Identity in an Information-centric Internet

A white paper prepared by the Privacy & Security Working Group within the MIT Communications Futures Program (CFP)

Participating CFP Companies

- BT
- Intel
-
- Comcast
- Nokia Siemens Network
-
- Motorola
- Cisco
-

We observe that many contemporary services in the Internet are moving away from being centered around communication between individuals towards production, publication and retrieval of information, and this trend is foreseen to continue. This trend poses many interesting questions on essential issues in the architectural foundations of the current and future Internet. One such key question is the nature of *identity* and its surrounding issues of security and privacy that will arise in an Internet that is driven by the access to information and its surrounding tussles.

This whitepaper intends to shed light on the relevant questions in this space of identity for a potential future information-centric Internet. It is not our intention to provide answers but to motivate a clear definition of a research agenda, a vision to some extent, which can provide fuel for future work for the Privacy & Security WG within CFP. For this, we provide suggestions for concrete steps forward in our identity discussion.

The intended audience for this whitepaper is, as usual with all CFP documents, the CFP sponsor community. However, it is not intended to delimit the readership to this circle. We believe that this document provides useful thoughts and input for a wider audience that is immersed in general questions around identity, its future architectural integration and its evolution beyond today's Internet. It is to this audience that we are reaching out with our work to build a research agenda around the question on *identity in information networking*.

Executive Summary

The Internet is changing. What used to be an environment for connecting a relatively small scale research community has grown into a vital part of the infrastructural foundations of many societies around the world [2]. Current discussions around the Future Internet question many of the fundamental foundations of the 30 plus year old design of the Internet that we know of today.

From endpoints to information. Apart from the tremendous growth of the Internet we can observe another more subtle shift away from black-phone-like communication, focused on connecting endpoints, towards forms of communication where information production, discovery and retrieval is central for end users of this pervasive infrastructure. This can be seen in changes brought to us by the World Wide Web (WWW) and new contemporary applications like sensor networks, social networks, context-aware computing and others. With the diminishing role for endpoints, the role of identity is likely to change dramatically, placing the nature of *identity* for information, its producer and potentially its consumers in the foreground while transitory identity relationships between delivery endpoints will likely loose current importance.

From architectures to a single execution environment. Communication needs are surrounded by *concerns* which govern the way we use technologies and perceive their usefulness and also threats. Conflicting concerns lead to *tussles*, requiring mediation, negotiation and eventually resolution. These concerns and tussles (as well as their resolution mechanisms) are mostly deeply embedded in architectural choices and solutions eventually being deployed, i.e., tussle resolution happens pre-dominantly at system design phase. We outline an architectural vision in this whitepaper that will lead to an increasing resolution of tussles in runtime within a single (Future Internet) architecture with the help of policies and their expression as well as mediation.

Understanding identity is key. Identity is a natural field of tussle. It is an embattled area already today, with initiatives like LibertyAlliance, Passport, OpenID, OAuth [6], and many others. And it will be even more tomorrow, in particular in the light of the projected changes of the Internet towards information. In order to better evaluate future propositions, it is crucial to understand the very nature of identity in the light of these changes, e.g., the identity of players acting on information, the identity of information pieces and the identity of policies governing information. It is our ambition to help our sponsors to increase this necessary understanding.

From a research agenda to a dedicated action. In order to make headway towards the architectural vision presented in this whitepaper, we intend to outline a research agenda as a first crucial step, formulating a list of research questions that is driven by our architectural vision and the context of the changing Internet that we project. This list of questions gives a first insight in potential areas of investigation. As a dedicated action to push forward the discussions, we propose to convene a workshop with the attendance of crucial stakeholders, past and potential future ones, in the identity space. This event is intended to provide answers as well as new questions as to the role of identity in the future information-centric Internet and pave the way for a consolidated set of actions.

The Context of a Changing Internet

This paper considers the identification required for security under changing Internet conditions. We examine first three kinds of evolutions in our underlying Internet paradigm and then consider the kinds of challenges this raises in order to provide security. In each case, we will identify key questions with respect to identity.

The *first evolution* relates to change of concerns surrounding the wider use of the Internet from its early days towards today. With the spread and integration of Internet technologies into our economic, social, and civil lives, the uses and therefore perceptions of benefit have become increasingly complex. In the original Arpanet and even early days of the Internet, the community of users and usage of networking were homogeneous, reasonably compatible, and shared a fairly simple model of trust in each other and the network facilities. The network technology was developed to support sharing of resources and communications in a collaborative and trusting world. Even in the early days of the World Wide Web, the common mindset was one of sharing and cooperation.

