Arguments for an Information-Centric Internetworking Architecture

Dirk Trossen
Non-affiliated
Colchester, UK
dirk_trossen2000@yahoo.com

Mikko Särelä
Ericsson Research
Helsinki, Finland
mikko.sarela@ericsson.com

Karen Sollins
MIT CSAIL
Cambridge, MA, USA
sollins@csail.mit.edu

ABSTRACT

The current Internet architecture focuses on communicating entities, largely leaving aside the information to be ex-changed among them. However, trends in communication scenarios show that WHAT is being exchanged becoming more important than WHO are exchanging information. Van Jacobson describes this as moving from interconnecting ma-chines to interconnecting information. Any change of this part of the Internet needs argumentation as to why it should be undertaken in the first place. In this position paper, we identify four key challenges, namely information-centrism of applications, supporting and exposing tussles, increasing accountability, and addressing attention scarcity, that we believe an information-centric internetworking architecture could address better and would make changing such crucial part worthwhile. We recognize, however, that a much larger and more systematic debate for such change is needed, un-derlined by factual evidence on the gain for such change.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design – *distributed networks, network communications.*

General Terms

Design, Economics.

Keywords

Publish-subscribe, information-centric networking.

1. INTRODUCTION

Communication is often based on the sharing or exchange of information. We underline this observation with a scenario in which such exchange is critical, in order to illuminate challenges to implement such scenarios in general.

Off the west coast of the US, an oil tanker has been leaking oil after a storm. As the oil spill approaches a wildlife preserve, several loosely coordinated emergency services step in, including staff of the preserve itself, containment and oil cleanup services, marine services to clean the water, rescue teams to treat animals and birds, and medical services to monitor and address health issues of the emergency teams. The preserve staff has deployed a large-scale (information) network that includes sensors for environmental data, cartography data, information about the tasks and locations of workers, and communication and coordination tools. Using a range of networking technologies from wired to wireless to ad hoc point-to-point communications, the rescue teams need to know what others are doing, what they have learned, and where they are, as a basis for both safety and application of limited resources. They also confer with worldwide experts to draw in additional competence for dealing with the emergency. One invaluable component of the rescue operation is that visitors are also encouraged to take pictures of the troubled animals and birds, as complementary information to share them with the team through the network.

The information-centric nature of our scenario leads us to challenges that are mainly driven by human interests like avoiding information overload, ensuring security and accountability as well as creating viable socio-economic environments .By placing information at the heart of our solution, we argue that an information-centric internetworking solution is best positioned to aid addressing these challenges holistically rather then application-specifically.

The first challenge is the observation that a scenario like ours is inherently about information and its exchange, largely independent from devices and networks used for its delivery. For instance, rescue workers must be able to attend to their information from any device, either mobile or a stationary desktop, and any network, either in their office or in the field.

The second is that there are policy disagreements, in other words tussles that require both exposure and mediation. For instance, governance of information that is gathered during the emergency may conflict with the interests of the public to be informed and the rescue teams to preserve privacy and confidentiality of certain information. Another tussle may derive from the use of network resources, e.g., restricting mobile phone calls to emergency ones only, conserving radio resources in the vicinity of the emergency.

The third is increasing accountability. The privacy of the workers' medical records is a prime example in our scenario. The challenge is to make them available as needed and track where they may be stored, cached, and used, in order to monitor usage in a situation in which strong medical record management may not be feasible.

The fourth challenge is how to address the problem of attention scarcity. In our scenario, individual workers must be able to control how much of which kind of information they receive. It may be the case that an animal rescue worker might need detailed information on the anatomy of a particular kind of animal rather than being pushed a full catalogue of animals. Furthermore, unsolicited information (spam) could pose a problem in this situation through visitors uploading irrelevant or even misleading information.

In this paper, we argue that the network architecture should contribute to solving a range of human-centric information problems rather than leaving the solution to the application. We also argue that an information-centric view on internetworking is the key to effectively contributing to solutions on such low level. To ground our discussion, we briefly outline an example for an information-centric architecture. This design combines known concepts into a new internetworking layer that operates entirely on the concept of information rather than (uniquely addressed) endpoints. Although we believe these challenges are centrally important in examining the effectiveness of any architecture, we recognize that there are many other issues that must be given attention in comparing the current architecture and its potential alternatives. Hence, this paper is intended to be the beginning of an architectural discussion and in no way the final word.

