owns an authorization for distributing the object; (ii) whether the data identity maps to a known *copyright infringing object*, *e.g.*, a “cracked” software.
- The data identity encodes the identity of the user (resolvable by the identity authority) who created the data object. Therefore, it poses a threat to anyone who wants to generate copyright infringing objects, because she must use her own user identity to generate the data identity.
- Warnings or penalties for violations of standards can be effectively applied via the unified incentive system. Compliance with standards will be honored also via the unified incentive system. Both the compliance and violations are detected and verified via auditing.

6 Related Work

Our identity-based-capability solution learns from related work on capabilities (*e.g.*, [10, 35, 51, 55]) and host identity & address schemes (*e.g.*, [8, 29, 38, 44, 50]). Unlike other capability solutions, our solution does not rely on any routing level (or link level) assumptions. It works totally independent of routing, thereby allowing maximum flexibility for routing architectures and address schemes. Our solution adopts permanent identity instead of using temporary ones as most host identity schemes do. In addition, our solution ties all services together with a unified identity. These two distinct features make our identity solution not only work well for capability, but also capable of serving as an effective threat against misbehavior.

Our research about controllability refers to a wide range of related work. The game theory based behavioral study and solutions (*e.g.*, [7, 33, 34, 41]) provide us experiences on developing solutions to behavior of service providers and the public. Related work on centralized administration (*e.g.*, [19, 20, 53]) provides us experiences on designing the administration plane that we propose. Still, our solution is in the context of a nationwide network instead of an enterprise network, and in the context of using government as a single trusted third party for administration. To address scalability of the administration plane, we will adapt proper ideas of related work on scalable online information processing (*e.g.*, [18, 23, 36, 39, 45]). Related work on reputation systems (*e.g.*, [5, 22, 52]), incentives (*e.g.*, [25, 28, 42, 47, 49]), information privacy and trust (*e.g.*, [2, 15, 24]), auditing (*e.g.*, [11, 27, 40, 43]), future Internet architectures (*e.g.*, [16, 37, 48]), routing schemes (*e.g.*, [6, 9, 46, 56, 57]), cryptographic tools (*e.g.*, [17, 21]), and network management and fault diagnosis (*e.g.*, [12, 14, 30]), provides us ideas on designing the incentives and auditing leverage points of our framework. Contrary to such approaches, which are targeted to specific services, we are designing infrastructures that support incentives and auditing in a more generalized way. Finally, our research also integrates ideas and principles from economics, sociology and law, *e.g.*, [26, 31, 32].

7 Plan of Work

Table 6 outlines the timeline of our research agenda.

Proposed Research Tasks	Year 1	Year 2	Year 3
T1. Distributed Announcement Database Centric Infrastructure (Sec. 4.1)			
1. Analyze possible uses of the DAD and evaluate candidate structural models (Sec. 4.1.1)	█		
2. Design a basic version of DAD (Sec. 4.1.1)	█	█	
3. Design and evaluate a scalable identity authentication architecture (Sec. 4.1.2)		█	
4. Design and prototype a common interface for real name systems (Sec. 4.1.3)		█	
T2. Network Auditing Information Platform (Sec. 4.2)			
5. Analyze possible uses of the platform and define detailed design guidelines (Sec. 4.2)	█		
6. Design and evaluate a quasi real time extension of the DAD for network auditing information (Sec. 4.2.1)		█	
7. Design and evaluate an information feed-in system for the platform. (Sec. 4.2.2)		█	█
T3. Semantics Aware Network Prototype (Sec. 4.3)			
8. Analyze possible services, build models, define invariant primitives (Sec. 4.3.1)	█	█	
9. Design network layer metadata and related processing interfaces (Sec. 4.3.2)		█	█
10. Design and prototype a generalized data-object-based authentication service for data delivery (Sec. 5.3)		█	█
T4. Unified Incentive System Fundamental (Sec. 4.4)			
11. Analyze possible uses of the unified incentive system and generate concrete models (Sec. 4.4.1)	█	█	
12. Design and prototype the reward point and its manipulation interface (Sec. 4.4.1)		█	
13. Design and evaluate a scalable "writing" extension of the DAD (Sec. 4.4.1)		█	
14. Design and prototype the compliance & contribution score and its manipulation interface (Sec. 4.4.1, 4.4.2)		█	█

