From Design for Tussle to Tussle Networking:

PSIRP Vision and Use Cases

TR08-0001

Document Properties:

Title of Contract	Publish-Subscribe Internet Routing Paradigm
Acronym	PSIRP
Contract Number	FP7-INFSO-ICT 216173
Start date of the project	1.1.2008
Duration	30 months, until 30.6.2010
Document Title:	PSIRP Vision and Use Cases
Date of preparation	23.05.2008
Author(s)	Dirk Trossen, Janne Tuononen, George Xylomenos, Mikko Sarela, Andras Zahemszky, Pekka Nikander, Teemu Rinta-aho
Responsible of the deliverable	Dirk Trossen
	Phone: +447918711695
	Email: dirk.trossen@bt.com
Target Dissemination Level:	PU
Status of the Document:	Final
Version	1.00
Document location	http://www.psirp.org/
Project web site	http://www.psirp.org/

SEVENTH FRAMEWORK PROGRAMME

European Commission
Information Society and Media

Production Properties:

Reviewed by:	Arto Karila, Janne Riihijarvi, Jarno Rajahalme, George Xylomenos, George Polyzos

Revision History:

Revision	Date	Issued by	Description
0.01	2008-4-22	Dirk T	First version from Wiki
0.02	2008-4-24	Dirk T	Re-formatted use cases and started expanding vision views
0.03	2008-5-6	Dirk T	Embedded comments received from Pekka N., Janne R. and Arto K.
0.04	2008-5-9	Dirk T	Embedded comments by George X., finalized editing and format
0.05	2008-5-16	Dirk T	Embedded comment by Jarno R. & George P.
1.00	2008-5-23	Dirk T	Changed to Final after PMC approval

This document has been produced in the context of the PSIRP Project. The PSIRP Project is part of the European Community's Seventh Framework Program for research and is as such funded by the European Commission.

Table of Contents

1 Background

The PSIRP (Publish-Subscribe Internetworking Routing Paradigm) project [1] is an EU FP7 funded project with a 30 months lifetime. Its ambition is to investigate major changes to the current IP layer, up to the point of replacing this layer with a new form of internetworking. To this end, PSIRP undergoes all phases of a *clean-slate design project*, from state-of-the-art survey over outlining basic design principles and understanding design choices through the definition of conceptual and actual architectures and their implementation. Architectural and technological choices are evaluated from the angles of security, socio-economic and quantitative design constraints. Results in the form of architecture design documents, evaluation techniques and figures of merit as well as actual running code will be made available to the public in an open manner (all major deliverables of PSIRP are open) in order to invite and foster future innovation on this potential new foundation of the Future Internet.

2 Objective of this Document

The ambitions of the PSIRP project are high. The intention of re-designing the fundamental principles of the Internet can quickly lead to loosing sight as to why we strive for such a goal. Proper argumentation is required that explains to ourselves and to outside observers (reviewers, other scientists or just leisure readers) why such endeavour is not only worthwhile undertaking but also critical for the evolution of the Internet. This document aims to provide part of this argumentation by presenting a vision for PSIRP along with use cases that reflect this vision and the life it might enable in the future. It is important to understand that, in the long run, the vision alone cannot provide sufficient argumentation needed to motivate the changes that PSIRP (and similar efforts in this space) will be striving to achieve. But it can outline a potential picture of the future that might create the desire for the Internet to evolve via such changes. However, other forms of argumentation will be required to amend the vision laid out in this document. We intend to follow suit with additional argumentation that looks more closely at today's systems and their inadequacy to implement our vision. This will be the thrust of another PSIRP paper to follow.

Our vision is presented from two angles. We start off with a user-centric view of how this vision could manifest itself. For a more scientific person, this view might read odd, futuristic, or even unrealistic and utopian. But one has to keep in mind the potential view an end user might have at what we are aiming. To complement this end user view, the document also outlines a more technical (or scientific) vision that is based on state-of-the-art discussions on system and architecture design. It formulates our visionary goal in relation to these discussions and will therefore serve as a technical counterbalance to our user-centric perspective.

Visions are not conceived in a vacuum and we are well aware of this. For that reason, the document presents relations to other visions, brought forward in the past and still influencing contemporary areas of research. We recognize that we cannot be exhaustive in listing all the previous great visions that called for great changes. But we do outline some major ones.

As mentioned before, we complement our vision by presenting a story of future life within such vision. These use cases are purposefully written as snapshots of the life of end users in order to bring forward the end user view of our vision. The more technical reader is asked to attempt to recognize our technical vision in these use cases. When presenting these short stories, we amend them with short background and technical essence information that might help the reader understand some of the assumptions made for that particular part of the story.

The intended audience for this paper is twofold. Firstly, this document is meant to be inward facing, i.e., towards the consortium itself, so as to clearly motivate what we are doing throughout the project. Secondly, the document intends to address the external reader, technical or not, in order to understand what the project envisions realizing in the long run. This is the reason that this document has been released to the public.

3 The Vision of Tussle Networking

In this section, we present our vision of tussle networking from two angles, giving a user-oriented and a scientific view.

3.1 A User View

3.1.1 From Breaking...

Have we not all experienced this? We want to show these great new pictures from our last holiday to a friend we are visiting. We think that taking our own laptop is already better than relying on the right configuration on his. But still, after joining his WLAN, our laptop won't connect to the newly set up media server in our home network since the connection is blocked by our standard home router security policy. Although we might know the solution to this particular problem, such lack of adaptability to our changing needs often becomes the reason for a technical solution that seemingly worked well in one situation (i.e., within our own WLAN in the above example) to utterly fail in another. And even more so, most people do not know the particular reason and even less so the solution to the problem at hand.

Take another real-life example. One project member's sister recently complained that her WLAN had ceased to work properly after having worked for 1.5 years without any problems. Consequently, she had resigned to use her laptop at the living room floor, right next to the ADSL modem. The problem here appeared to be a standard one: a neighbour had just recently bought a WLAN access point, both of the access points used the default channel 11 and neither worked nicely. Simply changing the channel fixed the situation, but even the idea that there might be such problems is well beyond the imagination of a typical user.