As networking has become both increasingly commercial and central to all aspects of society, it is no longer the case that there is a single shared mindset about who ought to be able to use what and under what conditions. We can consider this change to a lack of shared concerns or objectives along a spectrum. At one end, one might find simple disagreement about what should remain private and what might become more public. Further along the spectrum, this may take the form of economic competition, in which, for example, a customer might disagree with an ISP about the cost of transport or the quality of the service provided for transport. Toward the other end of the spectrum one might find behaviors that are generally considered malicious, whether spamming, phishing and pharming, or many other kinds of malicious behavior. The spectrum identifies a collection of behaviors in which the parties to the behavior have a disagreement about acceptable behaviors, the trust that can be placed in them, and the policies controlling them. This is a significant divergence from the original model of the environment for operation in networking, and in particular points to a significant change in the underlying trust model.

The *second evolution* we identify here is from the original model of a flow of bits from source to destination. This is the TCP/IP model, with its narrow waist of an hour-glass embodied in IP. We can distinguish four significant characteristics of that model, in which change is occurring. First, in the IP model it is assumed that bits are indistinguishable. The network mechanisms simply transport all bits equally. Second, the transport protocols are assumed to transport to the destination (as best as they can) exactly the bits provided to it. The bits, and in fact, the flows are assumed to be immutable. Third, a flow of bits is exactly a pairwise flow from a source to a destination. This currently remains the most common paradigm. Finally, the flow only occurs in real-time. There is no assumption of storage or significant delay in the network model. In that sense, the content of the flow is ephemeral.

For a variety of reasons, each of those characteristics of the traditional TCP/IP model is undergoing change. As protocol layering and organization have become increasingly complex, deep packet inspection has come into wider use, directly contradicting the model that the bits are indistinguishable. It may allow for improved handling in routers and other components comprising the network infrastructure. In addition, in the 1990's the advent of the World Wide Web, and subsequent work on searching, began the change forever of the network mechanisms from being only for transport of flows of bits and toward providing access to information or objects. What has become known as middleboxes, such as firewalls and NAT boxes are intermediaries placed along the network flow with the intention of mediating and perhaps modifying the flow. This development is a clear example of the transition away from transparent transport. Multicast and the web provide two different approaches to a transformation from pairwise to multi-destination traffic (the web being multi-destination by virtue of CDNs that transform the client-server web into a multi-destination overlay with proxies and caches).

Finally, the *third evolution* is that of the web, searching, and more recently pub/sub and content-based networks, that reflect a separation in time between source and destination. The overall paradigm shift is toward separation of source and destination in both time and identification, and, perhaps more importantly, the information or content taking on a first class role in the paradigm. With this in mind, we can identify several key aspects to such a system:

- An object has a source that made it available or "published" it.
- An object has two kinds of characteristics, those that are inherent and can be considered its "type" or "nature" and metadata or other attributes of it, which are not inherent in the same way. These may reflect ownership, various timestamps such as creation or modification, usage or access policy constraints, and so forth.
- A potential "user" of an object can express an interest through the same kinds of attributes.
- There exists at least one rendezvous mechanism that provides matching between attributes and interests.

Security and Privacy in this Changing Context: The Central Role of Identity

In the changing context of the Internet as described above, we consider the challenges this raises with respect to the broad area of security in terms of the following issues:

- The nature of an object: how can one provide confidence that an object as accessed by a "user" is in fact what the "provider" intended? This is a question of the preservation of integrity of an object.
- The identity of the provider: can the provider identity be authenticated, to what degree, and how?
- The provenance of the object: where has it been? Is there some means of tracking the trajectory of the object as seen by the user?

- The identity or other validity of the user: How and to what extent can access to the object be controlled and authenticated (i.e., authorized)?

- Verification of the rendezvous: Is the rendezvous valid? Is there a match between attributes and interests, and who is validating that?

At the core of these issues and questions is a set of issues surrounding *identification* and *identity* answering the question of who and what. There are potentially different requirements on identification to address each or at least some of these and therefore possibly different approaches required. Consider, for example, access policies that may be defined for an object. They may be in terms of a set of verifiable characteristics that a user must have, but that can be resolved and validated only when access is requested, but probably need to be signed by a uniquely identifiable and authenticable author or publisher.