2. THE ARCHITECTURAL CHALLENGES

The information centrality in our scenario is largely driven by human interests like security and economics, leading us to the challenges outlined in the introduction. We recognize the potentially many contentions in realizing these challenges but also the contention itself that these challenges ought to be addressed at internetworking level at all. We believe that this need arises when moving from interconnecting machines to interconnecting information, and it is the ability to address these challenges at this level that constitutes a major difference to the current internetworking solution. However, we recognize the many other challenges, such as scalability, security and manageability, but will defer their discussion to a deeper analysis of our solution at a later stage.

2.1 Information-Centrism of Applications

We recognize that information is at the heart of most if not all communication. Van Jacobson claims in [1] that "the overwhelming use (>99% by most measurements) of today's networks is for a machine to acquire named chunks of data (like web pages or email messages)". This is confirmed by the implementation of services like the World Wide Web, providing access to information by means of URLs, used by web servers, proxies and services such as Akamai. As suggested in forecasts such as [10], efforts to digitize material like books, journals and cartographical data as well as the advent of services like YouTube and BBC iPlayer [11], RSS feeds, podcasts and others will further increase the amount of information being produced. In addition, more and more new applications appear that focus on the provisioning of information without requiring any unique endpoint addressing of the underlying IP protocol. The rise of P2P applications is one such area. Most nodes in these scenarios reside within private IP networks. Information is disseminated in these structures without specifically addressing one of the participating nodes. Instead, a particular piece of information is requested in the hope a suitable provider is found by the (overlay) network. Diffusion methods in sensor networks are another example of a growing area in which global network location addressing is not of importance. Instead, an information-centric approach is used for disseminating the information, such as addressing based on geospatial location.

At the heart of the increasing role of information is not the mere production but the flexible and policy-compliant interconnection of information to useful services. This is usually done through a plethora of middleware techniques, which begs the question as to what generic functionality the network could provide rather than relying on many (yet similar) solutions overlaid on top of a generic bit transfer service – the current internetworking architecture. It is, thus, our view that a future network needs to provide an abstraction for applications that can be directly used to address information across various application domains.

2.2 Supporting and Exposing Tussles

Given the often adverse interests of crucial stakeholders in the Internet, conflicts of interests (or *tussles*) often occur, putting a strain on the underlying architecture in the attempt of fulfilling these adverse interests. Clark et al [13] outline the importance of defining clear tussle boundaries during architecture design for two major tussles, namely those of trust and economics. Clark and Blumenthal expand on the trust issue by postulating the *trust-to-trust* (T2T) principle [2], i.e., moving functionality to points of trustworthy implementation as a means to mediate tussles in the trust domain. There are, in fact, two forms of the trust tussle. The first is the issue of the extent to which communicants or end-points trust each other, share an interest. The second is between the end-point(s) and the infrastructure providers, reflecting the degree to which end-points and infrastructure providers must trust each other despite possibility contradictory interests. The economic tussle reflects similar contradictory interests between either communicating end-points or end-points and infrastructure providers on the level of resource usage. A future solution must provide means to mediate between existing and future tussles in a way that enables the tussles to commence without ossification of the underlying architecture through point solutions and patches.

2.3 Increasing Accountability

With the increasing information availability, problems related to data misuse, security breaches, and data loss are likely to increase. While the integration of security mechanisms can mitigate some threats, a growing effort, e.g., by Weitzner et al [15], focuses on building systems that account for the usage of information. In this approach, misuse or loss of information can be accounted for, possibly at the cost of some loss of privacy, performance, or other characteristics. The focus of this approach is not only on enabling accountability through tracing appropriate information but allowing for *posthoc consequences* with the potential to create a deterrent for unwanted behavior.

