PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

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

DELIVERABLE D5.7

Migration Plan

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	33 months, until 30.9.2010
Document Title:	Migration Plan
Date of preparation	22.6.2010
Authors	Hannu Flinck (NSNF) (ed.), Jane Tateson (BT), Andras Zahemszky (LMF), Perti Jokela (LMF), Janne Tuononen (NSNF), Dmitrij Lagutin (Aalto-HIIT)
Responsible of the deliverable	Hannu Flinck (NSNF) Phone: +358.50.483.9522 Email: hannu.flinck@nsn.com
Target Dissemination Level:	Public
Status of the Document:	Completed
Version	1.0
Document location	http://www.psirp.org/deliverables/
Project web site	http://www.psirp.org/

SEVENTH FRAMEWORK PROGRAMME

European Commission
Information Society and Media

Table of Contents

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. All information in this document is provided "as is" and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability. For the avoidance of all doubts, the European Commission has no liability in respect of this document, which is merely representing the authors view.

1 Introduction

This deliverable lays out a set of migration paths towards information centric networking and finally to the full instantiation PSIRP architecture. Technology wise the migration could follow either an overlay, and/or underlay approaches, but all the critical steps need to have the right incentives to justify the required investments of the involved parties. Therefore the presented migration plan is building on the results from D4.6 which studied deployment incentives and Business Models. Deliverable D4.6 concludes the following:

1. Application sector investment is critical to the commercial success of an Information Cloud, for the network operator.
2. Demand for the capabilities of Information Cloud exists and could be very attractive to the sectors explored: Government, Collaborative Business ICT and Content-Centric.
3. In order to optimise the timing of network investments and to stimulate demand, a network operator would work closely with customers from the key sectors, to build the market opportunities.
4. A shim layer or information overlay approach should be pursued, in the first instance, with lessons learned being transferable to a later native Information Cloud deployment.
5. There should be an early focus from the network operator on developing media-rights management, probably using Packet-Level Authentication. This is expected to be a strong early driver for the Information Cloud and PLA forms an important under-pinning of the wider Information Cloud.
6. Working with Collaborative Business ICT would be the most likely scenario for developing the ‘native’ Information Cloud, firstly in a single-tenanted way (or multi-tenanted but with heavy safeguards, i.e. not general internet); moving eventually to a mass-market native information cloud for multi-media immersive environments.

We do not expect that a planned systematic migration will be carried out by a network operator, but rather a transformation will take place, driven by a set of incentives and market opportunities. There is no single specific migration path, instead multiple parallel routes depending on the business environment and the starting points.

This document is structured as follows. At first we briefly discuss how the overlay and the underlay approaches could evolve from existing technology building blocks. This consideration doesn’t take into account business incentives, but shows that there are technological “hooks” for a migration path. We then lay out how these approaches could be applied for migration through service deployment. This part is based on the vertical segments that were identified and studied in detail in D4.6.

2 Technical starting points for migration

2.1 Overlay of Rendezvous Service

Overlay networks are used for enhancing the basic infrastructure with extra functionality needed for a community or group of applications. Historically, they have played a role in the dynamic evolution of Internet technology. For example, the whole Internet was an overlay over a circuit switched network. Similarly, DNS, now a part of the critical Internet infrastructure, started as an application overlay network with its own protocols and servers. Functionality that is missing in the current Internet may be first offered as an overlay for those users that most require such enhancements that may not be available in the general Internet. Overlays, as opposed to application-specific network solutions, are seen as the mechanism to introduce functionality into the Internet.

The layerless and recursive PSIRP architecture enables the PSIRP rendezvous concept to have independent migration paths apart from the other PSIRP architectural elements, namely the forwarding and topology functions. For example, the rendezvous concept can be implemented and deployed as an overlay information discovery system over existing TCP/IP as such. As a matter of fact this was done during the earlier phase the WP3 rendezvous node development activity. As identified in the state of art description [D2.1] there exists precursors of rendezvous systems, such as the Siena event system [SIE] and the Common Object Request Broker Architecture [COR], but these system are acting on the application layer and have very limited interaction with the network to guide how the accessed information is passed through the network. Naturally, these application level rendezvous solutions will evolve towards more generic information centric concepts, if suitable incentives exist.