Figure 6: Project Timeline

8 Impact

Our research will bring about two innovative insights on fundamental design guidelines of the future Internet. First, basic principles and experience in SusEco (*i.e.*, the traditional sustainable development) can help a lot for the design of the future Internet. In particular, a central guideline of SusEco on the correct role of government in sustainable development can be directly applied to the SusInet context to address businesses' incentives to conduct sustainable practices. This guideline can also be adapted to a unified incentive system that we propose such that it can directly address incentives for individual users as well. Second, we will expose the feasibility to exploit technical means of the Internet to foster *non-technical solutions* (that demand change in human values or ideas of morality), which are keys to problems unsolvable nowadays. Our research will explore in depth how we can achieve this by leveraging the four crucial elements (*identity, standards, incentives, and auditing*) in synergy. The impact of this part will go far beyond the Internet development area, contributing to the entire economic and social development.

Our research on the four proposed core systems will expose great feasibility and provide valuable experience on building scalable wide-area administrative infrastructures using a single trusted third party. We will provide experience on how to achieve comprehensive controllability without compromising high scalability, user privacy, data confidentiality, diversity of routing choices, and its balance with neutrality.

References

- [1] Emulab. <http://www.emulab.net/>.
- [2] LOAF. <http://loaf.cantbedone.org/>.
- [3] Planetlab. <http://www.planet-lab.org/>.
- [4] E. Aboujaoude, L. Koran, N. Gamel, M. Large, and R. Serpe. Potential markers for problematic Internet use: A telephone survey of 2,513 adults. *CNS Spectrums: The International Journal of Neuropsychiatric Medicine*, Oct. 2006.
- [5] T. B. Adler and L. de Alfaro. A content-driven reputation system for the wikipedia. In *Proceedings of ACM WWW '07*, Banff, Alberta, Canada, May 2007.
- [6] A. Akella, J. Pang, B. Maggs, S. Seshan, and A. Shaikh. A comparison of overlay routing and multi-homing route control. In *Proceedings of ACM SIGCOMM '04*, Portland, OR, Aug. 2004.
- [7] A. Akella, S. Seshan, R. M. Karp, S. Shenker, and C. H. Papadimitriou. Selfish behavior and stability of the Internet: A game-theoretic analysis of TCP. In *Proceedings of ACM SIGCOMM '02*, Pittsburgh, PA, Aug. 2002.
- [8] D. Andersen, H. Balakrishnan, N. Feamster, T. Koponen, D. Moon, and S. Shenker. Holding the Internet accountable. In *Proceedings of ACM HotNets-VI*, Atlanta, GA, Nov. 2007.
- [9] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Rao. Improving Web availability for clients with MONET. In *Proceedings of NSDI '05*, Boston, MA, May 2005.
- [10] K. Argyraki and D. Cheriton. Network capabilities: The good, the bad and the ugly. In *Proceedings of ACM HotNets-IV*, College Park, MD, Nov. 2005.
- [11] K. Argyraki, P. Maniatis, O. Irzak, S. Ashish, and S. Shenker. Loss and delay accountability for the Internet. In *Proceedings of IEEE ICNP '07*, Beijing, China, Oct. 2007.
- [12] P. V. Bahl, R. Chandra, A. Greenberg, S. Kandula, D. Maltz, and M. Zhang. Towards highly reliable enterprise network services via inference of multi-level dependencies. In *Proceedings of ACM SIGCOMM '07*, Kyoto, Japan, Aug. 2007.
- [13] H. Ballani, Y. Chawathe, S. Ratnasamy, T. Roscoe, and S. Shenker. Off by default! In *Proceedings of ACM HotNets-IV*, College Park, MD, Nov. 2005.
- [14] H. Ballani and P. Francis. Conman: A step towards network manageability. In *Proceedings of ACM SIGCOMM '07*, Kyoto, Japan, Aug. 2007.
- [15] L. Banks, S. Ye, Y. Huang, and S. F. Wu. Davis social links: Integrating social networks with Internet routing. In *Proceedings of ACM SIGCOMM LSAD workshop*, Kyoto, Japan, Aug. 2007.
- [16] A. Bavier, N. Feamster, M. Huang, L. Peterson, and J. Rexford. In VINI veritas: Realistic and controlled network experimentation. In *Proceedings of ACM SIGCOMM '06*, Pisa, Italy, Sept. 2006.
- [17] M. Bellare, D. Micciancio, and B. Warinschi. Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions. In *Proceedings of EUROCRYPT '03*, Warsaw, Poland, May 2003.
- [18] L. F. Cabrera, M. B. Jones, and M. Theimer. Herald: Achieving a global event notification service. In *Proceedings of IEEE HotOS-VIII*, Elmau/Oberbayern, Germany, May 2001.
- [19] M. Casado, M. J. Freedman, J. Pettit, J. Luo, N. McKeown, and S. Shenker. Ethane: Taking control of the enterprise. In *Proceedings of ACM SIGCOMM '07*, Kyoto, Japan, Aug. 2007.