In general, systems mostly break because their designers did not consider particular concerns and needs of individuals, organizations or even entire societies in their designs. Working around this lack of consideration might be possible albeit frustrating, but often it simply is impossible. This often leads to end users changing **their** behaviour in the light of system shortcomings (in the above examples, it is advisable to carry your memory stick around or use a wireline connection instead) rather than the other way around, i.e., humans adapt where their artefacts cannot.

3.1.2 ...To Adapting

Imagine a system that is designed around the ability to adapt its appearance under the changing social needs and concerns of its actors!

In other words, imagine a system that is designed to work in ways similar to how societies themselves work.

While this may sound futuristic, even utopian, we do believe that it deserves a closer look. Social structures inherently express notions of **who**, **what** and **why** with respect to actions and tasks associated with the people and organizations that are immersed within these structures. Task- or goal-oriented computing has been chartered for a while to program solutions under certain given constraints. However, we assert that the underlying structures of networks and systems only inadequately reflect these notions of who, what and why in their present design and deployment. We further assert that pushing the notions of *concerns* and *needs* deep into the system will enable us to adapt the system under the changes that might occur in the future.

These changes are executed in **runtime** (i.e., upon need) rather than at **design time** (i.e., upon perception of the need through the designer) of the system. Moving the ability to adapt from design to runtime would introduce the required flexibility since it would minimize the dependency on the foresight of the system designer and maximize the ability of a given system to keep working under changed constraints.

Hence, envisioning a system that enables the dynamic adaptation to evolving concerns and needs of the users in runtime is our aspiration in this project.

We believe that such system-level ability of adaptation will eventually be enabled through embedding a plethora of intelligence, implemented in agents acting not only on behalf of the end user but also interacting with him. While some of this intelligence will need to be deeply embedded into the system, the underlying system also needs to provide the necessary *hooks* for such intelligence to be able to perform the required adaptation. With this, we envision that the resulting system and network design will accommodate virtually any social structure (and therefore business structure as well) without breaking the basics of the design. It will allow for executing these changes in runtime, within a single architecture, rather than requiring the design of a multiplicity of alternatives.

3.2 A More Scientific View

3.2.1 Starting with Concerns

Our starting point is again end-user centric. To us, any particular communication scenario is essentially a form of production, dissemination and/or consumption of information. This communication is immersed within the users' and society's concerns surrounding the information and tasks embedded in this information. These concerns include, for example, the following:

- Who to share information with?
- Where to deliver information to?
- What to receive in return?
- What is the information used for? How long and where is it stored?
- How to receive exactly what I want and need for my purpose?
- How to choose community members or explicit business partners?

The concerns of different parties, such as individuals, communities, organizations, or societies within a given set of communication experiences are often in conflict, e.g., lawful interception vs. privacy. This leads to conflicts (**tussles**) between the parties involved.

Our desire is also to consider many of these questions deeply from a technical point of view. For example, how to define the "who" (to share with) or "where" (to deliver to) in technical terms? In today's Internet, the IP address and associated domain name, or extension thereof such as an URL or e-mail address, seem like the only real means of identifying both actors (who) and locations (where). In future networks, such rather network-oriented mechanisms seem inadequate if we want to properly support intelligence and intentionality. It seems desirable to base future networking solutions to better reflect the intentions of their users, and not the other way around, as it is the case today.

3.2.2 From Designing Systems for Tussle…

Tussles already exist in today's end-point oriented communication paradigm, as exemplified by both voice telephony and the Internet. For instance, emergency call (911 in the US) regulation is in conflict with the operator's desire to reduce costs, while lawful interception stands in conflict with individuals' right for privacy.

We assert that the increasing push towards information networking, i.e., application scenarios centred around information dissemination such as IPTV, Web usage, sensor networks, and many others, will increase the number of tussles due to the attached information and its surrounding tussles, such as digital rights management (DRM), privacy, security, controlled sharing of confidential information, and many more.

It can be observed that, for the resolution of tussles foreseen by designers of a system, parallel architectures and solutions are developed and deployed, inherently implementing a specific set of tussle constraints. As a consequence, in many deployments today the architecture itself reflects the tussles of the scenario the architecture was designed for (and therefore the involved parties).

Design for Tussle [2] offered a novel insight for the proper design of such systems, i.e., providing guidance as to how a system ought to be designed so as to withstand and even incorporate a wide range of tussles. The authors of [2] laid out principles for architecture design that would allow for tussles to commence within the particular architectural solution. It is therefore a rather architecture-driven view on the problem. The principles were held relatively informal in that hard metrics for *successful tussle design* were missing; a situation often criticized when referring to Design for Tussle. It was further recommended that a separation of concerns should be possible although it is hard to specify how to achieve this. The authors did separate the problem into design and runtime phases, i.e., resolving tussles at design time (through means of standardization and proper requirements engineering) and at runtime, although the latter was not underpinned by any particular recommendations on the architectural level.

Given the Design for Tussle methodology, we can observe the following reality in system design and development:

- Today, a potential tussle in the marketplace finds its entry into the architectural solution mostly at design phase, if even then, and is typically represented more as a *set of nuts and bolts* in a technical solution, e.g., firewalls, rather than as a conceptual solution, e.g., middlebox control.
- Tussles that become apparent after the original design are often incorporated into the architectural solution by means of evolutionary extensions to the original design, using the same design means, e.g., standardization. In other words, tussles that the designers and standards holders of an architecture fail to recognize upfront often seep into the architecture through incremental, non-standardised features or network elements, leading to gradual erosion and eventual ossification of the architecture; e.g., consider the history of Network Address Translation (NAT).

This reality results in the appearance of multiple parallel architectural deployments, each of which represents a set of tussles for the set of parties that were involved in the design of the particular architecture. An example of this is the area of VoIP deployments with parallel IETF SIP compliant deployments and 3GPP deployment, both underpinned by a highly similar conceptual architecture, yet implemented in parallel deployment solutions.