Migration to PSIRP rendezvous could start from a specific service deployment or from a separate vertical industry segment with very loose integration to the existing inter-domain network architecture. The notion of scoped information (Scope ID) provides a viable starting point from a legacy resolution system towards a PSIRP rendezvous system. The scope of the information can be used as a trigger to divert the information request to a specific rendezvous system built for that scope of information. An end user could try to access certain information first through the traditional means, but when the access request is identified to refer to a scope that has another resolution mechanism then the request is redirected to the PSIRP rendezvous system that resolves the request and matches it to the right publication. For example, when an end user tries to resolve an URL referring to certain site and content that is supported by the PSIRP rendezvous system the request can be identified during the DNS name resolution process (by the resolver or by the DNS server) as belonging to an information scope that should be resolved by the PSIRP rendezvous. Then the request is redirected (e.g. by use of the SRV resource record of DNS system) to a rendezvous node that acts as an entry point to the PSIRP system. This rendezvous node would then route the request according to PSIRP principles, based on the Scope ID, to the right rendezvous node hosting the rendezvous point for that particular information scope that was requested or subscribed for. The use of scopes of information offers a clean stepping stone towards the PSIRP rendezvous service. As the example above shows, a rendezvous service can work in parallel with the legacy systems to serve only specific types of information. Referring to figure 1 some of the “edge” RENE clouds could be any of the legacy resolution systems, enhanced with a capability to redirect all requests that match the specific scopes served by the PSIRP rendezvous system to the PSIRP “Interconnection overlay”. The scope ID could either an explicit separate attribute assigned according to PSIRP principles or the scope cloud be concluded implicitly from the request (e.g. URL or associated metadata) itself.

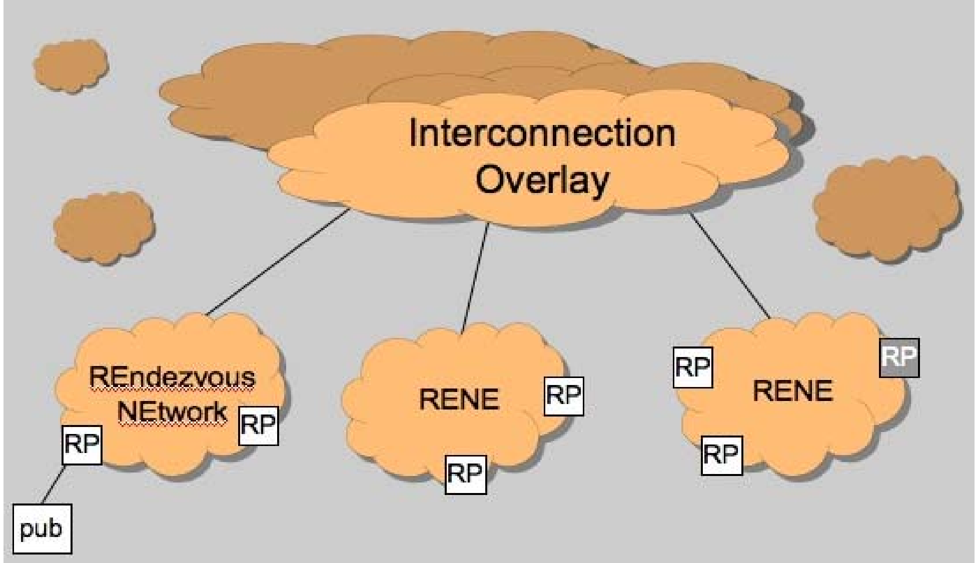

Figure 1 - Three tiered overlay rendezvous architecture [D4.5]