In examining the issues in designing identification schemes, there are a number of factors to consider. To do this, we distinguish between identifiers and the identities that they denote. Further we recognize three basic functions for which identifiers can be used: equality, access, and nature. The function of *equality* is defined as clarifying whether or not two identities are equal by some definition of equality, based on their identifiers. The function of *access* defines how an identity is utilized. Finally, *nature* provides some information about the character of an identity. This is often embodied in the character string of the identifier, which may, for example, describe the identity in some manner. With a definition of whether and how identifiers should support each function, we often begin by determining whether the identifiers should be unique over the whole set of identities or relative to some context. One can then ask whether identifiers are assigned only once or can be reused over time. One can also determine whether an identity can have more than one identifier. In each of these cases one is considering the mapping between identities and identifiers, and for different sets of identities the answers may be different, as we will see when we discuss our examples later in this paper. Finally, in considering designs for identification schemes, there are a set of mechanisms that must also be determined. Specifically these are the approach to generating and assigning identifiers and the approach to resolution or mapping, where needed. Design questions here have to do with how centralized or distributed authoritative entities, and replicated these may be.

The Intention of this Paper within this Changing Context

Because we propose an enriched information based communications paradigm for the Internet, a thorough study of identity in that context is a critical starting point for our further work, especially in the areas of security, privacy and trust in such a paradigm. For this, we will outline issues related to security and privacy at large as well as identification and identity in particular. For this, we will outline an architectural vision for the Future Internet, formulated under the observation of these evolutions. This vision will help us formulate research questions related to security and the nature of identity in the Future Internet. With this, we intend to formulate a potential research agenda, not only for the future of this working group but also beyond.

The remainder of the paper is organized as follows. We proceed with a set of use cases, including one in particular to be used for deeper examination and illumination. We then discuss architectural challenges of tussle and a broader information-based architectural vision. That is followed by a deeper examination of the application of the architecture to our demonstration use case. The paper concludes with our vision and list of key challenging questions that arise in moving toward a research agenda, including running a workshop to bring together some of the key participants in this area, especially those interested in the broad set of questions and problems derived here.

Sample Use Cases of the Future Internet

We consider many of the scenarios for the Future Internet to be truly cross-value chain type of scenarios, crossing today's relatively closed sectors of communication, content, retail, traffic and others. In all our use cases, the notion of identity (WHO and WHAT) and identifying (HOW) play an important role. We will outline this relation to identity later on while we merely present the use cases in this section.

As an example for this, consider a highway accident during rush hour involving a truck and a chemical spill. There are a variety of high level responses or functions necessary, all important. One is to clear the accident. A second is to provide medical care to any victims in the accident. A third is to alert and to remove any residents in the neighborhood who may be affected by the chemical spill. Finally, it is necessary to clear the highway of other drivers, to remove and redirect those caught in the traffic jam and to prevent others from entering it. In this example, there exists a flow of information from car manufacturers through emergency response through community services and mobile environments with different actors, different policies for retrieval and storage of information as well as interception. We will reconnect to this example later in our discussion around identities.

Other examples for use cases are:

1. News aggregation: news is all about information from trusted providers of news items (e.g., agencies). Aggregation puts forward this information according to the goals underlying the aggregation (e.g., summary, scoping, widening, …). This requires means to potentially authenticate the provider of a particular news items, often defining the trustworthiness of the particular pieces of information. The aggregated information is provided to an end user, e.g., in form of news readers.

2. Sensor data resembles the closest match to natural perception-based information consumption. Similar to naturally perceived information (such as a particular color in your surrounding environment), identifying the particular form of information source is less important rather than identifying the WHAT (i.e., the type of information) provided. Extending this perception from local to remote environments potentially adds identity requirements like location (where do I perceive?) and surrounding environmental information (who else is present?) to the use case.

3. The ‘Internet of Things’ has been touted as enabling the integration of all things around into the communication experience. RFID (see [3]) has been seen as an enabler for such kind of propositions. While typical RFID value-chain scenarios are still largely living single value chains, true cross-value chain scenario could see the usage of RFID in scenarios such well-being (e.g., using food consumption, enabled by RFID information, to recommend better lifestyles) where information flows from retail information providers (e.g., supermarkets) to health providers for recommendations to potentially search engines (for providing proper information of surrounding stores), all of these flows embedded in the privacy concerns of the individuals, regulatory requirements for the organizations (such as health providers) and ROI concerns of investing companies.