However, the proposed architecture in [15] largely leaves the actual delivery infrastructure out of scope. In other words, accountability for information traversal and delivery but also general support of the underlying delivery infrastructure is not considered. In our opinion, this is largely due to the disconnection between information semantics at the application layer and opaque data in individual (IP) packets. This disconnection places a significant burden on integrating accountability mechanisms into an overall architecture. Point solutions like deep packet inspection or lawful interception intend to restore this broken link between the actual information (semantics) and the scattered data in individual packets. However, this is achieved at a relatively high cost and is therefore only applied for particular imminent problems such as law enforcement. We recognize the ongoing debate as to whether or not a networking architecture should provide any means to ensure accountability. We believe that the outcome of this debate will be an increase of requirements to provide some sort of accountability and any future internetworking solution must be prepared for this increase of requirements. It is this view that raises this challenge to the level of an architectural one as to whether or not a different

internetworking architecture can change the tradeoff between increasingly demanded accountability and required scalability.

2.4 Addressing Information Scarcity

It is recognized that the attention of humans is limited. Focusing on what matters is crucial to avoiding information overload, potentially leading to cost, stress and mistakes being made. The increased availability of information brings with it the danger of overloading many individuals' ability to attend to the right things under specific conditions.

Mark Weiser outlined this problem in his work on Calm Computing [12]. His main conclusion was that mechanisms are required that allow end users to attend to less information (although more information might be needed to determine the right attention). This requires mechanisms to express what is 'right' under what conditions, i.e., it needs methods to express intent and concerns (represented through policies). Weiser's example at that time was the then relatively novel IP multicast, a technique, as he saw it, that greatly supports applications in providing to the enduser what really matters.

An additional challenge arises from the possibility that interests and concerns could change throughout the lifetime of the scenario, requiring flexibility in the underlying information structures that represent these interests and concerns. The attention scarcity challenge is also one of control. Consumers must be able to control their expressed interest in content and producers must be able to identify when there is no interest in their production. This requires a balance of power between producers and consumers. We believe that a future internetworking solution must provide network level techniques for implementing intents and concerns, aiding the application development rather than merely leaving it all to the applications.

3. A STRAWMAN PROPOSAL

Many relevant pieces of work exist in the area of information-centric approaches, including the recent work of Van Jacobson et al. [17]. However, none of them attempts to replace the internetworking layer as a whole. In the following section, we outline a strawman proposal of how such (information-centric) internetworking architecture could look like.

Intuitively, we start from the viewpoint that all network operations shall be based on information being the primary named entity across all layers. We believe that this aids the consistency of concepts across the layers as well as enables common cross-layer policy statements. We further expect efficiency gains in operating over a single concept, namely that of information, across all layers. With that in mind, we assume that each piece of information has a statistically unique name and that applications can request the network to deliver named information. Hence, the primary function of the network is to locate and deliver information rather than to locate hosts and arrange communications between them. In order to make the vast amount of information manageable, we introduce a concept called *scope*. From the application's perspective, a scope groups related data together. From the network's perspective, it denotes the party being responsible for locating a copy of the data. It creates a point of control, which enables e.g., access control and usage policies related to a set of data[1]. This supports composition of higher-level concepts like social networks, organizations, or cross-corporate relations (e.g., sub-contracting chains). As the underlying service model, we assume a *publish-subscribe model*, i.e., information is published by any provider while it is subscribed to by anybody who is interested in it. Data exchange only occurs when a match in information item and scope has been made.

This intuitive perception of our architecture is more concretely underlined by the following key design concepts:

A1: Everything is information: The architecture is based on information throughout all layers, including in particular the internetworking one. We define an information item as the simplest unit transmitted by the network and identify each with a rendezvous identifier (RId), a statistically globally unique identifier. Such identifier can, for instance, be created through a strong cryptographic link with an endpoint-or application-related identifier (e.g., a human-readable name) through a cryptographic hash with such higher-level identifier. The information item itself can be any array of bits including service information or other rendezvous identifiers, allowing for constructing larger information items from a collection of smaller ones. The latter enables linking between information items and introduces the concept of metadata, which in turn can be used for defining access control policies or quality of service parameters for particular information items. However, the particular usage of metadata as well the potential enforcement of policies is highly application-specific. Hence, the architecture is neutral to any semantics of the information, and only attends to the bits, the RId, and application of (network-level) meta-data.