The major issues raised by this reality in system design are:

- Architectural rigidities through feature interactions of the “nuts & bolts” type of extensions occur increasingly.
- Parallel architectures are difficult to traverse, leading to breaking scenarios as presented in the first example of Section 3.1.1.
- Parallel architectures are costly to maintain. For instance, developing and maintaining two VoIP stacks (catering for two different markets) can be costly for device manufacturers.
- Parallel architectures are costly to design and deploy due to the increased costs for standardization and development.
- Parallel architectures typically lead to diminished networking externalities, thereby reducing welfare compared to a system with a single common architecture.
- Parallel architectures are often specific to certain business cases (e.g., 3GPP) with little flexibility to changes. Should such changes occur, the architectural solution can

often prove to be costly or impossible to modify for the new conditions (the authors in [2] refer to the ability of a solution to bend under changing conditions).

3.2.3 …to Tussle Networking

Given this state of affairs in architecture development and deployment, it seems desirable to remove the need for parallel architecture developments and deployments as a result of tussle delimitation, moving instead towards a design of the network itself as an execution environment for *tussle policies*. These tussle policies express the concerns of the players within a particular tussle space (individuals, organizations, public bodies, etc.). Within such an execution model, tussles are resolved on the basis of runtime policies, avoiding the need for design and deployment of parallel architectures, i.e., *the resolution of tussles moves from the currently predominant design phase to the runtime phase*. It is worth noting that a vision for a single execution environment might sound rather far fetched and unrealistic. But even a small reduction of the number of potentially emerging parallel architectures through increased runtime ability to adapt to new tussles is likely to yield some improvement over the current situation.

Attempting to formulate an architectural vision for such a runtime-driven tussle environment, we begin from our starting point in Section 3.2.1, i.e., information and its surrounding concerns. With that in mind, our vision (as shown in Figure 1) assumes a basic plane for provisioning information at a large scale, delimited by the surrounding concerns of the particular communication scenario (bottom plane of our architectural vision).

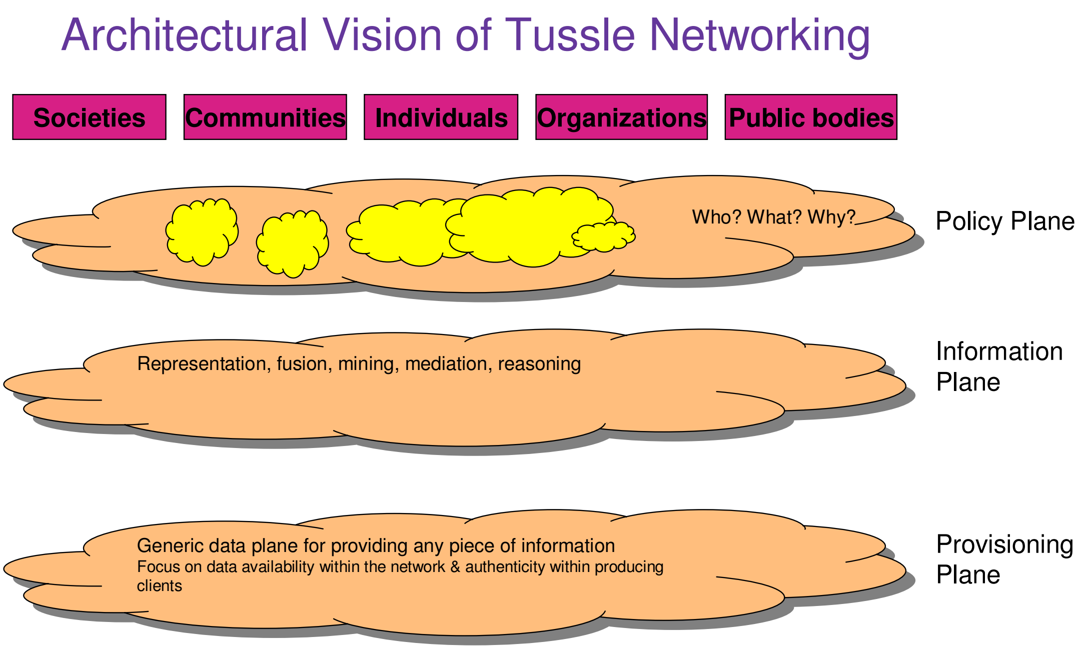

Figure 1 Architectural Vision

It further assumes another plane dedicated to information representation, mediation, reasoning and mining, as a kind of toolbox for building information systems (middle plane in Figure 1). The (upper in Figure 1) policy plane is envisioned as the one putting the communication scenario in the required context of concerns, expressed as policies that govern the communication scenario. Means for negotiating and mediating conflicting concerns are envisioned. This can lead to the isolation or partial overlap of communication scenarios, depending on the given set of concerns that governs them.

Many parts of our architectural vision have their counterparts in contemporary fields of research. Figure 2 illustrates how work on the Semantic Web [3] or policy runtimes, as well as in agent technology fit into the architectural implementation of our vision. While the provisioning plane could conceivably be considered as already being in place in the form of today's (IP-based) Internet, we see a publish/subscribe internetworking based plane as more suitable to implement our vision, not least due to its more apparent information-centrism. However, this claim requires further investigation and argumentation beyond the scope of this paper.