The key innovation of the PSIRP rendezvous system is its architecture and ability to function in the inter-domain context across multiple domains. The architecture divides the system into a generic anycast-like scope based discovery system and to a mechanism that allows the existence of separate scope-specific rendezvous systems. The scope specific rendezvous systems would likely use their own protocols and rendezvous nodes, while the generic rendezvous system would glue them together ability wise. In such a system, it would be the responsibility of the rendezvous point hosting the scope of a specific rendezvous system, to implement gateway functionality between the generic and the specific rendezvous systems. From the end user perspective the rendezvous concept would define a generic client-side protocol and an API to access the generic rendezvous system that in essence would enable the existence of multiple scope specific rendezvous points and forward rendezvous signals to the closest one

For example, Google-search, which enforces client-server principle, could instantiate ‘Google search’ scoped rendezvous points into different parts of the PSIRP rendezvous system. The PSIRP rendezvous system would be responsible for providing the ‘Google-search’ scoped rendezvous signals to the closest rendezvous point hosting the scope. The actual Google search parameters would be part of the signal. Based on the information in the signal the ‘Google search’ specific rendezvous point would know which real Google server should take over the processing responsibility of the signal.

The migration towards PSRIP rendezvous would add one of the missing parts that would make the Internet architecture more suitable for information centric networking. What is good in the rendezvous concept is that while there are applications (strongly client-server oriented applications like web banking) for which the data-oriented model may not fit well, the concept can co-operate with them and in some circumstances may even provide lower latency for the initial access.

2.2 Underlay migration of path selection and forwarding

Underlay migration refers to migrating lower layers of the Internet protocol stack. In order to introduce information centric PSIRP architecture this means replacement or extension of path selection, routing, forwarding and traffic engineering functions of the current IP/MPLS protocol stack.

2.2.1 Zfilter-based forwarding in the PSIRP architecture

In PSIRP, the forwarding is done by utilizing in-packet Bloom filters (iBF), named also as zFilters in the architecture. In a nutshell, identifiers of the links to be traversed by the packet are added into a small Bloom filter that is placed in the packet's header as an iBF. The intermediate nodes can quickly check whether their outgoing links are present in the packet's iBF by performing parallel bitwise AND operations between the iBF and all its outgoing interfaces' identifiers. If a match is found, the node forwards the packet on the corresponding interface.

To increase security, the link identifiers can be changed periodically over time as described in an enhancement called zFormation. The Bloom filter is bound to some information in the packet header and to a changing key at the forwarding node. In practice, this means that when the node performs the forwarding decision, it has to compute the link identifier for each link based on information in the packet, the incoming interface, and on the current key. The calculated link identifier is further compared with the packet's iBF, and if a match is found, the packet is forwarded out on the interface.

iBF-based forwarding is multicast-friendly, can prevent DDoS by acting as a capability, and supports also fast reroute [Zah2009]. Although it is a compact source routing-style representation of the path/tree, it inherits the probabilistic properties of Bloom filters, meaning that false positives can occur with a controlled probability. False positives mean unnecessary packet transmissions over links that are not included in the delivery tree. However, when utilizing opportunistic caching, this is not considered to be harmful. More information on zFilters and zFormation can be found in [Jok2009, Est2009, D2.3, D2.4]

2.2.2 Multiprotocol Stateless Switching (MPSS)

2.2.2.1 MPLS-TE and GMPLS

MPLS-TE (Multiprotocol Label Switching-Traffic Engineering) is a protocol set allowing the operator to control the resources of their network. This is achieved by balancing the load of the network through performing Traffic Engineering and offering high resilience to failures by allowing rerouting of traffic to backup paths right after detecting the failure. It acts as an enabler for implementing Virtual Private Networks (VPNs). MPLS uses LSPs (Label Switched Paths) to achieve these goals. In the border of the MPLS network, the packet gets a short MPLS label. On each forwarding node the label is inspected, and before forwarding it further, the node may change the label (label swapping), push another label in the header (label stacking) or remove the outermost label. GMPLS (Generalized MPLS) extends MPLS by allowing the generalized label to be e.g. time-slot (TDM), wavelength (WDM) or fiber. GMPLS thus enables the interworking of different data plane technologies by providing a unified control plane.

