PSIRP
Publish-Subscribe Internet Routing Paradigm
FP7-INFSO-IST-216173

DELIVERABLE D4.3

First report on deployment incentives and business models

Document Properties:

Title of Contract	Publish-Subscribe Internet Routing Paradigm
Acronym	PSIRP
Contract Number	FP7-INFSO-IST 216173
Start date of the project	1.1.2008
Duration	30 months, until 30.6.2010
Document Title	First report on deployment incentives and business models
Date of preparation	21.5.2009
Author(s)	Dirk Trossen (BT), Mikko Särelä (LMF), Paul Botham (BT), Trevor Burbridge (BT)
Responsible of the deliverable	Dirk Trossen (BT)
	Phone: +44 7918 711 695
	Email: dirk.trossen@bt.com
Reviewed by	Janne Riihijärvi (RWTH), Hannu Flinck (NSNF), George Xylomenos (AUEB), George Polyzos (AUEB), Trevor Burbridge
Target Dissemination Level	Public
Status of the Document	Completed
Version	1.0
Document location	http://www.psirp.org/publications/
Project web site	http://www.psirp.org/

SEVENTH FRAMEWORK PROGRAMME

European Commission
Information Society and Media

Table of Contents

1 Introduction 3
2 Deployment Incentives 4
2.1 Introduction 4
2.2 Rendezvous Function 5
2.2.1 Deployment Incentives 5
2.2.2 Interconnection Incentives 6
2.2.3 Interconnection Trust 7
2.3 Inter-Domain Topology Formation Function 8
2.3.1 Conceptual Architecture 8
2.3.2 Driving Needs 8
2.3.3 Scenarios 11
2.3.4 Incentives 13
2.3.5 Conclusions & Future Work 14
3 Business Models 15
3.1 Used Methodology 15
3.2 Results for Rendezvous Function 16
3.3 Results for Inter-Domain Topology Formation Function 18
3.4 Next Steps 19
4 Conclusions 20
5 References 21

1 Introduction

This document reports the first results from the considerations on deployment incentives and business models of propositions provided by the developing PSIRP architecture. Hence, this deliverable must be seen as a first extension of the socio-economic work that was originally presented in D4.2 [Psi2009b], based on the current PSIRP architecture as described in D2.3 [Psi2009a].

This deliverable focuses on the deployment incentives of two major components of our architecture, namely the **rendezvous** and the **inter-domain topology formation** functions. Both are inter-domain functions, requiring a deep understanding of the incentives for deploying the proposed technologies in the future marketplace due to the required collaborations of a multiplicity of actors. For this, we base the presentation of our first considerations in this document on the socio-economic work we perform in the larger (market) evaluation of the PSIRP architecture. In addition, we derive potential business propositions, even entire business models, for some of the existing and some of the expected new players in this new future market.

However, we appreciate that this understanding is something that needs deep investigation. Hence, the results of the present deliverable can only be seen as very preliminary ones in this direction, i.e., one must not expect a set of fully fledged business models within this report. Two main reasons for this preliminary nature can be given. Firstly, the considerations are based on an evolving architecture, the latest version of which has just been released as D2.3 [Psi2009a] a few months before this deliverable. Secondly, the wider socio-economic considerations necessary to create the understanding on deployment and business model level are ongoing activities themselves, as being reported in D4.2 [Psi2009b]. These considerations are now more deeply tuned towards understanding the future markets that are created by the PSIRP architecture, finding and explaining incentives of players to participate in these markets, and outlining business models within these new markets. We expect these investigations to continue and accelerate throughout the second year of the project with a final presentation of the overall results in D4.5.

The rest of the document is structured as follows. Section 2 presents our current considerations on the potential deployment incentives for the rendezvous and topology formation functions. Then, in Section 3, we present our first considerations on the level of business models, outlining first results of our socio-economic study. Section 4 concludes this report with some outlook to future work.

2 Deployment Incentives

This section outlines our current considerations for deployment incentives of two major components in the PSIRP architecture, namely the rendezvous and inter-domain topology formation function.

2.1 Introduction

The PSIRP socio-economic work will evaluate different architectural deployment options with a view to creating sustainable value chains, emphasising opportunities for enabling new businesses by architectural and technological design choices. PSIRP broadly seeks to replace currently deployed IP technology with something new and significantly different, potentially causing marked socio-economic changes such as:

- Creation of new markets (leaving out the necessity for infrastructural change to make this happen)
- Emergence of new players
- New ways of creating and maintaining communities
- Development of radically different business models

The PSIRP network evolution will be driven by many factors including:

- Technology advances (memory performance, new access techniques)
- Degree of collaboration between new PSIRP-compliant and legacy ISPs
- Regulation, which imposes (or removes) privacy and competition constraints
- User demand (based on perceived coolness, usefulness, reliability, security)
- Price pressure and investment strategies

Investigation of likely scenarios will give insight into the relative importance of these factors and corresponding strategies for migration. The models will help in:

- Defining and evaluating market scenarios
- Identifying new players to emerge due to the network design choices
- Devising relevant business models for potential players in these markets

The focus will then be on the likely impact of changes in markets, business models (operator/customer policies), regulation etc. These will all affect patterns of technology consumption (operator deployment and/or user take-up).

The eventual design of any new PSIRP function is crucial with respect to the socio-economic changes it is likely to cause. Vice versa, the likelihood of particular relevant socio-economic changes happening will influence the viability of certain design choices and, therefore, the deployment incentives for players in the market with respect to these choices. Hence, an understanding of such viability is crucial.

Challenges

Any new inter-domain technology, in order to be deployed, must be able to plug into the current technological and economical structure of the Internet. However, having merely an understanding of what can be easily deployed is likely to result in an ever-increasing number of quick fixes that each solve a particular problem but also lead to more complexity and make solving further problems even more difficult. For such reasons, clean-slate work is important, providing landmarks of what is technologically feasible (and desirable).

Migration is an important aspect in building the right incentives. For instance, the rendezvous system, described in D2.3, can be deployed either on top of the current Internet, or on top of the PSIRP topology management and forwarding functions. This makes the rendezvous system a promising candidate for migration from the current Internet to the publish/subscribe based internetworking architecture.

Perhaps the best-known techno-economical challenges for internet-scale routing systems are *efficiency* and *scalability*. Fir instance, efficiency can be measured as the average latency and AS-level stretch caused by the inter-domain rendezvous system. These must be low enough to enable the use of, e.g., interactive applications. Another important aspect is the efficient use of local resources - using already paid for resources whenever possible to maximize the operational efficiency.

Likewise, with scalability we mean that the communication, memory, and processing overheads must scale to the growing global Internet namespace sizes with reasonable cost, considering the likely technology development trends [ITRS2008]. Poor scalability implies higher costs, which will limit the types of objects the system will support, cutting off the long tail [And2006] of the system.