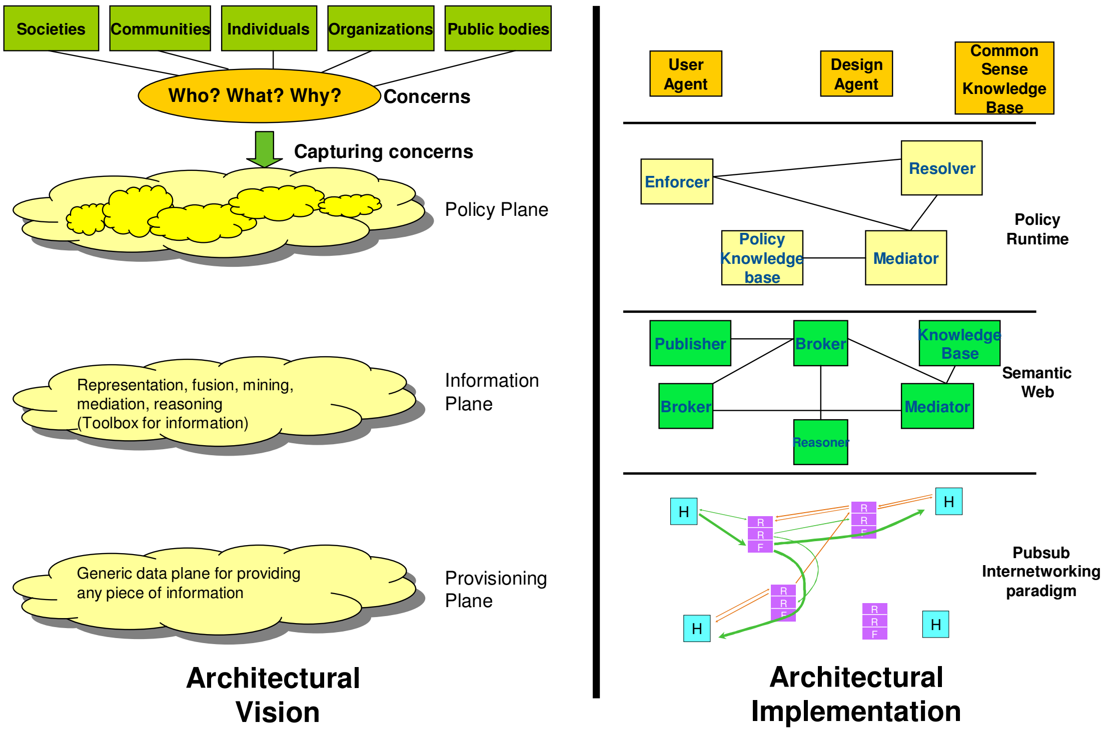

Figure 2 From Vision to Implementation

3.3 Relation to Other Visions

There are several other visions that we can relate to. The most obviously related vision is that of *Design for Tussle* [2]. As discussed above, this vision outlines a design process for system solutions in which the solution is built for a variety of potential conflicts (tussles) that can occur in a system deployment due to conflicting interests of the actors involved. With this in mind, we can see our own vision as an extreme solution under the vision of Design for Tussle; that is a vision to resolve **any** tussle within our execution environment without requiring a re-design of the system. This relation to the Design for Tussle vision determined the name for our vision, i.e., *Tussle Networking* [4].

A more subtle relation exists to Marc Weiser's vision of *Calm Computing* [5][6]. With concerns and the resolution of conflicts rooted in these concerns being the foundation of the tussle networking vision, we envision a system improving end users focus on what matters (the *calm* part of Weiser's vision). With this, we see our vision directly addressing the problem of *attention scarcity*, i.e., helping end users to cope with their limited attention span; a problem which cannot be solved by merely increasing the availability of information (i.e., solving the information scarcity problem). While we can claim that the current Internet solves many of the information availability problems that we saw before the take-up of the current Internet, our vision provides the next step towards allowing for more focus towards what matters, along the vision of Marc Weiser, by deeply embedding the concerns of end users (i.e., what matters) into the way the system is designed and (re-) configures itself.

3.4 Scope within our Work

The aspirations outlined above **go far beyond the scope of a single project like PSIRP**, as also argued in [3] and [7] as well as being illustrated in our architectural implementation of Figure 2. It includes technical aspects around policies, HCI, information representation & reasoning, mediation but in particular also aspects in economics and regulation.

The scope of this project is focused, within our vision, in the provisioning plane (see Figure 1) and therefore concentrates our efforts on the aspects that relate to the more traditional "routing and forwarding fabric" under the light of this vision. The current model of IP used in the Internet establishes a routing fabric centred around a topological notion of the network, in which packets are delivered end-to-end between two explicitly named endpoints. Changes in this topology as well as changes in forwarding behaviour are difficult, albeit not impossible (e.g., firewalls). This leads to inflexibility not only on technical (see the photo example above) but also on business (e.g., inter-provider peering) level. Hence, the ability to adapt to changing needs on the technical, social and business levels is limited and usually achieved with a more evolutionary extension model rather than with a fundamental conceptual model for an adaptive network.

> We aspire to change the routing and forwarding fabric of the global internetwork so as to operate entirely based on the notion of **information** (associated with a notion of *labels* to operate the fabric on) and its surrounding **concerns**, explicitly defining the **scope** of the information. While we do not embed the higher level semantics of information into the network, we intend to devise means that will enable the higher levels to embed concerns and social structures surrounding this information deeply within the architecture. This will be reflected by selected work items in our project such as forwarding and rendezvous.

It is worthwhile pointing out that the envisioned operation on **information**, as formulated above, indicates an important move away from the current **endpoint-centric** (data) networking model. While a model such as in IP networking enables a stream of **data** between two explicitly addressed endpoints (with total transparency as to the information represented in this sent data and the communication surrounding this exchange of data), the model envisioned in our aspiration elevates information onto the level of a *first class citizen* in the sense that data pieces are explicitly addressed and therefore directly embedded into the network, unlike in today's IP-like networks. This explicit addressing of data instead of identifying endpoints consuming this data can be seen as the most significant step on this level of the overall architecture of moving from a data to an information view, i.e., establishing correlations between data through explicit (data) addressing. This is aligned with, e.g., Van Jacobson's view (as formulated in [8]) on the evolution of communication from connecting *wires* (public telephone system) over connecting *endpoints* (current Internet) to connecting *information* (our aspiration above).

4 Use Cases

The following use cases are intended to reflect the benefits that our formulated vision could potentially bring to us and the society as a whole. It is to be kept in mind that the use cases are not about painting a shiny picture of our future or to outline all the possibilities that our vision and its realization could bring to us. We are neither in the position to correctly predict this nor do we intend to. Instead, the use cases are about bringing out the essence of *crucial aspects* of our vision. This should be kept in mind when reading through the text.

In order to better outline these aspects and their benefits, the use cases are presented as a **story** (the main text) in the life of John, Mary and their family, from a user-centric perspective. In addition to this view, each story section is annotated with some **background** (presented as a white box next to the story), such as assumptions on future developments of hardware and other accompanying factors. Further, the **essence** (presented as a light grey box next to the story) of the case is pulled out. This essence is more technical and relates to the particular part of our vision showcased by the story of the case.