A2: Information is scoped: Information exists in a context called scope. This concept supports grouping information that is relevant to specific application domains, as well as reducing the space to be searched for a RId and the application of access control policies enforced by the scope. It therefore supports composition of information, enabling the mapping of higher-level concepts onto the concepts of information items and scopes. In order for information to be found it must reside in at least one scope, but is not limited to only one. Scopes themselves are also information. Their RIds are called scope identifiers (SId). The data of the scope includes the network control information (such as subscriptions, policies, etc.) for the set of RIds that are assigned to the scope. While applications may attach specific semantics to scopes, only the structure of information is revealed to the network, in the form of scoping and metadata. The underlying network is agnostic of the application-level semantics and treats them all alike, as it does with all information at the base level of the architecture. Hence, it is not the goal to construct complex ontologies on internetworking level or even provide a unifying single ontology for all applications. Instead, the structuring of information items under particular application-specific semantics is enabled by virtue of a naming (A1) and structuring (A2) principle. Applications can create their own naming schemes and ontologies on top of this (information) naming structure.

A3: Equal control: Publishing information is sender-controlled while retrieval of information is receiver-controlled, provided access has been granted. Thus, communication will not take place

[1] Note that the same piece of data can exist in multiple scopes and the network system does not by default prevent a user from republishing data it has in another scope.

without both parties having agreed through a trusted party. With that, our architecture provides a balance of power between publisher and subscriber, offering a new set of network services that shifts the network from send-receive between endpoints to a publish-subscribe model of information [7][8]. In addition, endpoints do not require unique identifiers to be addressed. Instead, it is information being identified with endpoints potentially hosting this information. This enables to control the information exposure (and therefore mitigate the ability to attack) by quickly creating new information identifiers, potentially with strong cryptographic binding, in a case of attack.

3.1 Conceptual Architecture

Based on our design concepts, the following information-based architecture relies on basic labeling (cf. A1) and grouping of information (cf. A2), while providing a publish-subscribe service model (cf. A3). The main objective of the architecture is to provide the required mapping of these concepts onto concrete forwarding relations between endpoints, producing and consuming information. This keeps the network architecture simple, while enabling more complex application-level naming structures, as suggested in [17]. We can only provide a glimpse here; more can be found in [9] or similar work like in [3][6][8][17].

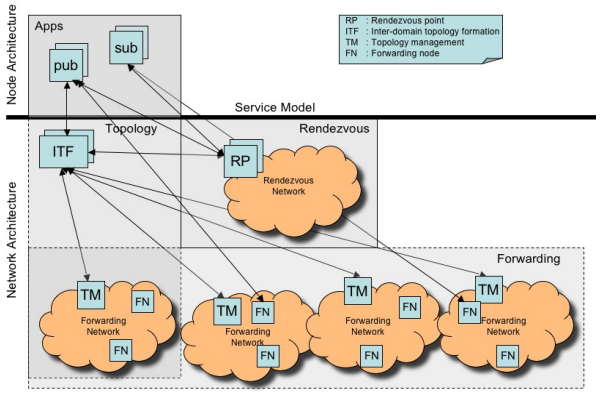

Figure 1: Conceptual Architecture

Figure 1 presents the main architectural components. The pub and sub components at the application level implement applications based on basic publish-subscribe network services, enabling publications and subscriptions towards information items within particular scopes. Transactional services, operating in request-reply mode, can easily be supported through a publish-subscribe model, with the server subscribing to receive requests. From this basic mode of communication, we can bootstrap internal network operations as well as offer a new information-centric service API, similar to [5]. Such new API replaces in many ways the role of traditional middleware layers since it conflates low-level information discovery as well as location determination of publishers and subscribers into a single network service. Therefore, the need for such mapping functions to exist on application level is largely eliminated. However, there is still a need for mapping application-level information concepts onto the basic concepts provided by our architecture.

The network architecture itself consists of three main functions, rendezvous, topology and forwarding. Generally, the rendezvous function implements the matching between publishers and subscribers of information items, each identified via a RId. Information items logically reside within at least one scope. Each scope is identified via a SId, which is in turn provided by dedicated rendezvous points. Hence, rendezvous points match the semantic-free information items within the scope they are serving. There is at least one rendezvous point per scope, each of which subscribes to the SId through a global rendezvous system. Upon subscription to an information item in the scope, the request can be routed either to all or to the 'best' rendezvous point, using anycast-like functionality. Furthermore, rendezvous points implement policies associated with the matching, such as access control.