There is an implicit tension between efficiency and scalability, as the means for maximizing one influences the other. The balance between them is a matter of a run-time tussle [Cla2005] and needs run-time policy support to match the incentives within the evolving internetworking market (assuming that an approach of highly specialized architectures for limited scopes is not sustainable within an environment on increasing number of scopes – in other words, building an increasing number of highly vertical systems is not something desirable). As evidenced by the success of BGP [Rek1995], policy mechanisms are crucial in enabling local control of the system behaviour according to stakeholder interests [Sno2004].

2.2 Rendezvous Function

This section outlines the considerations on incentives for deploying the developed rendezvous solution, presented in D2.3 [Psi2009a], and for interconnecting various islands of rendezvous providers.

2.2.1 Deployment Incentives

Deployability is an interplay of several dimensions, including at least incentives and functional and economical feasibility of the system. It is important to understand which parts of a multi-participant architecture can be designed in advance and which will depend on the deployment and run-time incentives of the participants.

As it is not very likely that all parties in the global Internet will deploy the new system in the foreseeable future, a deployment strategy that requires everyone to adopt the system is doomed. It seems desirable to aim for the incremental deployment requirement: any party deploying the system by itself should get some utility out of the system. Additionally, any parties who have deployed systems should be able to interconnect their systems to gain the possible network effect benefits.

Whether or not deployment happens is determined by stakeholder incentives. If a party is required to invest in the system, that party should also be able to expect benefits at least in proportion to the (required) investment. Requiring investment without any obvious benefits is a recipe for technology failure. Even the chance of operational disincentive (such as possible loss of future revenue) is enough to make the system a hard sell. Dependence on third party deployment should be kept to a minimum in order to lessen the effect of operational disincentives. However, the other major objectives (scalability and efficiency) may require some level of functional specialization within the overall system.

From the above we can deduce that when applied to the current internetworking market, not all autonomous systems (ASes) will participate in or even support a rendezvous solution [Rat2005]. Therefore, while clustering deployment along the AS adjacencies are beneficial for efficiency reasons, it cannot be assumed to happen universally. Especially, as argued in [Raj2008], we should not expect large transit providers (the so called Tier-1 ASes) to support a system that may take traffic away from their networks, e.g., through caching information items on Tier-1 level. Some less obvious disincentives may apply at any AS, so the architecture should be prepared to bypass any AS, and only rely on support of the ASes that actually want to take part in the solution.

2.2.2 Interconnection Incentives

The faith-sharing principle [Cla1988] suggests that the local ISPs are natural candidates for rendezvous provision. They also have a built in incentive to provide the rendezvous service, as it will enable a high degree of localization of data-oriented traffic [Xie2008]. The desire for maximal utilization of already paid for resources may also lead to policies where information reachability is advertised over the existing inter-domain transit and peering links, whenever the provider or peer in question is willing to store such information reachability state. This may also expand over multiple AS hops.

Rendezvous service is a two-sided market [Roc2004] with the providers and users of the objects in their respective sides of the market. Two-sided markets are defined by cross-market network effects. In the case of rendezvous these happen between the publishers and subscribers of information. This, combined with the socio-economic complexity of the Internet, makes deployment a challenge. As an example of economic uncertainties, let us assume that the rendezvous service has network effects, i.e., the utility grows super-linearly, say O(N log *N*), where *N* is the size of the rendezvous network [Bri2006]. If there are two rendezvous network providers of sizes N_1 and N_2, should they co-operate by interconnecting their networks or not?

The utility that each operator gains from settlement-free interconnection compared to the status quo is on the order of $N_i \log ((N_i + N_j) N_i)$. But in a competitive market, one may not expect the status quo to be maintained. If either (or both) operators expect to gain market share, or even to capture the full market for itself, it may be beneficial to refuse interconnection or to offer it only on a monetary settlement basis. The beliefs that market participants hold about their future with or without interconnection, shape their responses and evolution of the system. Interconnection cannot be forced, and thus takes place only when both participants believe they gain from it compared to their beliefs of the future without interconnection.

To show how expectations shape outcomes, consider a simple interconnection decision between two rendezvous providers A and B. The providers may either continue separately or decide to interconnect if they can agree on appropriate compensation. If they continue separately, the sizes of their respective networks change due to competition. They form expectations of the future size of their own and the total network, as well as the effects of the possible interconnection on these factors. Each provider approximates its expected interconnection utility and compares it to a baseline of no interconnection.

Interconnection serves to allow a provider's subscribers to be reached from foreign publishers, but also serves to allow the provider's publishers to find out about foreign subscribers. Thus settlement increases the value to both networks for both publishers and subscribers. If we consider a larger and a smaller provider, the smaller provider's publishers gain more from interconnection as they see a proportionally higher increase in the number of potential subscribers. Similarly, the smaller provider's subscribers will see a greater benefit. However, when considering whether any settlement should occur, we need to also consider that although the smaller provider's individual customers benefit more, there are numerically less of them to benefit. Thus, the larger provider may still seek a settlement free solution. We should realise that any settlement solution can consider different settlement schemes for publications and subscriptions in both directions.

With more than two rendezvous providers and the AS structure of the Internet, the situation becomes more complicated. As an example, consider rendezvous being deployed by ASes in a transit relationship. The transit customer and provider can utilize their combined resources more efficiently if the provider also offers a rendezvous transit service, i.e. only the provider makes the full rendezvous state explicitly available and the customer relies on its provider for those items that do not lie within its own area. Similarly, peers in the AS structure may benefit by directly interconnecting their rendezvous state, instead of using their transit AS as rendezvous transit between them. This means that we can utilize parts of the AS structure in a manner of [Kop2007] as a model for forming rendezvous networks.

However, as the rendezvous related traffic will only represent a small fraction of the overall traffic, other issues, such as scaling costs, dependency on the transit provider, etc. may hinder arrangements such as those described above. Therefore, we do not expect the whole network to form one homogeneous rendezvous network, but assume a market-based evolution to determine the economical sizes of rendezvous networks, and the type of interconnectedness between them.

A rendezvous provider serves interested parties by matching interests, i.e., acting as a middle man and a point of (mutual) trust between publishers and subscribers. For example, a provider of an information object may register it to the rendezvous function through a subscription to receive requests for the object. Subsequently the provider of the information will be a publisher to deliver the content to the user. If there is only one rendezvous provider, it will need to be reachable by all potential customers, independent of their location in the inter-domain graph.

Such a situation draws attention to the situation that the rendezvous market is potentially a two-sided market with publishers and subscribers making the two sides with some participants acting on both sides of the market. The optimal pricing strategy (to maximize the size of the network) often means subsidizing one of the market sides. In the CDN market, we can see the network operator, the consumer of information, and the content provider may all share the costs. Some CDN operators offer free peering to ISPs, while others require settlement from ISPs to offset price reductions to the content providers. Users pay the ISP (through limited or unlimited subscriptions and additional volume charges), and may also pay the content provider directly. Content providers seeking lower costs can move to CDN providers that charge ISPs, but ultimately must be aware that these costs are passed onto their consumers (either as increased ISP fees or as reduced service quality for those ISPs who do not interconnect directly with the CDN).