Our Architectural Vision and Implications

The Internet of today is focused on passing messages from a party A to another party B with the end points seemingly being in control of implementing the policies with respect to the delivered data, according to the End-to-End principle as postulated in [5]. This fundamental architectural principle is increasingly questioned in information-centric scenarios in which the information, not the delivering party is put in the foreground. In order to provide a better insight in the nature of a potential evolution of identity in an information-centric future of the Internet, we outline our architectural vision for such future in the following section.

From Design for Tussle to Tussle Networking

In their seminal paper on **Design for Tussle** [1], Clark et al. laid out principles for designing architectural solutions that would allow for deployment in different market settings without breaking the underlying architectural solution. *Tussle* as depicted in the paper is characterized as conflicting concerns of involved players in the particular deployment, e.g., the conflict between reception of any IP packets by any party B and the concern of corporate IT departments that such reception would lead to endangerment of corporate (IT) security. A prominent example in this paper for a typical *tussle space* is **identity** and its surrounding concerns and potential conflicts between players involved. The principles laid out in the paper are defined relatively informally, i.e., hard measures of ‘successful tussle design’ are largely missing (which is often criticized) although the separation of concerns is identified as being important, yet hard to specify.

With the Design for Tussle concept in our minds, it is interesting to investigate architectural designs in practice. When doing so, it can be recognized that a potential tussle in the marketplace finds its entry into the architectural solution that is being deployed by a set of nuts and bolts in the actual technical solution. For instance, the conflict of interest in receiving IP packets by end nodes and the desire of corporate IT department to delimit that very reception of ANY packet led to the introduction of a particular architectural element afterwards, the firewall. Generally, it can be observed that tussles are often incorporated into architectural solution after the design. But also other means, like standardization and proprietary deployment, are methods to delimit the architectural solutions towards a set of (more or less well identified) set of tussles. Examples of this can be found in the VoIP space where the 3GPP IMS (Internet Multimedia Subsystem) can be seen as a particular tussle implementation of a technically similar design, namely the IETF-based SIP VoIP.

Design for Tussle in practice leads to parallel deployments of seemingly rather similar architectures, significantly increasing costs of deployment through specifications and implementations of these solutions. But also feature interactions and architectural rigidities of these manifold solutions unnecessarily add to the complexity of the overall ICT landscape.

The Vision of Tussle Networking

Based on our discussion on Design for Tussle above, it seems to be desirable to remove the necessity of parallel architecture developments and deployments as a result of tussle delimitations and instead move towards a vision in which **the network itself acts as an execution environment for well expressed tussle policies**. These tussle policies express the concerns of the players within the particular tussle space (individuals, organizations, public bodies, …), enabling the system to resolve commencing tussles on the basis of policies not deployment of parallel architectures. *Such resolution can be based on mediation, regulation or delimitation of networking solutions within the system*. In other words, the architectural vision of tussle networking intends to shift the resolution of tussles increasingly from the design to the runtime phase, as expressed in [1].

The following picture shows a high level architecture for this vision, outlining the different planes from the policy plane above the supporting information plane down to the generic provisioning plane.

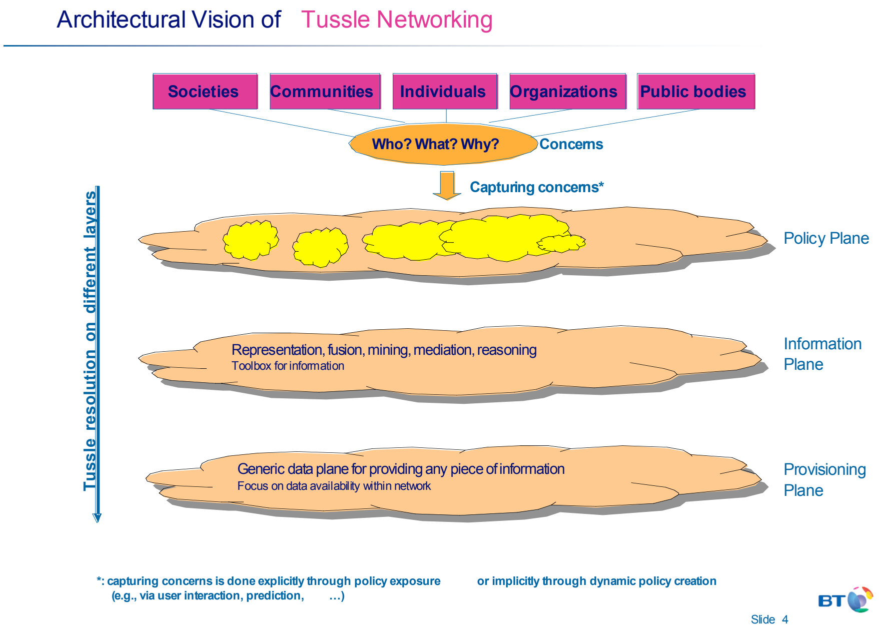

Architectural Vision of Tussle Networking

Societies | Communities | Individuals | Organizations | Public bodies

Who? What? Why?