Once the rendezvous point has matched a publication and one or more subscriptions, the forwarding topology is created in negotiation with the inter-domain topology formation (ITF) function. This is based on the publisher and subscriber "locations" on the level of autonomous systems (ASes), the applicable policies and the ITF information that includes peering and transit relationships among ASes. This is similar to BGP or (G)MPLS PCE, but the underlying networks forward information, not (opaque data) packets. Hence, there exists a rich set of policies attached to potentially every information item. Unlike BGP, this approach also allows for multiple ITF functions, each offering different sets of peering and transit opportunities that were exposed to them. This establishes the potential for peering markets, with the ITF function being similar to routing service providers [14]. Such choice is achieved by ASes publishing peering and transit relations to various ITF functions, usually constrained by policies governing these relations, while particular (sets of) ITF functions are chosen for topology formation. The choice of ITF function can either be implemented via a clearinghouse (similar to [14]) or through pre-existing knowledge, e.g., contractual business relations. The desire to separate the tussle of (policy-based) inter-domain path selection and inter-domain forwarding requires that transit ASes cannot make additional policy-based decisions on traversing packets. We believe this can be achieved by, e.g., in-packet bloom filters as proposed in [16].

In addition to building inter-domain paths between the forwarding networks to which the publisher and subscribers are attached to, appropriate intra-domain paths need to be constructed. This is done in collaboration with the topology management function that resides in every AS. This function is responsible for instructing its local forwarding nodes (FNs) to establish paths to local publishers and/or subscribers or to serve as transfer links between ASes. Publisher and subscriber locations are hereby identified as mere local *link identifiers* (themselves representing information) that only require local (AS-level) uniqueness while the inter-domain path ensures AS-level forwarding of information.[2] As in the current Internet architecture, our approach does not prescribe any particular intra-domain forwarding mechanism, with the one constraint that the local mechanisms should support the traffic policies chosen by the ITF function.

Compared to the approach taken by Van Jacobson et al. in [17], the presented strawman architecture differentiates in two crucial issues. Firstly, information is labeled with identifiers and organized into scopes, each of which is identified as information

[2] Hence, a "location" of endpoints is a stack of AS-level forwarding identifiers, created by the ITF function, together with a stack of link identifiers within ASes.

itself. This differs from the globally unique naming approach that is taken in [17] since it allows for a variety of naming approaches to be layered on top of the internetworking architecture. For instance, hierarchical naming approaches, such as proposed in [17], can be tied to the identifier-based scheme through cryptographically binding the (unique) name to an identifier via a hash function. The lack of enforcing a particular naming scheme, however, allows for implementing other application-based identification mechanism on top of the identifier scheme of our architecture, e.g., using localized identification. In addition, the scoping mechanism provides further separation and organization mechanisms within the architecture, allowing for rendezvous functions to scale to Internet size through limiting the relevant search space for information items.

Secondly, although laying our strawman solution over the current Internet is possible, it is the declared goal to finally replace the current internetworking. This leads to a focus on inter-domain functions, which is not found in previous work. For instance, the rendezvous and topology functions are particularly considered for inter-domain operation – although their detailed operation are left out due to the lack of space. Approaches like [3][6][8][17] aim either specifically at operating as an overlay or do not consider large-scale operation. The solution outlined in [17], for instance, applies a flooding mechanism for finding named content, a mechanism that is hardly scalable in any larger environment.

4. REVISITING OUR ARGUMENTS

We now return to our original challenges. We review them in the context of our proposed architecture alongside our scenario that we presented in the introduction. With that, we intend to shed some light on the question as to whether a change of the crucial internetworking function would be for the better or worse, i.e., is it worth the enormous undertaking? We appreciate that many contentious issues will remain unanswered after our presentation. We highlight some of these throughout the section.

4.1 Revisiting Information-Centrism

In our information-centric architecture, application-level information structures are reflected through simple and highly scalable structures on network level, in forms of items and scopes. The linking of information through metadata enables to reflect governance and provenance of information on low level. Furthermore, the publish-subscribe model is compatible to many event models that are currently implemented on middleware level. Additionally, the topology and forwarding functions in our architecture largely implement the location determination of the producing and consuming endpoints, without explicitly exposing this functionality to the application, making many of the usual transport mappings of middleware platforms unnecessary.