In summary, the deployment and interconnection incentives presented here outline the relation to the inter-domain topology formation function presented next, in particular on the interconnection incentives.

2.2.3 Interconnection Trust

Along with economic incentives to interconnection strategies, we should also consider the trust relationships that exist. The user (publisher or subscriber) must have some trust in the network provider to which they attach. This attachment network will offer the initial forwarding and rendezvous capabilities and thus forms the initial link in a chain of trust between the user and more distant rendezvous providers. When considering whether to interconnect, therefore, a rendezvous provider must consider not only the direct economic benefits, but also the desires (and hence value) to their users. Interconnection can be seen to increase the risk of many attacks including phishing (rogue subscribers) and Denial-of-Service (rogue publishers). Thus whether interconnection takes place will depend upon the security and interconnection policy of the network we are connecting to.

2.3 Inter-Domain Topology Formation Function

This section focuses on the inter-domain topology formation function. The design of the developing inter-domain topology formation solution is driven by a set of identified *needs*, based on a set of *scenarios*. The following sections present these needs and scenarios. With a solution space that fulfils these needs, we can expect to create a **marketplace** with players having an incentive to invest in the deployment of the necessary technology. We round off our section on inter-domain topology formation by presenting potential incentives for such market participation. But first, we briefly repeat the underlying conceptual architecture for our considerations.

2.3.1 Conceptual Architecture

In this section, we briefly review the conceptual architecture relating to inter-domain topology formation, as detailed previously in [Psi2009a]. A *topology management* function is assumed to exist within each autonomous system (*domain*). This function implements the local topology management and communicates the relevant peering information to the *inter-domain topology formation* (ITF) function.

Publishers and *subscribers* come together in the rendezvous process within the *rendezvous point*, representing the particular SId in which the information items (labelled via an RId) are located. The arrows in Figure 2.1 show the relations of these components and are not meant to illustrate the exact message and information exchange between them. However, dashed arrows indicate relations stemming from the rendezvous process while solid arrows show topology formation relations.

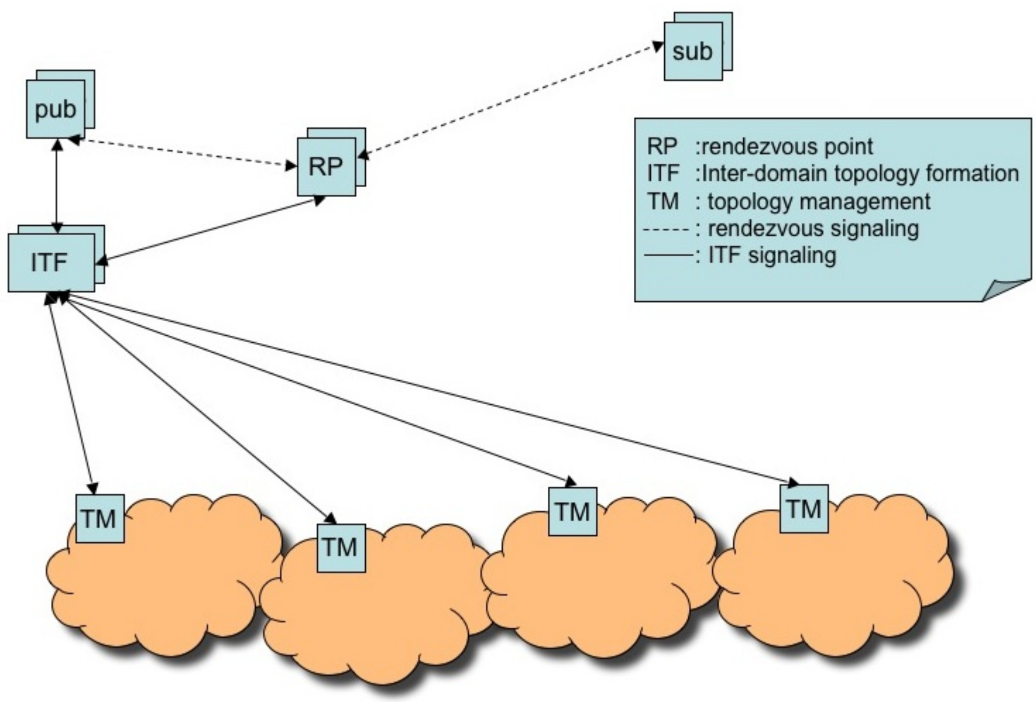

Figure 2.1: Conceptual Architecture for the Inter-Domain Topology Formation Function

2.3.2 Driving Needs

From a socio-economic perspective, we first identify the fundamental network functionality/attributes required by operators and/or users. These “driving needs” serve to motivate any subsequent PSIRP deployment in specific service scenarios, as developed below.

Resilience

One of the fundamental requirements in high capacity networks is "survivability", referring to the ability of a network to recover affected traffic in failure environments and to provide different services continuously. Conventional resilience helps by offering protection against network downtime, eliminating single points of failure through the use of duplicate network components (circuits and routers), in conjunction with special redundancy protocols.

High availability networks can be costly to deploy and maintain, so it is important for an organisation to identify the optimum level of resilience for its requirements. This is achieved by balancing the potential cost of network downtime against the cost of any proposed expenditure on resiliency measures. The degree to which an enterprise needs resilience built into its network is obviously directly related to the importance of network applications to its core business processes and will vary from industry to industry. However, costs incurred from network redundancy can provide benefits other than resilience. For example, multi-path routing can enable traffic to be balanced more evenly across the network and lessen the load at choke points.

PSIRP resilience is largely derived from the DHT-type architecture in the rendezvous solution [Psi2009a], with an ability to self-organise in response to failures. The extent to which PSIRP can support various levels of resilience will be crucial for its deployment prospects.

Policy Compliance

Any new Internet architecture needs to take business relationships between ISPs into consideration. The contractual relationships between domains (inter-domain policies) define the internet topology, not the underlying physical connectivity. Bilateral relationships (transit, peering) between ISPs effectively prune the physical network topology, restricting the path choices available for end-to-end traffic. This policy-constrained topology must then be covered by the routing architecture, the policy view of the network leading naturally to the valley-free routing model.

In PSIRP, the inter-domain topology formation process must resolve potential conflicts between policies of the sending and receiving hosts. For example, they might disagree about desirability of paths traversing a particular AS. But also, policies of ISP being traversed during transit of particular information must be taken into account. For instance, certain information being carried of certain ISP might be considered illegal by the jurisdiction covering the legal entity of the ISP.

Today, blocking unwanted traffic depends on configuring access control lists (ACLs) or null routes at various points inside the network. Ideally, unwanted packets should be discarded close to the sender to reduce bandwidth consumed. Traffic can be discarded based on a wide variety of policies, such as a combination of source and destination prefix, source and destination port numbers and protocol. Crucially, criteria might also include information content, permitting a much finer degree of discrimination.