4.1 The Network as My Memory

It's year 2013, John and Mary are in the middle of their two weeks hiking holiday in the wilds of Kuusamo in the northern part of Finland. John has brought with him his new multipurpose PDA (personal digital assistant), *Sparrow 2013*, which is equipped with all the latest innovations on PDA: display and wireless/wired communication technologies, such as full high-definition (HD) digital video camera, hundreds of GBs flash memory and a multi-access radio module supporting radio access technologies beyond current 3G, such as long-term evolution (LTE) and new forms of WiFi (e.g., 802.11n).

John and Mary are really amazed by the beauty of the nature surrounding them. In order to be able to record the memories of the trip both for themselves and for their friends and relatives, John has been recording lots of HD video material with his Sparrow. The device has recorded dozens of hours during their walk, since the camera also has a very efficient image stabilizer. Since John is concerned about loosing the recorded videos, e.g., through damaging the Sparrow during the hikes, he configured the device so that it will automatically establish a connection to the Internet and upload the data to the mass storage system provided by the network. If a connection cannot be established the device will notify John about it on the contact lens display.

> We assume a future in which network bandwidth is generally available in abundance (either fixed or mobile) and computing resources have reached a point where processing and memory is generally ubiquitous.

> End users will be able to store and access their data **anywhere** in the network without the need for purpose-built storage server farms. *The network therefore becomes a resilient mass storage* in itself, being able to reconfigure its operation based on the changing needs of end users (e.g., transferring heavily used data to physically closer data storage or spreading data more widely for increased security). With this, end users are able to exploit the available mass memory anywhere in the world and enjoy the benefits of resilience, which brings increased security for their precious data. End users can also sign up some of their equipment's capacity to contribute to this ubiquitous memory.
>
> With that, we can even envision that the current trend of increased mobile device memory for storage purposes will change to either not requiring any dedicated storage on mobile devices at all (therefore driving down the prices for such devices) or morphing mobile devices into pure consumption and/or recording devices (such as along the lines of Gordon Bell's LifeBits efforts). Any persistent storage will be provided by the network (and its attached devices).

4.2 My Memory Under My Control

John and Mary have returned home from their holiday trip. A week later, John sits at his home office desk and connects his Sparrow 2013 to his local ISP access network, which provides a 2 Gb/s Ethernet connection. What John wants to do is to edit the material he has filmed during the holiday. According to his calculations there

> Similar to case 1, we assume ubiquitous access to the network through any form of device. Devices can be mobile, personally owned or borrowed from friends (e.g., for demonstration purposes).

should be several hundred gigabytes of compressed raw material. His idea is to compile a 1-2 hours HD streaming video for his friends and relatives to be watched from his personal video blog. Before he can start editing he needs to retrieve the raw data, which is accessible only to him due to the original “personal” access scope that John attached to the raw material.

Editing goes by swiftly, and the 2 hours HD video of their holiday in Kuusamo is soon ready. Now John needs to publish it somehow so that his friends can watch it, if they like. He therefore publishes the edited video with an appropriate *friends* scope.

A couple of days go by and John starts to get lots of voice, video and text messages from his friends, many of them containing a pointer to a certain part (subpublication) in his new holiday HD video publication. The messages typically state something like “Did you watch this yourself?” When John follows the link, a stream of spectacular events start to unravel: while hiking John had accidentally managed to film the big meteorite that fell down on the peninsula of Kuola in the northwest part of Russia (no casualties other than some pines and lemmings) and his footage happens to be the only existing recording of this event that the people in the community are aware of. Therefore they urge John to make his publication available with a wider scope. John understands the importance and value of his “lucky shot”, so he updates the scope for the specific subpublication, containing the meteorite fall scene, to everyone and sends a notification of the link from his blog, which as metadata containing the news about the meteorite fall, to the BBC news hint portal that he is subscribed to.

My memory is an essential part of my life. It stores experiences such as photos, music, personal data, even dynamic data such as context information. **Controlled access to my memory** is crucial, in particular in conjunction with use case 1 which envisions the network itself as being my memory, therefore entrusting my data to machines that are likely to be out of my personal control.

The notion of **circles of trust** becomes important in this. We envision that it will be possible to easily define access to certain partitions of our data within these circles of trust. Examples for these can be myself only, my family, my friends, my visitors, my company and many others.

Next morning the link to the meteorite fall footage on his blog, which now directly allows subscription to anyone, has had 100 000 hits and, even though John doesn’t know this yet, the footage itself has thousands of cache entries around the world cumulating the total hit count for his content to more than a hundred million within 24 hours.

4.3 The Network as a Computing Cloud

John has had an enormous success with his holiday video (something he never expected). It is the story of many parties that he and Mary are visiting. At one of these events, he is asked (again) to show the rare event but also some more video snapshots surrounding the event. Being at a party, John did not carry a video-enabled device with him. But he of course did not leave the house without his PSIRPIM (PSIRP Identity Module), a small device that allows access to his scoped information from anywhere. Luckily, his friend who organized the party has a large flat panel available and the party guests insist on seeing the video in its full beauty. John inserts the PSIRPIM into his friend's panel slot, which results in the media environment

We assume an increased availability of computing resources everywhere, similar to case 1, whether fixed, mobile, end user operated or 'provider-operated', personally owned or borrowed on the spot. Abilities of the computing resources may vary from small devices (with little power and memory) to server-like machines with vast memory and an abundance of computing power.

being configured according to John's preferences (which are themselves stored in the network under his personal scope). Quickly John retrieves the video in question and demonstrates it to his friends. After he finishes (and retrieves his PSIRPIM), another friend approaches him asking to see some more of John's vacation highlights on his friend's video-enabled handheld. Using his PSIRPIM again, John retrieves some more video content. But he soon realizes that his friend's handheld is not HD-enabled. He therefore asks the system to include a transcoding function in the transmission, thus allowing him to show his vacation highlights. What he does not realize is that the system uses his friend's media system (the one he used a little while ago to show the meteor video to everyone) to perform the transcoding.