Within our scenario, we expect the application to mainly focus on the mapping of human-understandable concepts like readably names or ontologies onto our concepts of items and scopes. For instance, concepts like (rescue) organizations, location, or usage context could be mapped onto different scopes, all of which in turn are mapped into a single scope reflecting the current emergency – this scope could be dissolved after the emergency with the included scopes and items being assigned to a different purpose. Information is published within these scopes, each of which has assigned policies for granting access to other scopes or items within. Functions for locating or exchanging information are directly implemented by the internetworking architecture.

Some of the remaining contentious issues relate to the generality that can be achieved with the proposed labeling and scoping approach of our architecture. Is it really possible to map (any) application layer concept almost directly onto the structured approach of our addressing? Another issue relates to the required functionality of finding and exchanging information. Are these functions generic enough for future applications or will overlays occur in various forms, similar to today's Internet? Only an understanding of various applications can shed light on these issues.

4.2 Revisiting Supporting and Exposing Tussles

Within our architecture, we support the resolution of important economic and trust tussles through a dedicated choice of modularity of crucial network functions. This is reflected in the separation of forwarding and topology management as well as the choice of rendezvous functions.

More specifically, the tussle of economics is addressed by the forwarding fabric providing a simple and efficient (information) packet delivery service, while the inter-domain topology function constructs a policy-compliant path between the publisher and subscriber(s) across domains. These policies can be specific to possibly each scope and even information item being transmitted. For this, we utilize the concept of metadata, facilitating a low-level system of policies that can be incorporated in the matching process (e.g., for access control or pricing) as well as in the construction of forwarding paths (e.g., avoiding particular domains for security reasons). The policies are used to expose requirements for the rendezvous as well as the topology formation process.

With that, the tussle resolution related to (runtime) path selection is concentrated in the ITF function, which can take into account both operator policies as well as end-user and regulatory requirements. Furthermore, the potential for co-existing ITF functions enables the creation of a peering market with each player, addressing particular socio-economic goals of importance to the communication scenario. In our scenario, we envision the utilization of an ITF function that is specific to emergency services, i.e., restricted towards other, non-emergency, services. We also envision a scope-dependent topology formation, e.g., utilizing high priority links for emergency scopes while utilizing best effort links for, e.g., the visitor scopes and uploading of photos by bystanders. This constitutes a process of aligning incentives in runtime that is seen as crucial in our architecture.

The tussle of trust is addressed in the rendezvous point through ensuring the balance of power between publisher and subscriber by implementing proper authorization methods for the exchange of information. The choice of rendezvous points provides leverage for publishers and subscribers to ensure that the T2T principle is adhered to, i.e., non-trustworthy rendezvous points can be changed, e.g., through establishing new scopes with other providers or publishing information in existing other scopes. In our scenario, the particular scopes are administered by trusted emergency services under possibly regulated governance rules. Distrust in these services, e.g., by non-emergency providers, can be mitigated by publishing information items to emergency as well as public scopes, when publishing information.

The trust tussle with respect to resources is addressed in the topology formation function, by integrating policies in the

creation of the forwarding path. Policy violations in the forwarding space can be countered by changing ITF providers, since multiple ITF functions can be supported. In our scenario, we assume a specific ITF function for emergency services while 'public' ITF functions can be utilized for non-critical emergency services. This allows for protecting scarce resources while ensuring a public service to others.

The most contentious issue with respect to supporting tussle resolution relates to the modularity of the chosen network functions. As outlined in [13], such modularity is crucial in a 'design for tussle'. We appreciate that a proper understanding of such modularity is not an easy task and difficult even for our current Internet. However, progress in analyzing the workings of the current Internet, e.g., through applying algebraic methods [18] to large-scale networking systems, promises to provide such insight. Since such understanding is necessary both in the current and any future networking architecture, this cannot be seen as particular for our proposal. Work is currently ongoing to analyze the chosen modularity in our design with respect to its ability to mediate a variety of potential tussles.