The PSIRP "black box" model of the network generally enhances privacy by making internal structure of a network private and employing routing based on data (rather than destination) identifiers. Furthermore, an explicit rendezvous system matching the wishes of publishers and subscribers provides a single point in the architecture for controlling tussle over resources, resulting in a "trust-to-trust" (rather than traditional end-to-end) design.

Quality-of-Service

Network quality of service (QoS) is the ability to provide to different applications, users or data flows certain level of performance. QoS guarantees are important where network capacity is a limited resource, particularly for real-time streaming multimedia applications such as voice over IP, online games and IP-TV (since these often require fixed bit rate and are delay sensitive).

Where the expense of mechanisms to provide QoS is justified, network customers and providers typically enter into a contractual service level agreement (SLA) which quantifies the ability of a network/protocol to give guaranteed performance/throughput/latency bounds based on mutually agreed measures.

Today, ISPs offer only coarse-grained SLAs for traffic staying within a single AS. Providing fine-grained QoS over an end-to-end path is extremely difficult but PSIRP could address this through a combination of wider view and an information perspective:

- Today's effectively deployed Internet, as opposed to proposed Internet protocols that have not gained wide acceptance and deployment, does not provide end-to-end signalling for reserving resources along a path traversing multiple institutions, or offer commonly agreed in-path QoS signalling. This makes it difficult to offer performance guarantees to end users. However, the inter-domain topology formation process could negotiate strict QoS guarantees with individual ISPs then stitch virtual links together to provide end-to-end QoS to customers.
- An ISP can provide QoS for highly aggregated traffic, but offering performance guarantees for individual flows is extremely challenging, in terms of signalling overhead and need for fine-grained packet scheduling. In contrast, the ITF process could reserve bandwidth across several ISPs for aggregated traffic and manage the division of these resources across individual flows. Reserving bandwidth does not necessarily mean that others cannot use the bandwidth, only that the ITF has priority over other traffic in times of congestion.
- Today's ISPs determine which traffic should receive priority service based on bits in the packet headers but cannot easily classify packets based on finer-grained information, or direct packets on different paths based on their performance requirements. In the PSIRP architecture, packets could be classified based on diverse customer policies and be assign to a sequence of virtual links with the necessary performance properties for each flow.

In principle, the ITF process could decide what performance guarantees to offer and how. As discussed above under policy compliance, criteria might also be information-based, enabling much finer differentiation. This allows for differentiation by offering special QoS services to customers, opening the way to new market opportunities.

Trust

Trust plays an important role in the inter-domain topology formation process. Design choices for this process must take into account different assumptions for trust relations between the relevant parties, i.e., publishers, subscribers, rendezvous point (RP), and the ISPs' topology management functions. Such trust assumptions must not prescribe particular relations requiring, e.g., ISPs to solely rely on a third party for creation of the inter-domain topology.

We can start from the trusted relationship between the user and the forwarding network to which they attach. This network is also likely to offer rendezvous and internal Topology Management facilities. From this point, we need to consider how the ITF is reached, and what information is returned to the user. For example, if the ITF is reached through one or more rendezvous providers, then there is an assumption that the chain of trust through the rendezvous providers is sufficient to allow them to make the correct choice of ITF function. Alternatively, the ITF may be contacted through the attachment forwarding network and its Topology Management services.

Considering trust in the other direction, forwarding providers trust the ITF with topology and policy information that they publish towards the ITF or that will be provided to the ITF in the process of topology formation. We need to consider whether this information may be shared with the rendezvous providers, or even with end users (who may try to discover topology information to attack the network).

Summary

On the basis of the above discussion, the driving needs for PSIRP deployment may be classified as fine-grained (information-based) support for

- Resilience
- Policy compliance
- Quality-of-service
- Trust

In a closely related study, the authors of [Lak2004] propose solving the internet routing tussle between end users and ISPs by outsourcing route computation to third-party routing service providers (RSPs). These would buy virtual links (with well-defined SLAs) from current ISPs and be able to offer specialised routing services (such as avoiding certain ASes for policy reasons, blocking unwanted traffic and QoS routing).

Technically, many of these advantages derive from a more global view of network topology and greater ability to aggregate customer demands. The PSIRP rendezvous point (RP) shares these characteristics, with the additional advantage of much finer-grained differentiation based on packet information content, offering an even higher level of control. Ultimately, the case for PSIRP/RSP technology rests on the expected demand for these specialised services.

Having identified the “driving needs” motivating PSIRP deployment, we next consider several likely service scenarios for deeper socio-economic analysis, particularly with regard to deployment incentives.

2.3.3 Scenarios

It is important to understand potential scenarios in which inter-domain topology formation takes place. These potential scenarios drive our view of evaluating design options for this function but also the considerations of deployment incentives for its implementation. The following sections give examples of such scenarios. For each of these scenarios, we give a brief description of its essence and describe three types of information:

- *End user identity* information that is involved in the scenario
- Information related to the *physical domain*, describing the domains (including devices) that need to be traversed during the delivery of the end user experience. This information serves the purpose of determining what aspects need to be considered when attempting to build an inter-domain delivery graph.

The space of information required to ensure a given user experience is meant to illustrate the space of interaction among different parties when forming suitable delivery graphs for the scenario at hand. It also shows that potentially several formation functions need to (co-)exist in order to implement the given scenarios.

Social TV Service

Description

The Social TV service aims at providing an enhanced TV experience by combining traditional TV-like services with social networking techniques. With this, opinions and recommendations on programmes can be exchanged with friends. Widgets on computers allow for viewing which friends in your social network are currently watching what programme. TV experiences are enhanced through extra information on shows provided before, during and after the shows. Furthermore, targeted purchasing offers can be provided, e.g., through medication reminders when watching a health show. Also, optional video conferencing services can be provided with your friends to exchange impressions and talk directly.

In essence, participants will be able to orchestrate simultaneously watching a programme with friends from their social network. Thus, some intersection of relevant information will be the basis of that shared experience. Furthermore, if some of the viewer's social networking friends are “regular” TV provider customers and others higher-level “platinum” members, the shared experience might include only regular membership level information.

Information

User Information: Facebook ID, profile, email etc. Comcast IPTV customer ID, membership level (affecting information availability), serial number of viewing card, TV guide recommendations (e.g. time, content, trailers etc.), personal favourites/history

Physical Domains: Internet, Comcast IPTV network, mobile network, set-top box, mobile device/WLAN, mobile device/Bluetooth

Payment Service

Description

In a financial service scenario, a secure and reliable exchange of relevant financial information is desired, e.g., for purchase orders in the field (via mobile technology) or for store purchases. Dedicated devices, similar to PIN-based card readers nowadays, or standard mobile devices (enhanced by security software) can serve as a terminal to initiate the transfer. Information relevant for the purchase, such as product information, barcodes, etc., can be exchanged via other services similar to today's Web services.

Information

User Information: secure login ID, profile (type/class of service), email, bank financial information etc.