- [20] M. Casado, T. Garfinkel, A. Akella, M. J. Freedman, D. Boneh, N. McKeown, and S. Shenker. SANE: A protection architecture for enterprise networks. In *Proceedings of USENIX Security Symposium*, Vancouver, B.C., Canada, Aug. 2006.
- [21] D. Chaum and E. van Heyst. Group signatures. In *Proceedings of EUROCRYPT '91*, Brighton, UK, Apr. 1991.
- [22] S. C. Corporation. TrustedSource: The next-generation reputation system for enterprise gateway security. *White Paper*, July 2007.
- [23] F. Dabek, M. F. Kaashoek, D. Karger, R. Morris, and I. Stoica. Wide-area cooperative storage with CFS. In *Proceedings of ACM SOSP '01*, Banff, Alberta, Canada, Oct. 2001.
- [24] P. A. Dinda. Addressing the trust asymmetry problem in grid computing with encrypted computation. In *Proceedings of ACM LCR '04*, Houston, Texas, Oct. 2004.
- [25] J. R. Douceur and T. Moscibroda. Lottery trees: Motivational deployment of networked systems. In *Proceedings of ACM SIGCOMM '07*, Kyoto, Japan, Aug. 2007.
- [26] K. C. H. et al. *The Natural Advantage of Nations (Vol. I): Business Opportunities, Innovation and Governance in the 21st Century*. Earthscan/James & James, Jan. 2005.
- [27] A. Haeberlen, P. Kuznetsov, and P. Druschel. PeerReview: Practical accountability for distributed systems. In *Proceedings of ACM SOSP '07*, Stevenson, WA, Oct. 2007.
- [28] C. Huang and K. W. R. Jin Li. Can Internet video-on-demand be profitable? In *Proceedings of ACM SIGCOMM '07*, Kyoto, Japan, Aug. 2007.
- [29] M. Karsten, S. Keshav, S. Prasad, and M. Beg. An axiomatic basis for communication. In *Proceedings of ACM SIGCOMM '07*, Kyoto, Japan, Aug. 2007.
- [30] R. R. Kompella, A. Greenberg, J. Rexford, A. C. Snoeren, and J. Yates. Cross-layer visibility as a service. In *Proceedings of ACM HotNets-IV*, College Park, MD, Nov. 2005.
- [31] L. Lessig. *The Future of Ideas: The Fate of the Commons in a Connected World*. Random House Inc., 2002.
- [32] L. Lessig and L. Lessing. *Code and Other Laws of Cyberspace*. Basic Books, Inc., 2000.
- [33] H. C. Li, A. Clement, E. L. Wong, J. Napper, I. Roy, L. Alvisi, and M. Dahlin. BAR gossip. In *Proceedings of OSDI '06*, Seattle, WA, Nov. 2006.
- [34] R. Mahajan, D. Wetherall, and T. E. Anderson. Mutually controlled routing with independent ISPs. In *Proceedings of NSDI '07*, Cambridge, MA, Apr. 2007.
- [35] B. Parno, D. Wendlandt, E. Shi, A. Perrig, B. Maggs, and Y.-C. Hu. Portcullis: Protecting connection setup from denial-of-capability attacks. In *Proceedings of ACM SIGCOMM '07*, Kyoto, Japan, Aug. 2007.
- [36] P. Pietzuch, J. Ledlie, J. Shneidman, M. Roussopoulos, M. Welsh, and M. Seltzer. Network-aware operator placement for stream-processing systems. In *Proceedings of IEEE ICDE '06*, Atlanta, GA, Apr. 2006.
- [37] S. Ratnasamy, S. Shenker, and S. McCanne. Towards an evolvable Internet architecture. In *Proceedings of ACM SIGCOMM '05*, Philadelphia, PA, Aug. 2005.
- [38] P. N. Robert Moskowitz. Host identity protocol architecture. *RFC 4423*, May 2006.
- [39] I. Rose, R. Murty, P. Pietzuch, J. Ledlie, M. Roussopoulos, and M. Welsh. Cobra: Content-based filtering and aggregation of blogs and rss feeds. In *Proceedings of NSDI '07*, Cambridge, MA, Apr. 2007.
- [40] M. A. Shah, M. Baker, J. C. Mogul, and R. Swaminathan. Auditing to keep online storage services honest. In *Proceedings of IEEE HotOS-XI*, San Diego, CA, May 2007.