In the (G)MPLS protocol family, RSVP-TE is used both for establishing LSPs satisfying bandwidth and other constraints, as well as for signalling backup paths. OSPF-TE is used for information distribution, where link characteristics, such as unallocated bandwidth or switching capabilities are advertised via Opaque LSAs. Finally, LMP (Link Management Protocol) is a management protocol for maintaining control plane connectivity between two nodes and for providing link property correlation and fault detection.

In (G)-MPLS based Layer 3-VPNs (L3VPNS), customers' IP traffic is routed via the service provider's network using MPLS LSPs. The PE (Provider Edge) router receiving the IP packet from the CE (Customer Edge) router performs an IP-lookup for the corresponding VPN, determines the remote PE where the packet should be sent to, and adds two labels to the packet. The outer label will be used to forward the packet inside the operator's network while the inner label will be processed by the receiving PE. Each PE advertises the inner labels used to separate the customer networks at that edge using BGP. With this separation, the system scales with the number of VPNs as the core of the network is VPN-unaware.

However, when multicast communication is needed, the number of point-to-multipoint LSPs required grows exponentially with the number of PEs. To overcome this problem current practice is to use ingress replication (sending copies of the packet unicast to multiple PEs), inclusive trees (trees to be used only by one or more VPN's all multicast traffic), and selective trees (trees supporting a set of multicast groups for one or more VPNs). Data transmission can be done by IP or by MPLS forwarding on LSPs built by LDP or RSVP-TE. Some solutions involve bandwidth waste, as PEs get packets they do not need, others involve too much state in the core, and require too much complexity in dynamic conditions (e.g. changes in the subscriber set). Anyhow, a complex optimization algorithm should be run by the operator to find an appropriate point of operation.

2.2.2.2 Partial iBF migration using MPSS

It is clear that migrating to a full Bloom-filter based forwarding solution cannot happen overnight, as zFilters are not compatible with IP. Thus, we need to identify the areas where iBFs can be used, and create a plan showing how the network can incrementally evolve towards full iBF support.

MPSS (Multiprotocol Stateless Switching) is a short-term deployment possibility for iBF-based forwarding. MPSS operates in a similar way as MPLS, but with a significant difference in packet forwarding; in MPSS the MPLS labels are replaced by small Bloom filters, encoding the path, or the tree, that the packet needs to follow. Thus, in the default case, the state related to the Label Switched Paths (LSPs) is drastically reduced, because in MPSS, the iBF already holds sufficient forwarding information about the whole path, or tree, in the MPSS-enabled network.

For further details on the MPSS architecture we refer to [Zah2010], but in the following we briefly sketch two scenarios to emphasise the flexibility of MPSS. In the first scenario, consider a situation where the tree is requested by the source node with requirements, such as bandwidth constraints, from a remote Path Computation Element (PCE). The PCE computes the tree satisfying the constraints and replies with the tree information to the source. Using the received strict source routing information, the source initiates an RSVP-TE process, where the resources are reserved from the nodes on the path and the iBF is calculated hop-by-hop according to the forwarding decision. Optionally, the calculation can also be based on some flow information.

In the second scenario, the source node can compute the iBF directly using link identifier information received from an extended OSPF-TE. The extension adds the link identifiers in the advertisements. While OSPF-TE information is exchanged between all nodes in the network, any node can compute the tree and the corresponding iBF. If resource reservation is not needed, the iBF can be immediately used for communication, without any additional signalling delay (cf. RSVP-TE explicit routes with zero bandwidth reservation, where the hop-by-hop path setup is still needed to configure the forwarding tables).