Concerns

Capturing concerns*

Tussle resolution on different layers

Policy Plane

Representation, fusion, mining, mediation, reasoning

Toolbox for information

Information Plane

Generic data plane for providing any piece of information

Focus on data availability within network

Provisioning Plane

*: capturing concerns is done explicitly through policy exposure or implicitly through dynamic policy creation (e.g., via user interaction, prediction, …)

Slide 4

Figure 1 Architectural Vision of Tussle Networking

Policies within this architectural vision are expressed in terms of the "who", "why" and 'what' of the particular task at hand, i.e., it needs to express identities of parties involved, task descriptions and usage policies (such as retention and distribution policies). It is important to emphasize however that the policy plane in our architectural vision is not meant to be restricted to particular technological means of feeding policies into the execution environment. It is rather meant to express the necessity to capture concerns of players in general (likely through some form of policy approach). Many concerns of players in particular scenarios are given through a wide variety of means, such as written law, regulation, common sense, standardization procedures and many others. Ultimately, many if not all of these concerns are somehow to be fed into the execution environment in clearly expressed form, e.g., through standardized ontologies but also through user interfaces capturing the intention and concerns of the end user.

With this, the **policy plane** defines the particular context of a given scenario to be executed by the system, this context being represented by the set of WHO/WHAT/WHY. However, one must not constrain the policy plane as a static set of 'game rules', which instantiate at runtime. As many real-life scenarios are dynamic in nature by virtue of negotiations, a crucial aspect of the policy plane is the enablement of *negotiation* and *dynamic policy composition*. For instance, the encounter of a (policy-based) web browser with a website, both having conflicting policies, e.g., for the disclosure of some user data, might eventually end in the negotiation of some information disclosure that will result in some tangible benefit for the user. This negotiation can either be automatic or end-user feedback based.

This high level description of networking scenarios through policies translates onto an **information plane** that supports representation of, mediation between and reasoning over the information that is stemming from the policy space. Much work on the Semantic Web [4] influences our vision as formulated here. Concepts of ontologies, mediation of ontological differences and reasoning over information provided within ontological contexts are of utmost importance here. In relation to the policy plane above, the information plane can be seen as toolkit to work with information (represent, mediate, reason about and even trade) within the context as defined within the policy plane.

On the lowest level, the **provisioning plane** is seen as a generic data plane that is optimized towards availability of information throughout the entire system. It is important to note that authenticity of the information is left to the original provider of the information by means of, e.g., encryption or self-signed public key identifiers. At this layer, information is merely published by an authorized source of information and retrieved by a set of interested parties, who can then verify the authenticity of the provided data (not necessarily the sender). Related work on publish-subscribe systems, even on Internet scale, is ramping up to provide useful input into this vision.

Within this vision of tussle networking, our aforementioned example of (controlled) packet reception, i.e., the firewall in today's networks would be implemented as a runtime scenario in our vision in which the policy governing the reception of packets within a certain set of receivers would be expressed with the help of the policy plane, being enforced on the provisioning plane. Changing the policies governing the reception would constitute a dynamic reconfiguration of our receiver set, something that is rather difficult to achieve nowadays. We will continue on this aspect of *scoping* later.

What is to be identified and expressed?

In the context of this model of tussle networking with information as a first class element central to the communication paradigm, a **central question is what needs identification**

and expression, reflected in our architectural vision through the central role of WHO/WHAT/WHY within the policy plane.

This question addresses the WHO part in the policy plane as described above throughout all the planes, i.e., what are the key elements in this? Information entities, sources, brokers, aggregators, and consumers? Who of these players needs identification (and identities) on what level and in what interaction. Also the WHAT part needs to be addressed with respect to identification, e.g., the resource such as profiling data. And last but not least, the WHY part with its expression of concerns, e.g., storage and retention policy, distribution, …, needs to be addressed properly.