Grid computing has been gaining acceptance for a while now, i.e., purposefully assembling resources in a grid of available hosts. **Policies** play an important part and so does the publication and subscription to availability, information about the resource (metadata), and actual invocation of the service such as information retrieval. All this can be envisioned to happen dynamically, freeing the end user from specifying in advance use cases for particular devices (e.g., watching your movies on your iPod since you cannot transfer them to a nearby large screen).

4.4 Emergency Services: Scope is Important

John decides to go on a business trip to New York. New safety regulations require him to subscribe to emergency services, local and nationwide, to warn him about dangerous situations that might occur. His intention is to visit a client just outside New York, on Long Island, for some consultation for a new product. Unfortunately however, a minor fire situation on Long Island unexpectedly gets out of control and turns into a major wildfire, threatening major parts of local communities, also affecting Greater New York. The local authorities, totally surprised by the turn of events, are apparently overrun in their capacities and turn to the federal authorities for support and help.

Our vision of executing many services within a single network does not stop at emergency services. We assume control functions being in place that allow authorities (and users in general) to define the access and scope of reachability for their messages to be relayed to the public.

Quickly reacting, the federal government extends the scope of its information system so as to make it accessible to the local government for sharing important coordination information, thus allowing them to inform John on his mobile that he is about to move into a danger zone and that he should avoid given areas. While cancelling his business visit, he receives another message from his own government telling him what he can do to receive help from the local consulate (John had not yet updated his travel plans, reflecting the recent cancellation, so that this advise was still triggered). Since the escalating situation has already made it into the worldwide news, Mary receives a news item about the wildfire in John's journey area (luckily, her constant worries beforehand made her subscribe to New York related items before John had left for his trip). Immediately, Mary sends a message to John, asking him if everything is fine.

Services and their **reachability** can benefit from our vision of scoping information with respect to the set of receivers. Unauthorized senders cannot perform send operations, in the same way that unauthorized receivers cannot read information that is not intended for them. However, this use cases does not consider the problem of negatively affecting, even attacking, the actual scoping mechanisms (such as the rendezvous system).

In the meanwhile, a local group of protesters tries to take advantage of the escalating situation by attacking the local government's communication system with an intention to severely obstruct the rescue measures, with the possibility to later blame the authorities for severe failures. Fortunately, the lack of authorization to send messages under the name of the local government makes this attempt fail so that measures to control the wildfire can continue towards the restoration of law and order.

4.5 Emergency Services: Reaction is Important

The New York trip proves to be quite an adventure for John. After the shock with the wildfire disaster (luckily it was eventually brought under control), he manages to get out of New York, rent a car, and drive up to Maine to see an old friend of his. Unfortunately, the rush out of the city due to the earlier panic causes a major accident just up north on the I-95, involving a truck with chemicals. Some of its content spills on the highway, leading to a major hold-up and John is stuck right in the middle. He carefully observes the emergency services and their performance. Quickly, the scene of the accident starts to getting cleared up by local police. For this, it is necessary to understand what size tow-trucks might be needed, whether other heavy equipment is needed to either remove or restabilize vehicles, and so forth. These decisions are initially made by some combination of police and fire responders.

No particular background for this case apart from a choice of local access technologies and consumption devices.

Unfortunately, there are injuries this time. This requires the paramedics to gain access to information about the victims as well as regarding the side effects of the chemicals that were spilled. This may include medical records that, for legal reasons, would otherwise remain private. Information about the chemicals turns out to be proprietary, so clearance is also required for the medical workers to access it for emergency purposes, despite intellectual property ownership issues.

Emergency situations are a nightmare for everybody. They often require ad-hoc decisions to be made yet following a framework of actions, defined by often contradicting policies under often competing agencies. However, in actual cases of emergencies, the communication infrastructure often proves to be too inflexible to cope with these conflicting objectives. It therefore often breaks rather than adapt to the particular situation, in times when it is most needed.

Hence, the **mediation** of these conflicting policies (conflicts often unknown beforehand) is crucial.

While John is starting to move out of the danger area through the excellent work of the local police clearing the highway, other drivers are still approaching the scene of the accident, worsening the traffic situation. The local police decide to set up an emergency notification channel, using the local government communication system. In addition to merely contacting the approaching drivers, the system provides them with maps and directions for alternate routes. This information changes with time as traffic is routed off the highway progressively increasing the congestion in the initial alternate routes.

John is amazed by the quick reaction of the emergency services and the resolution of the different competence conflicts. He uses his Sparrow 2013 to access the necessary information, provided to him by the local authorities, and manages to continue his journey north.

4.6 The Future of SatNav

After a stopover in Boston, John changes his rental car to a new fancy sensor-equipped and -enabled model that the agency recommended. As a basic commodity, the car provides a routing system, called FutureNav, which is touted as being far superior to the competition in dynamic route planning. John leaves Boston with the destination Bangor in North Maine. The system arranges a scenic route to the destination, leading John over costal highways, avoiding the larger highways (John's New York experience did not put him in the mood to try the interstate again - the system, with John's permission, has knowledge of the past incidents and integrates it in the route planning). The FutureNav system is an ahead-planning system and keeps amazing John. A minor hold-up on a beautiful coastal road (a wedding ceremony blocked the road for some minutes for a photo session) led to a re-calculation of John's route, based on the information that was published by some of the cars that got stuck in the hold-up. The system also uses current and historic road-side sensor data in the route planning. When John decides that he would like to make a stop at a Plymouth mall, the system advises against it since historical data shows increased traffic in this area for the last two days. A click on the *Why* button reveals that a local megastore just opened its doors two days ago, leading to a flood of customers seeking out new special offers. When arriving in Bath (Maine), John decides to take some detour inland to enjoy some of the landscape. The FutureNav system, however, brings up a warning that sensor information in that area is not available for the package that John purchased at the car rental. However, he is offered a purchase for a one-day pass for accessing the necessary information (therefore extending the scope of his system). Since he is new in the area, he decides to foot the bill and go on with his trip.