Physical Domains: VPN, Internet, mobile network, mobile device/WLAN, mobile device/Bluetooth

Information-specific Policy Enforcement

Description

Certain information within a given service, such as Web browsing or video delivery, needs to be treated in accordance with local jurisdiction of the country where the delivery takes place. Also, local ISPs might need to comply with local policy rules on a fine-grained level. For instance:

- Some countries impose restrictions on video content that can be accessed in their jurisdiction (e.g., in Thailand, disrespect of the King is illegal, and videos violating this law need to be blocked there).
- Offences to religion can be treated in certain countries as a criminal offence. Not complying with a ruling to remove single information items can lead to blocking the entire service in that jurisdiction.

Such material can currently only be blocked entirely worldwide, or at single ISP (e.g. You-Tube) level. The PSIRP architecture offers the prospect of control at any/all of pub/sub, rendezvous and topology management levels due to forming an inter-domain topology that is compliant with these policies.

There might also be commercial reasons for, e.g., not carrying Provider A transit traffic via a competitor's network. Hence, an ISP might want to have certain policies enforced in the decision to construct certain inter-domain topologies.

Information

User Information: ISP login ID, profile (type/class of service), email

Physical Domains: Internet, mobile network, mobile device/WLAN, mobile device/Bluetooth

Quality Content Service

Description

Future extensions to video delivery, e.g., ala YouTube, might consider video-on-demand like functionality in which video delivery for particular items, wanted (or even sponsored), is performed with a certain quality of service for the information item in question. Such on-demand and catch-up services will exist alongside live streaming content. There may exist differentiated pricing per information item, either through the information provider, the transit provider, or the rendezvous provider.

Information

User Information: ISP login ID, profile (type/class of service), email

Physical Domains: Internet, mobile network, mobile device/WLAN, mobile device/Bluetooth

Emergency Service

Description

Services similar to today's 112/911/999 emergency services will require ensured quality of service for all information items required for implementing the desired functionality. This is likely to go beyond pure voice transmission. In a scenario, which incorporates necessary patient information, sensor data as well as audio-visual information, all the various information items being retrieved in a distributed manner need the preferential treatment for their delivery. Potential regulatory requirements, e.g., privacy regulations for patient records, may introduce particular requirements for the underlying links being used in the transmission of this particular data, while other data can be retrieved via normal (current) Internet connectivity.

Information

User Information: Mobile phone number, location, (ideally, patient medical history)

Physical Domains: mobile network, mobile device/WLAN, mobile device/Bluetooth

2.3.4 Incentives

As stated in the introduction, we see the incentives for participation in the created marketplace being generated by the potential to fulfil clearly expressed needs of resilience, policy compliance, and quality of service. Our approach here is to look at the actors in the conceptual architecture and see how their participation in the market would fulfil these driving needs. For this, we describe for each actor the particular implementation of the needs. This implementation of needs gives a first introduction to the potential value creation through this actor and therefore builds the foundation for future business models.

However, it is important to understand that, quite apart from technical merit and feasibility, it is most important to establish the viability of the various scenarios in the face of likely commercial and regulatory pressures. Accordingly, each scenario has been assessed regarding the driving needs (hence incentives for participation) for prospective deploying parties of the technology, in terms of the broad classification introduced above. The result is shown in Table 2.1.

Service	**Need driving deployment incentive**
Social TV	Resilience, Policy Compliance, QoS
Payment	Resilience, Policy Compliance, QoS, Trust
Information-specific policy enforcement	Policy Compliance, Trust
Quality Content	QoS, Policy Compliance
Emergency	Resilience, Policy Compliance, QoS, Trust

Table 2.1: Driving Needs for Each Scenario

Given the driving needs and based on our conceptual architecture, we can outline the market player incentives for deploying technologies pertaining to the inter-domain topology formation process. These incentives are shown for the functions *topology management* (local ISP & transit ISP), *inter-domain topology formation* (ITF) function and *rendezvous* in Table 2.2.

Player	**Incentive for deployment**
transit ISP	peering/transit link differentiation, attraction to transit customers, regulatory policy compliance, potential optimization of resources through choice of ITF
local ISP	peering/transit link differentiation, attraction to local access customers through differentiation, regulatory policy compliance, potential optimization of resources through choice of ITF
ITF	brokering of inter-domain resources, trusted 3rd party for topology creation
Rendezvous	attraction to publisher/subscriber through differentiation, regulatory policy compliance, offering QoS differentiation

Table 2.2: Incentives for Market Players

2.3.5 Conclusions & Future Work

This section outlined the incentives for players in potential future markets for inter-domain topology formation within the PSIRP architecture. These incentives were based on a high-level conceptual architecture for this important process within the PSIRP architecture, as first presented in D2.3 [Psi2009a]. Although particular design choices for this process have not yet been designed in full, our preliminary analysis outlines incentives for players in these markets that show positive incentives for the deployment of the final PSIRP architecture in general and the topology formation technologies in particular.

Our first results on incentives, however, will need further exploration with a deeper socio-economic analysis as to the driving factors for adoption, the identification of potential obstacles and showstoppers and the development of exemplary business models for selected players. This will be the focus of our future work, building on our developing socio-economic models that were presented in D4.2 [Psi2009b].

3 Business Models

Another important aspect of our socio-economic considerations within the project is to develop potential business models for existing and novel players in the value chain(s) to be created by the PSIRP architecture. The following sections give an overview of the current work in this area. It is important to note however that this work is still in its infancy. More detailed results are expected for the second version of this deliverable, D4.6, in M28.

3.1 Used Methodology

D4.2 [Psi2009b] introduced the methodology underlying the socio-economic work within PSIRP. This methodology is based on original work at the Communications Futures Program at MIT [CFP2009] and was significantly improved for the first results in D4.2. The following Figure 3.1 provides an overview of the used methodology.