In this, it is interesting to postulate a potential grouping of categories with sources "posting" to those categories, possibly under some security/privacy constraints (how might this be managed, controlled, limited?) and consumers "receiving" events within these categories, when something is posted to a category in which they are interested. Other interesting questions revolve around issues as to how this might be managed, controlled, or limited.

Scoping: information rather than topology

Scoping plays a crucial role in networking today through means of topological constructs such as subnets and (network) domains. Information, however, is not necessarily bound to topological constructs such as today's subnets. For instance, the scope of today's music delivery systems is ultimately given by the often underlying DRM (Digital Rights Management) system, not the particular topological constraints of the actual delivery. Hence, the notion of **information scope** is crucial in the architectural vision of tussle networking, this scope being defined by the expressed concerns within the policy plane.

With this notion of information scope, another interesting field of investigation opens up in relation to identification and identities. For instance, what is a proper identity management structure to reflect and implement information scoping? What is the impact of information scoping on higher level on the means of identification on lower level? For these questions to arise, we need to properly reflect the concept of information scoping within the architectural vision as outlined above.

Identity Continuity

The information centrism as well as the runtime character of our architectural vision poses interesting questions on the notion of identity continuity. We can envision that forms of identity are likely to evolve from rather weak forms towards stronger forms of identity. This type of *trust gain* can already be observed in many web-based services that rely, for instance, on reputation systems to verify the strength of a provider's identity. The ability of expressing policies in our tussle networking vision is likely to accelerate such forms of continuous identity, such evolution expressed through policies.

But the question of continuity also relates to the abovementioned issue of information scoping, where continuity of identities can be reflected as a scoping of information accessible to a particular form of (continuously evolving) identities. With that, we envision the application of reputation-like mechanisms much deeper in the delivery system as we can see in today's Internet.

Impact on the Value Chain

It is apparent that our architectural vision of tussle networking will have a clear impact on the overall communication value chain. But even more so, we believe that it has an even stronger impact on related questions about the underlying economics of such a vision as well as regulation. This clearly makes our vision very appealing for a wider investigation in the context, for instance, of the Value Chain Dynamics WG in CFP.

However, value chain related questions also relate to the scope of this WG since identity is a natural tussle space and therefore any mechanism to resolve tussles naturally impacts identity-related questions on economic level, such as the appearance of new players and the mechanisms to resolve tussle in the identity space (e.g., using established real-life identification schemes or regulation).

In addition, the envisioned evolution towards an information-centric Internet is likely to open up new propositions (beyond currently existing ones) very specific to the information focus of such future, such as *information trading* or even *information banking*. This is likely to have an impact on the understanding of the nature of identity, from a technical and economic perspective.

It is therefore important that we gain a solid understanding of the identity-related issues that will impact the value chain most notably.

Returning to Our Example

In order to explore these ideas more concretely we reconnect to our example use case described earlier. From the networking perspective we can begin to consider at these high levels questions about provision of information:

- In the case of clearing the accident, it is necessary to understand what size tow-trucks might be needed, whether other heavy equipment is needed to either right or restabilize vehicles, and so forth. These determinations may be made initially by some combination of police and fire responders.

- In the second case, medical emergency workers will require information about the victims as well as the side effects of the chemicals. This may include medical records that otherwise, for legal reasons, must remain private. Information about the chemicals may be proprietary, so there may be intellectual property ownership questions. This information as well as information about local residents and means to access them will be needed for the third component.

- Finally, for the other drivers, not only must they be contacted, but they could valuably be provided with maps and directions for alternate routes. This final information may change with time as traffic is routed off the highway causing congestion in the initial alternate routes.

- Furthermore, as more drivers arrive at the highway with the intention of using it, they must be added to the set of people who need to be directed to go elsewhere and provided with alternate route information.

The above list of questions directly relates to the WHO/WHAT/WHY of our architectural vision, i.e., WHO is to be notified, WHO is notifying, WHAT is to be sent, WHAT is required to be distributed, WHY am I requesting access to particular information (in particular proprietary one)? We can assign these questions to all the above listed pieces in our example.

With this in mind, our example suggests a variety of information-based aspects:

- First, we can consider *two information flow directions*, information that is enriching the story about the current situation and the information that is needed by different parties in order to address the four components of mitigation of the situation. Information about the accident may derive from roadway monitors (e.g. cameras and other traffic sensors), cell phones (both actual calls and perhaps tracking lack of movement of phones across cells), and various other monitoring systems such as GPS navigation systems in some of the cars. This variety of capabilities represents a rather rich sensor net system, composed of several distinct sensor net systems, at least some in the commercial domain.