2.2.2.3 MPSS in Multicast VPNs

MPSS offers stateless multicast, which can be seen as a potential forwarding solution in the service provider's network provisioning Multicast VPNs. It has the promise of easing the trade-off and the difficult process of fine-tuning when setting up and managing the multicast trees. A small penalty comes though because of false positives, i.e. a controllable amount of unnecessary packet forwarding due to probabilistic reasons. By exploiting the features of MPSS, the signaling inside the operator's network could be reduced, and because of the zero-signaling possibility, changes in the network (e.g. new members joining a multicast group in one VPN) could be handled faster.

Similar to MPLS-based VPNs, MPSS could act as a multiplexer layer connecting remote sites of different VPNs. Furthermore, the customer networks are not restricted to IP-networks only. They could have the opportunity to migrate to a native PSIRP solution (see 2.2.4 for further details), meaning that the MPSS network would serve legacy IP and PSIRP networks simultaneously. The architecture is illustrated on Figure 2.

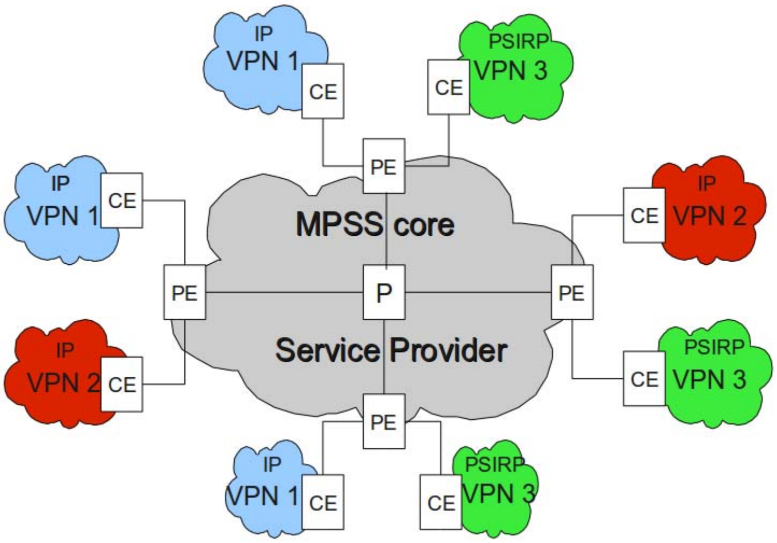

Figure 2 - An MPSS core network offering VPN services to customers. VPN1 and VPN2 are IP-based networks, while VPN3 is a native PSIRP network.

The ingress PE needs always to be aware of the receivers of the multicast groups. This means explicit tracking of members, which can be achieved using M-BGP or any other simple protocol that can transfer the information from the multicast protocol's join message towards the other PEs. This can be done by terminating the PIM process, sending a control message towards the PE closer to the source of the group and restarting the PIM operations at that PE. When a PE knows the receiver PEs, it can itself compute the iBF, or request a remote PCE to do the computation. Now, when a packet arrives, the ingress PE adds the inner MPLS label for identifying the VPN, and after that, the in-packet Bloom filter to route the packet through the operator's network to the receiving PEs. When an egress PE is reached, the iBF is removed and the VPN label is inspected. The reason for keeping MPLS as a service layer is to avoid cross-connection, i.e. the case when a false positive would cause the packet to enter a wrong VPN.

2.2.2.4 MPSS incremental deployment

iBF usage in customer networks

MPSS is mainly an intra-domain solution and could be introduced into service provider's networks offering VPN services without changing the IP layer in the customer networks. Furthermore, each customer can choose independently to migrate to in-packet Bloom filters-based forwarding. We can identify two scenarios; one scenario is to use a single iBF, containing the end-to-end tree, or as another scenario, we can use a stack of iBFs, where the customer network's iBF and the provider network's iBF are stacked in a similar way as MPLS labels in the MPLS based solution. The first solution has a drawback of the cross-connectivity threat and possible decrease in the performance because of increased number of links in a single filter. On the other hand, the second solution has longer packet headers.