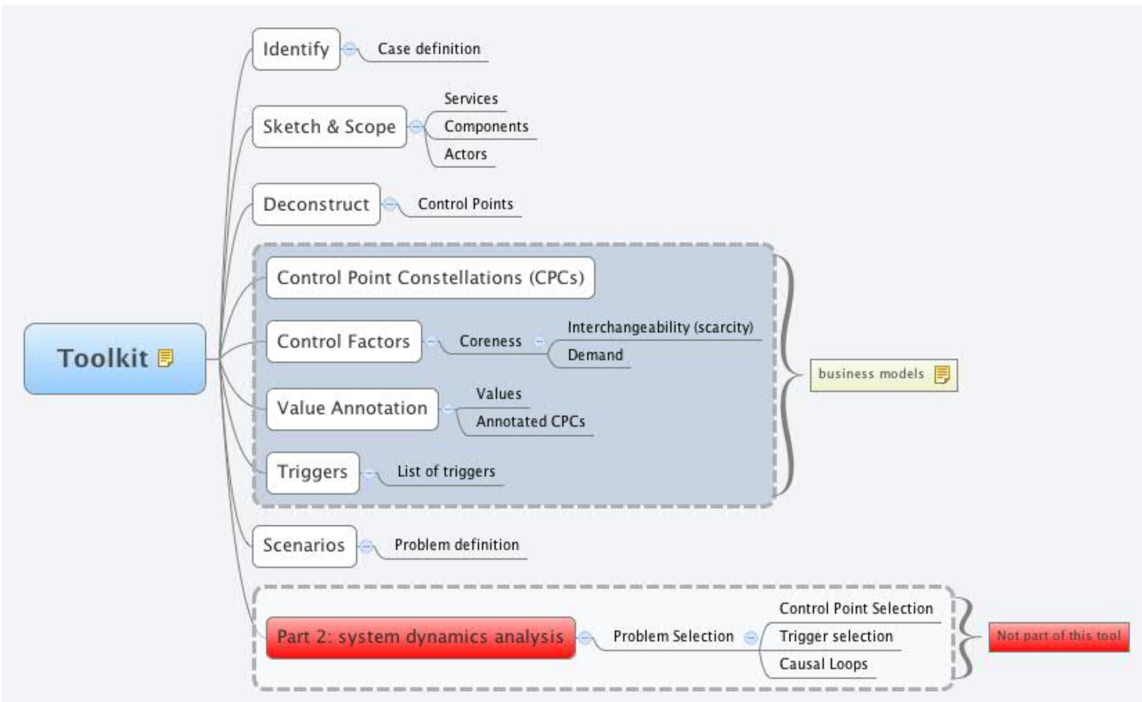

Figure 3.1: Methodology: Value chain analysis toolkit

D4.2 provides more insight into the different steps and their significance. For our considerations here at the business model level, the steps of identifying the services, actors and components (sketch&scope step) as well as the identification of control points (deconstruct step) are the most relevant ones. In addition, the triggers step provides information that relates to the sustainability analysis of identified business models.

In the following, we outline these three steps for two major architectural functions, namely the **rendezvous** and the **inter-domain topology formation** function.

3.2 Results for Rendezvous Function

As described in D4.2, we chose a relatively simple use case of publishing an information item within a Facebook-like social networking setting. With that in mind, we identified the following set of actors, components and services, shown in Figure 3.2.

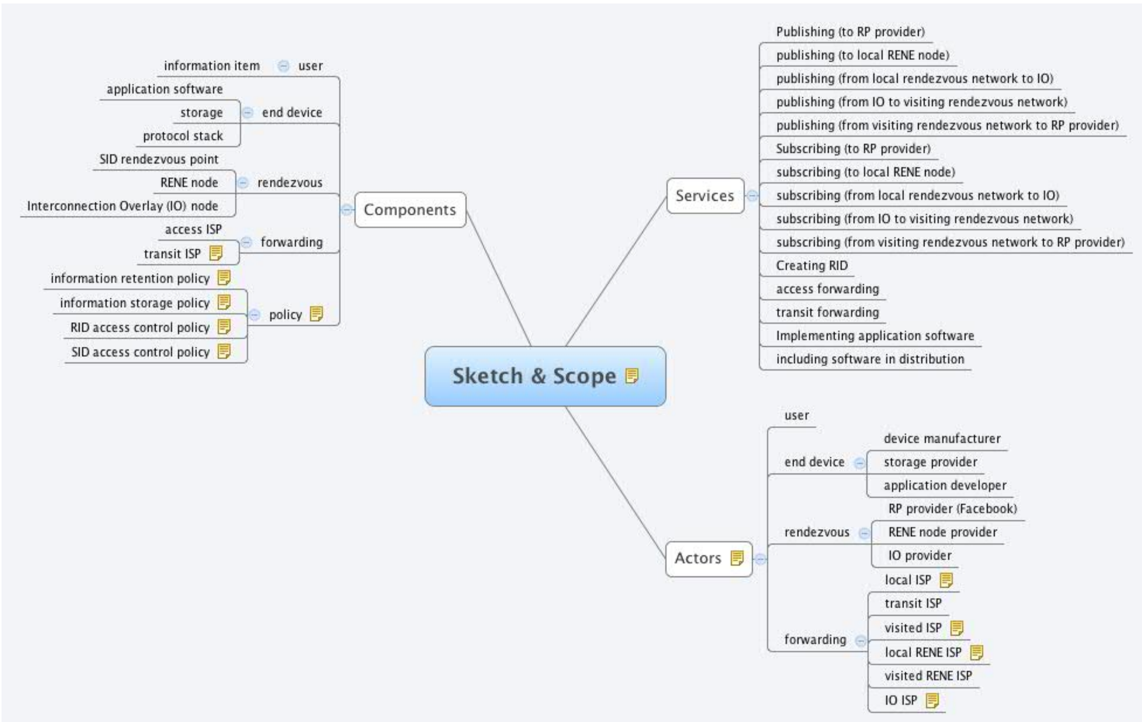

Figure 3.2: Sketch & Scope step

Our focus will be on the actors in the rendezvous space, i.e., *interconnection overlay providers* as well as *rendezvous network (RENE) node providers*. The ISP forwarding space will be addressed in a separate consideration on inter-domain topology formation.

The deconstruct step of the methodology gives some insight into the potential control points that are important for the construction of valid business models. Figure 3.3 shows the range of potential control points for the rendezvous use case.

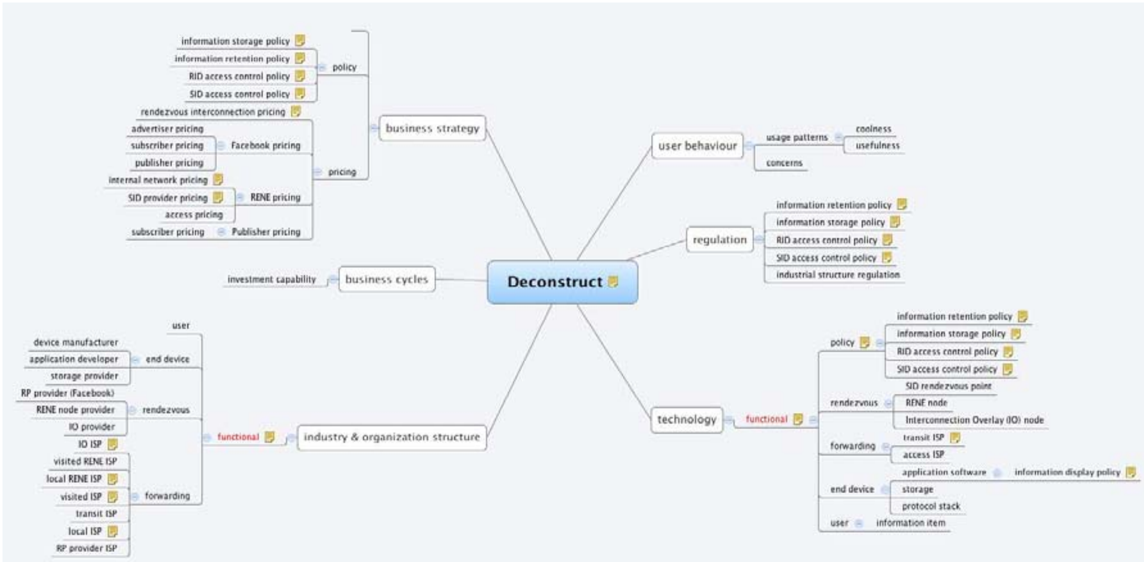

Figure 3.3: Deconstruct step

Apart from the obvious functional control points, a specific importance is seen in the regulatory and user behaviour control points due to the role that rendezvous plays when matching information item interests under given policies. This is also reflected in the trigger analysis, listing the influences of various kinds on the control points. Figure 3.4 shows the identified triggers for the rendezvous use case.