We assume a plethora of sensor solutions deployed in the real world. These sensors serve environmental, traffic monitoring, health and individual monitoring purposes as well as many others. They are integrated in future cars, bodily worn in textiles and other accessories, and placed in our environment. We assume access to these solutions through appropriate connectivity means, such as gateways, or even their direct integration in an overall system. The number of these sensors is vast, several orders of magnitudes higher than today. Even more so, we assume that the ability to control access and consider the concerns of the providers and users of these solutions, along the lines of our tussle networking vision, will lead to a massive integration of sensor solutions for wider usage.

Sensor networks are very close in nature to the systems considered in PSIRP. They are largely information-centric, apply similar concepts such as 'network as a database' (see TinyDB in SmartDust), and will benefit from similar advancements in power, memory and bandwidth as our vision. Integrating specific sensor deployment seamlessly is often an issue of considering concerns of users and providers. These concerns range from business concerns (e.g., return of investment) over individuals' privacy concerns to societal concerns (e.g., exploiting sensor information to improve traffic situations in many countries). **Dynamically integrating** the different systems requires a **mediation** of potentially conflicting concerns, automatically if possible or falling back to some form of human-computer-interaction (HCI).

4.7 One Potential Future of Broadcast

After John left Boston, he realizes that he just missed the opportunity to attend a major American football game of the New England Patriots. He confers with his navigation system to find a hotel with proper broadband connectivity; luckily one is found quickly. Arriving at his hotel, he configures the local system with the help of his PSIRPIM to subscribe to his SuperSportHD service, which licensed the event in full scale, i.e., it allows him to individually subscribe to all the HD feeds of the 30 plus cameras in the stadium. In addition, the service offers aggregated views, the commentator's choice as well as the player's choice. He always loves to see the latter since they are based on one of the best player's choice of feeds, presenting a unique view on the game. After he has enjoyed the game (vividly switching views), he discovers yet another aggregated version of feeds from another broadcaster who installed HD cameras on the individual player helmets, allowing him to see the action from the players' perspectives. While recalling his SuperSportHD recording, John enables in another window the particular player

We assume a high proliferation of HD (high definition) imaging technology through proper recording facilities being available at major events, covering a large number of angles (including aerial). Bandwidth is not an issue due to advances in backbones and access, e.g., through fibre. Major events are likely to be marketed through single broadcast providers with the ability to license the raw material to a variety of mass and niche providers (i.e., the virtual operator model applied to content).

Streams of events, produced through HD imaging machinery, are just another input of data into a PSIRP network. The virtual operator type of model outlines two major issues, namely **aggregation** and **scope**, taking into account the digital rights of broadcasters and event organizers. Aggregation comes into play by aggregating raw feeds into specialized feeds, potentially down to individual receiver level (i.e., an individual person chooses a particular setting or angle of view). Scope comes into play when offering the aggregated streams to a particular audience. This scope can be country-wide (due to different regulations), localized, social (friends and family) or others.

helmet view, bringing a new experience to the game that he's just seen.

4.8 Remote Perception

John has been on the road now for quite some time. He misses his family quite badly. During the last conversation, Mary told him that she and the kids will go to a theme park, and he is wondering what their experiences will be. He notices that they are supposed to be in the park right now. Being homesick, he calls up Mary and asks her to share their experiences with him. Luckily, John subscribed to the new Travel&Share service before he left. This service aggregates relevant image and context feeds of his family and any willing participants into their experience. Luckily, the theme park contributes to this service as well, feeding images and location maps into the service. Being a busy place, many surrounding people have subscribed to this service as well, making the experience a very rich one. John has the ability to watch his kid's (and Mary's) various rides on roller coasters, enriched with camera sweeps from the park. In one of the new chase games, he switches to views from some of the surrounding competitors in the game and watches how his guys are doing (Mary gets a big hit while John is watching!). One of the players suddenly drops out of the game and heads towards the toilets. The service automatically turns off the feed and provides John with an alternative one.

He's really enjoyed this new multi person aggregated video and location service. It offers a rich connection with his family and makes him feel somewhat better.

Proliferation of converged devices will continue. Integration will drive the adoption of high technology into more and more devices, in particular mobile ones. In addition, more and more separate devices, specialized for one function, yet networked to contribute to a larger networked machinery, will be deployed, such as high end panels, large processing facilities etc. Hence, we will see high end equipment, being densely available in certain localities.

Everybody can be a producer as well as consumer. There is no particular pre-defined role anymore, everything is specified at runtime, as our tussle networking vision describes. Feeds can be created, shared, and destroyed in the blink of an eye. Concerns that might conflict with the particular usage at hand, such as privacy concerns when contributing location or image information to sharing scenarios, are automatically taken care off and only obvious conflicts will be raised to the end user involved. The concept of scope becomes very important here since this use case reflects the issue of **geographical scope** as well as **temporal scope**, both of which are highly tied into the concept of **context**. Context, such as location, surrounding, even affective state might contribute to the inclusion of a particular feed into a particular scope. This inclusion in itself may highly depend on the timing of the contributor as well as the requester. This context dependence introduces a dynamic into the rendezvous system that will require consideration.

4.9 ...and files all over again

John is finally due to return from his US trip. After arriving at Boston Logan airport, he realizes that his personal photo storage service is down (although it provided him with such reliable service during his Finland trip). In the hurry of his departure, he does not have the time for opening an alternative subscription with another provider. Being afraid of loosing his valuable memories of the past days (still being stored on his mobile device), he decides to go back to basics by copying the files to his distributed filesystem. From his mobile device, he opens the local file manager (configured to be connected to his distributed filesystem) and moves the media files from his local storage onto his distributed filesystem by simply dragging the directory into another branch (off his mobile device). Satisfied that his files are now safer, he finally boards the plane and heads home.

Files are seen to be a basis for storing material on computing resources, such as text and media files. Seen as information bits, the PSIRP system can be seen as a gigantic filesystem that enables access to files across devices under a particular scope of access, similar to computer or filesystem boundaries. This will continue to be important, be it merely for providing raw access to these files in certain cases (e.g., when a proper service for consumption does not exist).

The availability of files can be seen as a publication to a PSIRP subsystem that implements a distributed filesystem, i.e., within a certain scope - the scope representing the particular filesystem. Access to a particular file is a subscription to the actual data. In this manner, distributed filesystems (with access rights) will become a standard form of sharing files among devices, obviating the need to manually replicate material to multiple devices.