MPSS and MPLS co-existence

So far we have discussed about converting the whole operator's network to support iBFs. However, we can also take a slightly different angle, where Bloom filters can act as generalized labels in GMPLS networks. It means that the LSPs can be built using different forwarding technologies in different segments of the network. Also, for inter-area and inter-domain MPLS one can build inter-area/inter-domain LSPs containing MPSS and traditional MPLS segments. MPSS segments can be used in areas where multicasting is needed and we can simplify the LSP creation and management process.

The inter-operability requires some additional operations on the routers connected both to the MPLS and MPSS segments. There are multiple ways how the interoperability can be implemented. When packets arrive from the MPSS area, we can define a special "link identifier" which actually means that the iBF should be removed from the packet. The forwarding decision can be done by a lower MPLS label, if it exists, or based on the IP header. In another variant, the special link identifier identifies directly the outgoing MPLS label the packet should get. Yet another variant is that the whole iBF identifies the MPLS label. Also, it is possible that the Bloom filter should not be removed; rather the MPLS label is pushed into the stack. This allows the usage of the same iBF also in other MPSS segments in the network, after removing the MPLS label. These scenarios are shown on Figures 3. and 4.

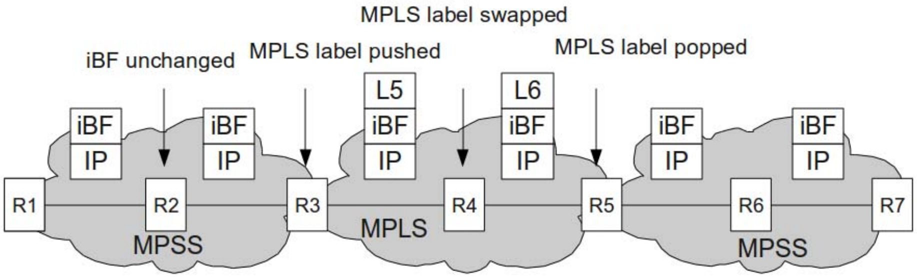

Figure 3 - MPSS and MPLS inter-operability. R3 keeps the iBF and pushes an MPLS label into the packet. R5 pops the MPLS label and the packet is forwarded by the Bloom filter in the second MPSS segment.

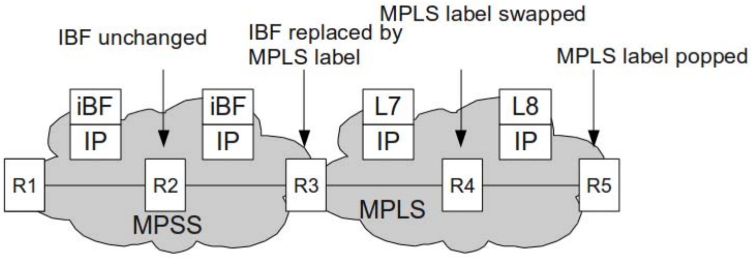

Figure 4 - MPSS and MPLS inter-operability. The in-packet Bloom filter is replaced by an MPLS label in the border of MPSS and MPLS areas.

For traffic from the opposite direction, the MPLS label can identify a pre-stored Bloom filter that should replace the MPLS label, or be pushed into the packet header, depending on the chosen solution. These scenarios are illustrated on Figure 5 and 6.

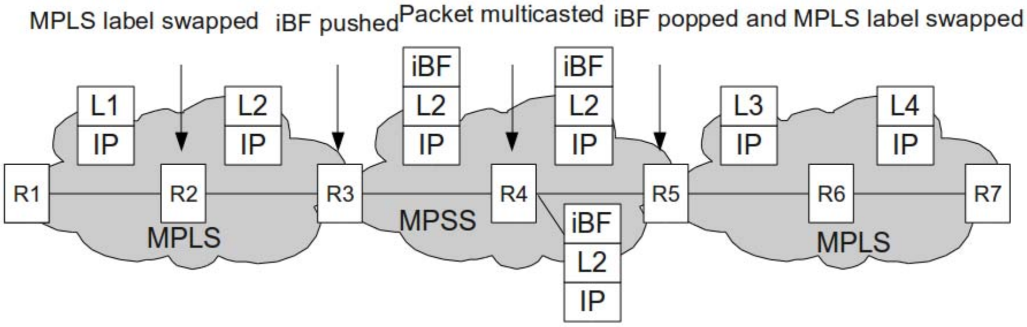

Figure 5 - MPLS and MPSS interoperability. MPLS label is kept by R3 and an in-packet Bloom filter is pushed into the packet. R5 removes the iBF and makes the forwarding decision based on the MPLS label.