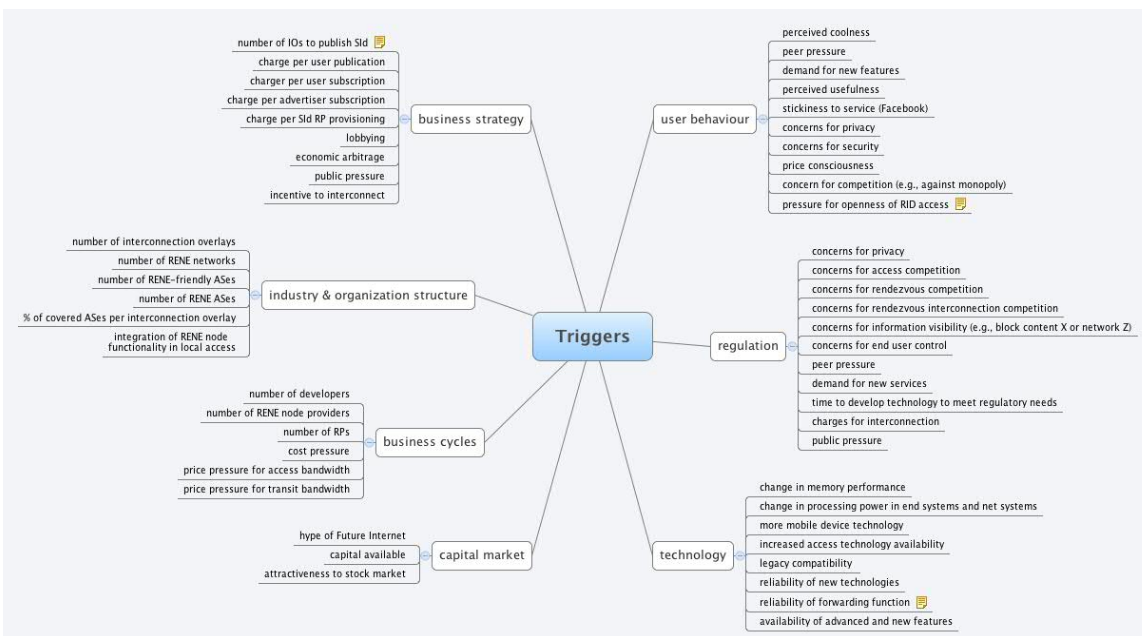

Figure 3.4: Triggers step

Knowledge from other areas of disruptive services, such as Voice-over-IP or digital music, tells us that triggers in the space of user behaviour (e.g., perceived coolness and concerns of various kinds), as well as industry and organization structure usually have a significant impact on the viability of developed business models and their uptake in the market. Hence, our focus is likely to be on the influence of these triggers.

The methodology presented in Figure 3.1 outlines two important steps for the consideration of business models, namely the **value annotation** as well as **control factors** steps. These serve the purpose of assigning a certain value to functional as well as non-functional control points; value which can be extracted by a particular market player in forms of business models. The control factors step outlines the trends as to which control in certain control points can be sustained and executed over time (e.g., technology development can lead to less scarcity of control points, therefore making its centralized control more difficult). At this stage, we have not yet executed these two steps.

3.3 Results for Inter-Domain Topology Formation Function

The work performed in the inter-domain space has a different focus compared to the rendezvous work. Here, we intend to develop business models for the main players around these particular architectural functions, i.e., business models are likely to revolve around peering markets, exposure of peering links to various ITF providers, and offerings for policy compliance, resilience and quality of service (see Section 2.2).

With the use case centred on the retrieval of an information item, i.e., after matching at the rendezvous point has been performed, the following set of actors, components, and services can be identified (based on the conceptual architecture for this function, as shown in Figure 2.1).

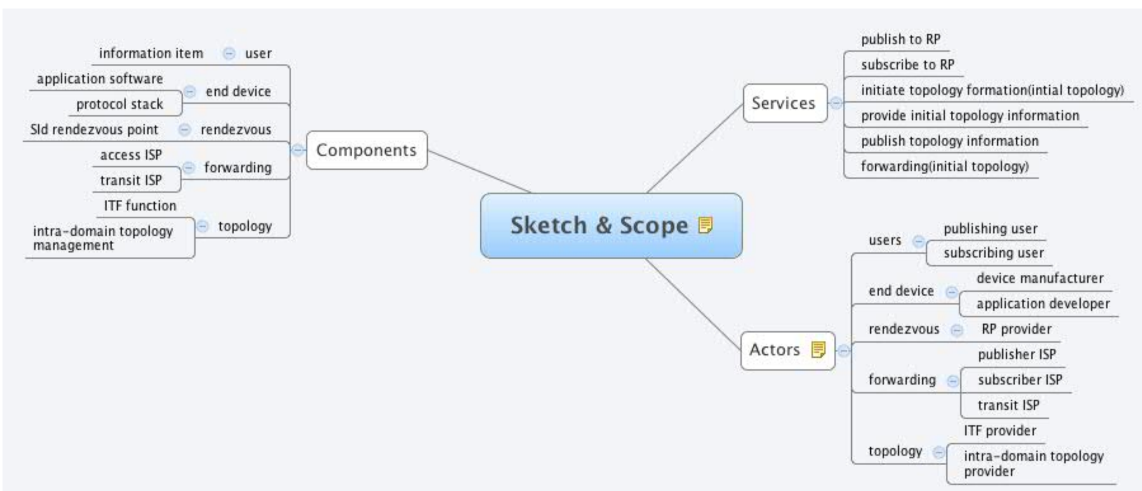

Figure 3.5: Sketch & Scope step

While the components in Figure 3.5 are largely similar to that of the rendezvous case in Figure 3.2, differences come in when considering the actors and services on the right hand side. The focus of our business model considerations will be on the different ISPs, the rendezvous provider as well as the intra-domain topology provider (implementing the topology management for a particular AS).

Similar to our rendezvous case, the user behaviour as well as regulatory control points, shown in Figure 3.6, can be seen as crucially important in the consideration of certain business models. For instance, control points in the regulatory space around access and transit regulation will have an impact on the expected market structure and therefore the viability of certain business models.