4.3 Revisiting Increased Accountability

In our architecture, we address the problem of accountability through providing a network where information itself as well as its structures is better visible throughout the delivery infrastructure – without relying on the knowledge of the particular application semantics at hand. Hence, rather than providing an opaque bit transfer service, our architecture realizes an information provisioning service, exposing the underlying structure of the provisioned information without revealing its semantics. Furthermore, the notion of scopes and the introduced ability to provide composition on lowest level enables the identification of single entities as well as organizations. With this in mind, we believe that the inclusion of accountability mechanisms in the actual delivery architecture is not only easier to realize but can also be done on a more architectural level rather than through point solutions like DPI. In our scenario, we envision possible policy enforcement points for utilizing particular resources like wireless links only through authorized entities. Our proposed approach to topology creation based on policies expressed as metadata provides usage and access accountability. This accountability is directly tied to the information items and scopes of the scenario, which in turn are tied to human understandable concepts, like names or context, through appropriate application layer techniques. Hence, accountability can be achieved regarding utilizing critical resources, like links or storage, in critical scenarios, here an emergency.

As pointed out in Section II, accountability needs a proper tradeoff against scalability of solutions for this problem. Accountability requires storing information about information (and its usage), which comes at an additional cost. However, we believe that the identification structures utilized by our architecture will allow for more pointed solutions rather than blindly applying, e.g., DPI techniques, in order to look for the needle in the haystack. Hence, we expect our architecture to change the tradeoff between required accountability and scalability through the information structures embedded into the network. But only careful study of accountability frameworks will solve this contention.

4.4 Revisiting Attention Scarcity

Within our information-centric architecture, we provide mechanisms to define intent and concerns through our information as well as service model. Thus, it becomes possible to abstract from simple bits to simple forms of information already on the network level, eventually enabling a layering of abstraction that can directly support the application in 'tuning in and out' these information structures by subscribing to (or unsubscribing from) relevant information. In case these structures change at application level, the direct connection to internetworking information concepts allows for a network-level reconfiguration to happen almost simultaneously, without the need to, e.g., establish tunnels. In addition, the networking functions to match publishers and subscribers as well as to construct policy-compliant topologies can be quickly re-configured through assigning new metadata to the information items and scopes. Last but not least, the publish-subscribe service model allows for shifting attention quickly through ceasing subscriptions or publications, i.e., avoiding information overload already at the network level.

In our scenario, these methods can be utilized to provide information to the rescue teams involved by organizing items and scopes alongside the application concepts implemented on top of them, e.g., a particular scope can relate to report of endangered animals. Rescue workers can utilize the pubsub service model akin to tuning into a radio station. This is similar to Mark Weiser's recognition of utilizing IP multicast for shifting focus of attention. In our case, however, this is implemented over a rich information structure rather than a mere multicast addressing scheme.

We recognize the significant contention regarding the role of the network in supporting applications addressing attention scarcity. Many believe that such functions should be left to the applications. However, we see such views as being largely driven by the endpoint-centric nature of today's Internet, where any support for applications would result in essentially embedding application functionality into the network. In our architecture, however, we see an alignment of application goals and network functionality as being easier to achieve due to the alignment of (information) concepts on both levels and the layering of abstraction enabled by this. Such alignment does not require embedding application semantic into the network. With that in mind, we believe that methods for addressing attention scarcity do have a place in a solution like ours. But only demonstrations of its usefulness in future applications will help resolving this contention.

5. NEED FOR CONTINUED DEBATE

We recognize that this paper can only be the beginning of a much larger debate on crucial challenges for a new internetworking architecture. Given the impact that any change of such crucial function would bring about and the position that the Internet has in our society, a purely technological discussion however cannot suffice. Hence, we must highlight end-user, economic as well as technology perspectives.

On the technological side, there are numerous challenges to be addressed. These include (1) *scalability*, e.g., the challenges of the rendezvous and inter-domain topology formation functions; (2) *security*, e.g., the challenges of potential new attack and threat models; (3) *impacts on future application development*, e.g., with respect to new service models or the need for new node

architectures, and (4) *manageability*, i.e., the potentially novel management principles for information-centric networking.