- In terms of *distributing information*, clearly the largest and perhaps most dynamic set of participants who need information are the drivers on or entering the highway. Again, the communications medium may be multifaceted and perhaps somewhat personalized. Clearly drivers exiting at each exit prior to the accident are likely to need directions tuned to the particular exit they take, and perhaps to their individual final destinations. The EMTs may need access to patients' medical records that must otherwise be kept completely confidential. The police and fire responders may need to know lots of detail about which sorts of equipment are available for the cleanup as well as information about the chemical and the neighborhood. One thing that is clear that both for efficiency and for policy reasons some information will be necessary for some participants and not for others.

- Thus at the *high level policies must be in place* that control and restrict which information is and is not available to whom under what conditions or for what purpose. Furthermore, the providers of communications media (cell phone, GPS systems, police, fire and EMT radio systems, WiFi and WiMax that may reach devices carried by people in the cars, etc.) may also be required to *impose usage policies*, in the context of some of them being commercial operations for which contracts for usage and payment will be necessary.

- Below the level of the policy control on flow and access of information, there will need to be the capability of *taking various low-level pieces of information*, and *analyzing sets of them* in the context of others, reasoning over them, mediating, based on the policy constraints and determining access and delivery of them. Supporting that we postulate a provisioning plane that captures, stores, identifies and makes available at a very low level of abstraction the pieces of information.

It also may reach and retrieve or collect new pieces of information as determined by the layers above it, either based on needs in the information plane (e.g. incomplete or out of date information, etc.) or the policy plane (e.g. facts that clarify the applicability of a policy).

The instantiation and role of information through these layers is perhaps reasonably clear, but the participants must also be distinguished through the layers as well. At the policy level, participants have rights and privileges or lack thereof. At the information plane level, participants have needs and capabilities to either use or produce information. At the provisioning layer, they are sources and destinations as pieces of information are either generated (sensed, learned, etc.) or delivered. They can be considered to be comparable to IP addresses or end-points at a transport layer.

We can recognize all the aspects of our architectural vision in this example due to the likely conflicts occurring and the need for real-time resolution of these conflicts. But most clearly, we can appreciate the challenge to enable a true system-crossing collaboration in real-time, requiring the mediation of potentially conflicting policies that govern the delimitation of each of these systems. Designing a system that would allow such (ad-hoc) system-crossing collaboration in real-time is the challenge, and the particular questions on the nature of identity in such system are of interest to us.

Towards a Research Agenda

As always, the way towards a vision is almost as interesting as the vision itself. In order to define a potential research agenda in this space, we first outline potential questions of relevance to the outlined vision. We then propose dedicated actions to go about some answers to these questions.

List of Potential Questions

Many research questions can be posed within the context of this paper. We restrict our list to questions directly related to identity and concerns, the main thrusts of this paper, with an attempt to group these questions accordingly. The list is not meant to be exhaustive.

The Future Nature of Identity

Most relevant to our paper are questions surrounding the future nature of identity. Some of these questions are:

- What is expressed on an entity level and what on a meta level? Can it be both? Does it depend on the usage, e.g., meta-level for service terms but entity level for actual transaction? Do certain design methodologies help, such as ontology-based design?

- How are ‘values’ represented and identified in this architectural vision? Do we need an identity framework for this or will it be inherent part of the overall design (e.g., policies implicitly reflecting individual and societal values)?

- How is the context-dependent nature of identity (who am I in what situation?) represented? Does it exclude a global solution to identity?

- What is potential layering of the identity information? How can we formalize the information layering?

- Is there a minimal form of identity? What would this identity look like? Do we need a central authority for this minimal identity?

- How to postulate groupings of information (and identities)? How do we form categories/groupings? How global should something like this be?

- How do we scope information (and therefore identity)? By virtue of topology? By virtue of publisher address dissemination to restricted set of people (i.e., key-based)?

- How can we ensure forms of identity continuity?

Architectural Questions

Apart from directly identity-related questions, there are wider architectural questions posed by our outlined vision, such as

- How to ensure a minimal data provisioning plane? What are the requirements for such generic provisioning plane? Is it some form of pubsub?

- How to scale? What identities need global scale, which ones don't? How to build scalable yet localizable system?

- How centralized or decentralized should access control authority be? Where does it sit architecturally? In discovery or provisioning?

- Where (architecturally) can we ensure forms of identity continuity?