- [41] G. Shrimali, A. Akella, and A. Mutapcic. Cooperative inter-domain traffic engineering using nash bargaining and decomposition. In *Proceedings of IEEE INFOCOM '07*, Anchorage, Alaska, May 2007.
- [42] M. Sirivianos, J. H. Park, X. Yang, and S. Jarecki. Dandelion: Cooperative content distribution with robust incentives. In *Proceedings of USENIX '07*, Santa Clara, CA, June 2007.
- [43] J. Sommers, P. Barford, N. Duffield, and A. Ron. Efficient network-wide SLA compliance monitoring. In *Proceedings of ACM SIGCOMM '07*, Kyoto, Japan, Aug. 2007.
- [44] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana. Internet indirection infrastructure. In *Proceedings of ACM SIGCOMM '02*, Pittsburgh, PA, Aug. 2002.
- [45] I. Stoica, R. Morris, D. Karger, F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for Internet applications. In *Proceedings of ACM SIGCOMM '01*, San Diego, CA, Aug. 2001.
- [46] L. Subramanian, M. Caesar, C. Ee, M. Handley, Z. Mao, S. Shenker, and I. Stoica. HLP: A next-generation interdomain routing protocol. In *Proceedings of ACM SIGCOMM '05*, Philadelphia, PA, Aug. 2005.
- [47] K. Tamilmani, V. Pai, and A. Mohr. SWIFT: A system with incentives for trading. In *Proceedings of ACM SIGCOMM NetEcon workshop*, Portland, OR, Aug. 2004.
- [48] D. L. Tennenhouse and D. J. Wetherall. Towards an active network architecture. *ACM SIGCOMM Compute Communication Review*, 37(5):81–94, 2007.
- [49] V. Vishnumurthy, S. Chandrakumar, and E. Sirer. KARMA: A secure economic framework for peer-to-peer resource sharing. In *Proceedings of ACM SIGCOMM NetEcon workshop*, Karlsruhe, Germany, Aug. 2003.
- [50] M. Walfish, J. Stribling, M. Krohn, H. Balakrishnan, R. Morris, and S. Shenker. Middleboxes no longer considered harmful. In *Proceedings of OSDI '04*, San Francisco, CA, Dec. 2004.
- [51] M. Walfish, M. Vutukuru, H. Balakrishnan, D. Karger, and S. Shenker. DDoS defense by offense. In *Proceedings of ACM SIGCOMM '06*, Pisa, Italy, Sept. 2006.
- [52] K. Walsh and E. G. Sirer. Experience with an object reputation system for peer-to-peer filesharing. In *Proceedings of NSDI '06*, San Jose, CA, May 2006.
- [53] H. Yan, D. A. Maltz, T. S. E. Ng, H. Gogineni, H. Zhang, and Z. Cai. Tesseract: A 4D network control plane. In *Proceedings of NSDI '07*, Cambridge, MA, Apr. 2007.
- [54] X. Yang, D. Clark, and A. W. Berger. NIRA: A new inter-domain routing architecture. *IEEE/ACM Transactions on Networking*, 15(4):775–788, 2007.
- [55] X. Yang, D. Wetherall, and T. Anderson. A DoS-limiting network architecture. In *Proceedings of ACM SIGCOMM '05*, Philadelphia, PA, Aug. 2005.
- [56] R. Yates, D. Raychaudhuri, S. Paul, and J. Kurose. Postcards from the edge: A cache-and-forward architecture for the future Internet. *National Science Foundation CNS-0626959*, Sept. 2006.
- [57] D. Zhu, M. Gritter, and D. Cheriton. Feedback based routing. In *Proceedings of ACM HotNets-I*, Princeton, NJ, Oct. 2002.