On the socio-economic side, many open questions remain. We highlight two of them here. The first relates to the *creation of future markets*. Any proposal for a new internetworking architecture will have to take into account economic realities, such as existing incentive structures [4]. It can be expected, however, that a wider deployment of any solution will gradually create a force of change that will transform current markets and create new ones. Such change will in turn have an impact on the (technology) solutions themselves, potentially altering their (original) viability. One example is the design of a global rendezvous solution, establishing a market for discovering information items. The relationships between localized rendezvous providers and global interconnection providers need to be understood in order to address scalability issues that might arise within the technical solution but also to identify potential opportunities for existing and new market players. *Migration strategies* are another area to be developed. Apart from migration issues on, e.g., the service level, many non-technological issues are at the heart of such strategies, such as the impact on privacy and governance, questioning current practice in these areas.

The overarching challenge of this paper is to begin the argumentation for changing a crucial part of our current Internet towards an architecture that takes information as the central entity of communication. However, these considerations can only be the beginning of such debate since we neither claim to have found the conclusive set of arguments nor do we claim to have found the compelling analytical evidence to support all of our arguments.

6. ACKNOWLEDGEMENTS

The design considerations presented in this paper have been the result of intensive discussions among the researchers in the EU FP7 PSIRP project (http://www.psirp.org), in particular within the architecture and implementation teams.

7. REFERENCES

[1] Van Jacobson, “A new way to look at networking”, at http://video.google.com/videoplay?docid=-6972678839686672840

[2] D. Clark, M. Blumenthal, “The End-to-End Argument and Application Design: The Role of Trust,” Conference on Communication, Information, and Internet Policy (TPRC), 2007

[3] T. Koponen et al., “A Data-Oriented Network Architecture”, ACM SIGCOMM, 2007

[4] J. Rajahalme, M. Särelä, P. Nikander, S. Tarkoma., “Incentive-Compatible Caching and Peering in Data-Oriented Networks”, ReArch08 workshop at ACM CoNext conference, 2008

[5] M. Demmer, K. Fall, T. Koponen, S. Shenker., “Towards a modern communications API”, 6th ACM SIGCOMM Workshop on Hot Topics in Networks (HotNetsVI), 2007

[6] S. Ratnasamy, M. Handley, R. Karp, S. Shenker, “Application-level Multicast Using Content-addressable Networks,” Lecture Notes in Computer Science 2233, pp. 14-29, 2001

[7] P. T. Eugster, P. A. Felber, R. Guerraoui, A.-M. Kermarrec, ”The many faces of publish/subscribe”, ACM Comput. Surv., 35(2):114–131, 2003

[8] D. Rosenblum, “A Tour of Siena, an Interoperability Infrastructure for Internet-scale Distributed Architectures,” Ground System Architectures Workshop, 2001

[9] D. Trossen (ed), “Architecture Definition, Components Descriptions and Requirements”, at http://www.psirp.org, 2009

[10] IBM, “Managing the Explosion of Information”, Report, 2008

[11] BBC iPlayer, available at http://www.bbc.com/iplayer, 2008

[12] M. Weiser, J. S. Brown, ’’Designing Calm Technology“, at http://www.ubiq.com/weiser/calmtech/calmtech.htm

[13] D. Clark, J. Wroclawski, K. R. Sollins, R. Braden, “Tussle in Cyberspace: Defining Tomorrow’s Internet,” in Proc. ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, 2002

[14] K. Lakshminarayanan, I. Stoica, S. Shenker, “Routing as a Service”, UCB/CSD-04-1327, UC Berkeley, 2004

[15] [15] D. Weitzner, H. Abelson, T. Berners-Lee, J. Feigenbaum, J. Hendler, , G. J. Sussman, “Information Accountability“, TR-2007-034, MIT CSAIL, June 2007

[16] P. Jokela, A. Zahemszky, S. Arianfar, P. Nikander, C. Esteve, “LIPSIN: Line speed Publish/Subscribe Inter-Networking”, Proceedings of ACM SIGCOMM, August 2009

[17] V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs, R. Braynard, “Networking Named Context“, Proceedings of ACM CoNext conference, December 2009

[18] C. Chau, R. Gibbens, T. G. Griffin, “Towards a Unified Theo-ry of Policy-Based Routing“, Proceedings of INFOCOM, 2006