Evolution

The outlined architectural vision potentially calls for radical changes in the current system design of the Internet. How to get from here to there poses questions such as

- How does this impact the Internet as a communications medium? How can we migrate from here to there? Can we do clean slate? Can we overlay (as the current Internet used to be)?

- Who are the future stakeholders in such information-centric world? What future role will identity play in such world?

- How does QoS look like in such world? Will today's notion of quality of service, which is largely provisioning-based (i.e., bandwidth, delay) move towards a notion of *quality of information*? What are measures for this?

Next Steps

The nature of identity has been widely discussed in the recent years, leading to various solutions in this space (such as LibertyAlliance, OpenID, …). Despite the various industrial thrusts behind all these initiatives, we fail to see clear solutions to the future question of identity, its governance and its architectural framework. The very nature of surrounding concerns adds another dimension of complexity to the quest for a global identity framework that is suitable for the Future Internet.

In order to shed some light on the question of Identity in Information Networking, we feel that it is of utmost importance to understand the shortcomings and lessons learned from previous approaches in this space. It is important to understand motivations of players in this field for attempts to establish themselves as an *identity player*. In this, expected value propositions (i.e., where is the money), the relation to context dependency and cultural difference in the identity question, and the reflection of organizational and individual concerns in the provisioning of an identity solution are examples of learnings we might want to derive. In this discussion, it is important to include identity providers that are not naturally sitting in the space of the current Internet but are likely to sit in the space of the Future Internet, namely providers of identity such as passports (i.e., governments), credit cards (i.e., finance industry) but also local institutions that today cater to the need for localized (and possibly temporary) identification of some sort.

Such understanding could potentially lead us to conclusions about future motivations to become (or not) a player in this space and therefore advise industrial and institutional partners on the future of identity.

A Proposal for Action: A Workshop on Identity for the Future Internet

In order to jumpstart this discussion and the desired insight, we propose a workshop to be held and organized by CFP which invites constituencies in this problem space to share their experiences and their expectations in this space in a neutral environment (i.e., academia environment). The workshop is guided by the context given in this whitepaper and will have a format of a discussion workshop. The objective is to summarize findings and insights from this workshop in a separate document, serving as a follow-up to this whitepaper. We will elaborate on the format of this workshop in a separate document.

Conclusions

The Internet is changing and this change affects many parts, one of which is the nature of identity and the viability of approaches to identity. Concerns of organizations and individuals are seen to become increasingly important in the nature of communication and will therefore affect any future solution in the identity space, as outlined in our architectural vision on tussle networking.

In order to understand the viability of future identity solutions, we propose a concrete step forward with the organization of a workshop, inviting current and future potential players in the identity space to engage in a dialogue that sheds light on important questions on the future of identity. The thoughts in this whitepaper serve as a guiding introductory document for this discussion, although it does not intend to restrict the discussions to the presented issues only.

Acknowledgements

Many of the ideas and concepts in this paper reflect the numerous discussions among participants in the Privacy & Security Working Group of the Communications Futures Program at MIT. Thanks needs to be given to the individuals involved in these discussions, in particular Karen Sollins (MIT), Dirk Trossen, Rob Beverly (MIT), Paul Kearney (BT), David Reed (MIT), Hannu Flinck (NSN), Jim Fenton (Cisco), …

It is to be made clear that opinions expressed in this paper and in the contributions leading up to this paper are drawn from the contributions of these individuals and do not represent official views or policies of CFP's sponsoring companies or universities.

More to Read: References

[1] D. Clark, J. Wroclawski, K. Sollins, R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet", IEEE Transaction on Networking, Vol 13, No 3, June 2005

[2] P. Mahonen, D. Papadimitrou, D. Trossen, G. Polyzos (Eds), "The Future Networked Society: A Whitepaper from the EIFFEL think tank", available at http://www.fp7-eiffel.eu/, December 2006

[3] C. Fine, N. Klym, D. Trossen, M. Tavshikar, "The Evolution of RFID Networks: The Potential for Disruptive Innovation", MIT Center for eBusiness Research Brief, Vol. VIII, March 2006

[4] T. Berners-Lee, J. Hendler, O. Lassila, "The Semantic Web", Scientific American, May 2001

[5] J. H. Saltzer, D. P. Reed, and D. D. Clark. "End-to-end arguments in system design", ACM Transactions on Computer Systems 2, 4, November 1984

[6] "OAuth: an open protocol to allow secure API authentication", available at http://oauth.net/