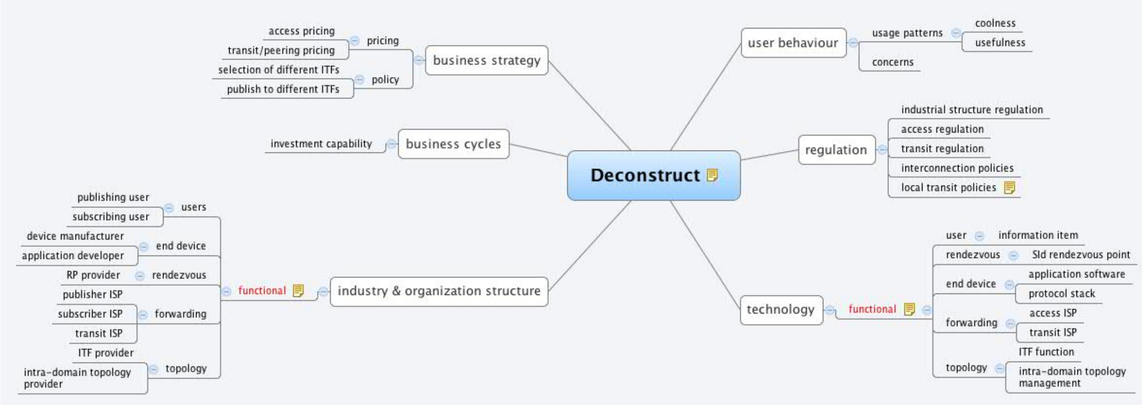

Figure 3.6: Deconstruct step

The step for the trigger analysis has not yet been performed, together with the value annotation and control factor steps. Hence, we are still at a very early stage of the analysis. This is also due to the ongoing design for the actual technical solution for inter-domain topology formation.

3.4 Next Steps

While the current results of our socio-economic work have not yet resulted in concrete business models, we have laid the ground work for doing so in the near future. The major control points and their points of influence, i.e., the triggers in our methodology, have been identified for two major components in our architecture.

Based on this, it will be possible to develop analytical system dynamics models for the sustainability of particular propositions, i.e., business models, based on certain assumptions of values that can be assigned to certain control points. The development of these models will be our next steps in this space.

However, in order to be able to assign proper value to particular control points, we need to better understand the market potential and size. This requires proper market models at a macro-economic level; something we are currently working on in the overall socio-economic evaluation (with the first results presented in D4.2 [Psi2009b]). In addition, we will require input into the steps of value annotation and control factors (see Figure 3.1); steps that reflect evolutionary trends in value during a transition from the current to the future (PSIRP) Internet. We intend to address the former issue by continuing work on macro-economic models, as outlined in D4.2, allowing us to make statements on markets and market developments in terms of players as well as market concentration and, potentially, size.

The second issue of providing proper input to the missing steps within our methodology is addressed through dialogue with business developers and strategists from current market players in the space of ISPs, manufacturers, content providers and others. This dialogue will take place in the form of a workshop in which scenarios and the parameters for their models will be developed in a collaborative way between PSIRP project members and an industrial audience that will bring the required experience to this exercise. We are confident that this will bring us closer to realistic business and market models for our architecture.

4 Conclusions

This deliverable sheds a first light on our evaluation of deployment incentives and the development of business models, based on our developing PSIRP architecture. However, the results presented here can only be seen as a first step towards this understanding. This is due to the ongoing work on the underlying socio-economic models that will allow for the development of the right incentives and business models. Furthermore, design work on crucial components such as the inter-domain topology formation is still ongoing, so that considerations on business models can only be preliminary. Most importantly, our work can only be the preparation for a dialogue with external experts that will bring the expertise to our models that will allow for a proper formulation of realistic business models.

Nevertheless, this deliverable has presented crucial results that lay the ground for further developing our work. A clear methodology has been developed that is used in overall socio-economic evaluation of our architecture and also for our incentives and business model work. This will give us a clear framework to continue working on. Furthermore, we presented first results within this methodology for two crucial architectural components, namely the rendezvous and inter-domain topology formation function, both of which we expect to be of the most interest to the business audience we intend to approach with our work.

The next steps of our work will further develop the models, based on our methodology, towards a point where we can seek the dialogue with business developers and strategists to collaboratively develop concrete business models that will outline the potential for the PSIRP architecture from particular players' point of view. This work will be documented in the upcoming deliverable D4.6 in M28.

5 References

[And2006] C. Anderson. "The Long Tail". Hyperion, 2006.

[Bri2006] B. Briscoe, A. Odlyzko, and B. Tilly. "Metcalfe's Law is Wrong". Spectrum, IEEE, 43(7):34{39, 2006.

[CFP2009] Communications Futures Program (CFP), http://cfp.mit.edu [accessed on 19th May, 2009].

[Cla1988] D. Clark. "The design philosophy of the DARPA internet protocols". ACM SIGCOMM CCR, 18(4):106-114, 1988.

[Cla2005] D. Clark, J. Wroclawski, K. Sollins, and R. Braden. "Tussle in Cyberspace: Defining Tomorrow's Internet". Networking, IEEE/ACM Transactions on, 13(3):462-475, June 2005.

[ITRS2008] ITRS. "International Technology Roadmap for Semiconductors". www.itrs.net, 2008.

[Kop2007] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica. "A Data-Oriented (and Beyond) Network Architecture". In ACM SIGCOMM'07. Proceedings, pages 181-192, 2007.

[Lak2004] K. Lakshminarayanan, I. Stoica, S. Shenker, "Routing as a service". Technical Report UCB-CS-04-1327, UC Berkeley, 2004.

[Psi2009a] M. Ain (ed.), "PSIRP Deliverable 2.3: Architecture Definition, Component Descriptions, and Requirements", 2009.

[Psi2009b] J. Rihiijarvi (ed.), "PSIRP Deliverable 4.2: First report on quantitative and qualitative architecture validation", 2009.

[Raj2008] J. Rajahalme, M. Sarela, P. Nikander, and S. Tarkoma. "Incentive-Compatible Caching and Peering in Data Oriented Networks". In ReArch'08. Proceedings, 2008.

[Rat2005] S. Ratnasamy, S. Shenker, and S. McCanne. "Towards an Evolvable Internet Architecture". In ACM SIGCOMM'05. Proceedings, pages 313-324, 2005.

[Rek1995] Y. Rekhter and T. Li. "A Border Gateway Protocol (BGP-4)". RFC 1771, IETF, Mar. 1995.

[Roc2004] J. Rochet and J. Tirole. "Defining Two-Sided Markets". Toulouse, France: IDEI, mimeo, January, 2004.

[Sno2004] A. Snoeren and B. Raghavan. "Decoupling policy from mechanism in Internet routing". ACM SIGCOMM CCR, 34(1):81-86, 2004.

[Xie2008] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz."P4P: Provider Portal for Applications". In ACM SIGCOMM'08. Proceedings, pages 351-362, August 2008.