4.10 Elderly Care in a Pub/Sub World

After returning home, John visits his father, Paul, in his new elderly living facility. This place has been constructed just recently with all amenities of modern assisted living technology. As an old techy, Paul is very excited about his new environment although he does admit that he is somewhat overwhelmed with all the possibilities that the new environment offers. For that reason, John had offered to be of assistance with setup and also monitoring the environment but also Paul's health and well-being in general before he left for his business trip to the US.

As a first step, John configured Paul's system with a health monitoring service that shares the main physiological parameters with the local health provider, enabling his GP (general practitioner) to check on Paul's general health but also to be alerted of critical changes. In addition to sharing the data with his doctor, Paul has agreed that John has the ability to check certain parameters such as heart rate and others remotely, giving both Paul and John the peace of mind that things are going well.

John further configured the system (in collaboration with his physician) to remind his father of medication to be taken, re-stock the medication and re-order grocery goods from the supermarket via the Internet. In addition, notifications of slip-and-fall incidents are automatically sent to the local facility management team but also to John directly.

When John meets his father, Paul tells him about the new diet that the doctor prescribed. John already knew about this change in his father's dietary requirements since he had received a notice from his health monitoring system while he was away. John recommends to his father to share the new requirements with the grocery ordering service to allow for tailoring the orders towards this new diet. Paul agrees and John makes the appropriate changes to the ordering system so that any changes in the dietary requirements will be reflected in any future order.

Ambient and assisted living environments have been envisioned for many years. We see increasingly the integration of monitoring devices, such as for movement, slip-and-fall but also for health scenarios, into modern elderly care facilities, aiming for a more self-sustained life by elderly people. RFID and similar technology will bring the ability of monitoring drug usage, food consumption (and restoration) together with ubiquitous indoor communication. We assume such modern deployment for an elderly living facility.

The essence of this use case is not the typical sensor-type deployment within an assisted living scenario. This is certainly within the scope of the PSIRP vision of a generic information provisioning plane (and in particular one that operates under a publish-subscribe paradigm).

The focus here however is on sharing sensitive data but also managing this sharing through trusted parties, here the son of the data owner. Hence, scoping the information is (again) crucial here but also the authorization of the scope changes. Acting on behalf of somebody requires techniques in authorization and authentication for **delegation** of these functions to other trusted parties.

4.11 Learning with PSIRP

After John's visit at his father's place, he decides to pick up his own fatherly duties. He checks with his son Darren on the recent progress in his school's homework. Darren is doing really well in school. John receives frequent status reports of Darren's work, using the new online system that the institution provides (he is also notified by all grade A exams his son passes – creating a reminder for him to take Darren to the next football game on this occasion).

The current exercise however is a difficult one for Darren. It involves a mixture of physical geotagging (the tagging of images with contextual information like location) and history research for his political science class. The history research part actually focuses on past events in his local

An increased proliferation of access and (mobile) devices in the education sector can be seen already today. Initiatives like OneLaptopPerChild aim at deploying large numbers of information production and retrieval devices to an increasing number of children throughout the world. Together with many initiatives to connect schools to broadband access, we can envision an almost ubiquitous access of young children to the information network of the future.

town that were never really reported, i.e., they are based on individuals' memories rather than news or reports. John recommends talking to his grandpa about it.

Excited about this, Darren visits Paul in his home (he is quite excited about his granddad's new cool home) and starts interviewing him about the event. Paul recalls quite much of the past events and Darren records the entire interview. In addition, Paul hooks Darren up with some of the other residents who provide very useful additions to Paul's recollections. Happy with his success, Darren returns home. His voice recognition tool at home translates the recordings into meaningful abstracts and presents them to Darren. He arranges them into meaningful pieces of event recollections under the constraints given by the teacher's exercise. Together with his geotagged pictures related to the event, he securely submits the exercise to the school's systems.

Information is constantly produced, to be used in learning and production or for sheer pleasure. Information can be produced by anybody, being relevant for anybody. Learning as a way to consume information can be seen to be complemented by producing information that can enrich others willing to learn. **Creation** and **publication** of information, preserving the **identity of the author** (for credibility), together with consumption is seen as key for future learning. The aspects of scoping, trusted communication, and information centrism are crucial for pushing forward new methods of learning.

Upon reviewing the exercise, Darren's teacher is getting quite excited about the boy's inventiveness in collecting and digitizing useful past memories. He asks Darren and his father (as his guardian) for permission to publish the essay in the local library system for preservation and further research by others. The local library system watermarks the essay properly, identifying Darren as the author of the information, and re-publishes it with a research scope.

Darren receives many responses from local citizens for his contribution. Finally, he is awarded with a young journalist award from the local council, recognizing his young talent.

5 Acknowledgements

The vision and use case work, summarized in this document, has been a team exercise. Many individuals have contributed to the use cases in particular. These individuals are Janne T. (case 1, 2 & 3), George X. (case 3), Mikko S. (case 4), Andras Z. (case 6), Pekka N. (case 7), and Teemu (case 9). Not specifically mentioned are all the valuable inputs provided directly to me. Case 5 is provided with courtesy of Karen Sollins (MIT), originally written for the whitepaper in [7], adapted to our own storyline. Cases10 and 11 have been inspired by Arto K.

6 References

[1] PSIRP (publish-subscribe internetworking routing paradigm) project website, available at http://www.psirp.org

[2] 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

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

[4] D. Trossen, K. Sollins, "A System Design Vision", Presentation to the Communications Futures Program at MIT, November 2007

[5] Mark Weiser, "Homepage of Mark Weiser", available at http://sandbox.xerox.com/weiser/

[6] Mark Weiser, "The Computer for the Twenty-First Century," *Scientific American,* pp. 94-10, September 1991

[7] K. Sollins, D. Trossen (Eds), "Identity in Information Networking", Whitepaper of the Privacy and Security WG at the Communications Futures Program (MIT), March 2008

[8] V. Jacobson, "If a Clean Slate is the solution what was the problem?", Stanford 'Clean Slate' seminar, February 2006, available at http://cleanslate.stanford.edu/seminars/jacobson.pdf