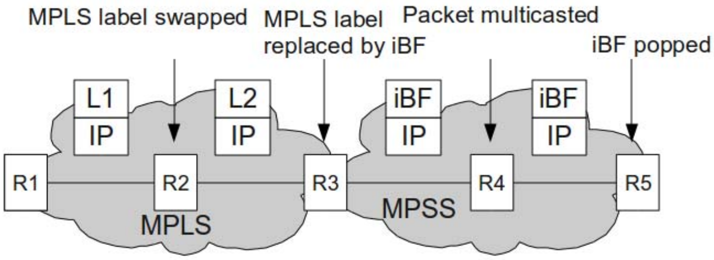

Figure 6 - On the border of the MPLS and MPSS networks, R3 removes the MPLS label (L2) and replaces it with an in-packet Bloom filter.

The abovementioned inter-operability scenarios hold the promise that MPSS could be incrementally deployed into existing MPLS networks. Also, MPLS and MPSS could co-exist in the same routers, and the operator can choose Bloom filters for multicast, and traditional MPLS for unicast traffic. Being deployed into the service provider's networks, it can act as an enabler of Multicast VPNs, by offering less complexity and less state than the current alternatives, with the small side-effect of some unnecessary forwarding due to false positives. As customer networks have the opportunity to migrate to iBF-based forwarding independently to others in their private networks, they could, as well, adapt other components of PSIRP, such as the rendezvous and topology modules.

2.2.3 Packet Level Authentication

Packet Level Authentication, as described in D2.3 and in [Lag2010], provides per packet public-key cryptographic operations at wire speed by virtue of new cryptographic algorithms (elliptic curve cryptography [Kob1987, Mil1985])

[Lag2010] discusses accountability issues of the current IP-based Internet. The described solution can also be applied in PSIRP. Basically, PLA introduces cryptographic identities and separates the accountability problem into two distinct parts: mapping the traffic to the cryptographic identity, and mapping the cryptographic identity to the real one. The PLA header ties the user's traffic to the cryptographic identity, and the trusted third party (TTP) ties the cryptographic identity to the real one. These tasks can be handled by separate entities increasing flexibility. While some major operators may also offer the TTP service, such a design would also allow lightweight operators that only convey traffic. Furthermore, the TTP service can be provided by other companies and organizations such as banks, credit card companies, and states.

Such an approach provides better security since it is based on cryptographic signatures. The privacy is also improved, since the user may possess multiple cryptographic identities and the user's real identity is stored by the TTP and not by every operator that the user utilizes. Finally, the burden of operators is reduced, since they do not need to store any data for accountability purposes. The operators should just verify that the user possess a valid certificate from a trusted TTP. In a case of misuse, the user's cryptographic identity is present in the PLA header along with the TTP's locator, therefore authorities can determine a user's real identity through the TTP and no co-operation from the operator is necessary.

PLA requires dedicated hardware for cryptographic operations in order to verify packets' signatures at wire speed. While the cost of mass produced dedicated hardware would not be high, the initial development and transition costs could be significant. However, PLA can be deployed gradually, and it would still offer benefits when only a small amount of routers would support PLA.

In the PSIRP architecture PLA is mostly used for securing control messages while using PLA for the payload traffic is optional, which simplifies migration to PSIRP. Furthermore, PLA was originally designed for IP networks; therefore PLA may be deployed independently of PSIRP. Some PLA-enabled infrastructure may therefore already exist before the wide-spread deployment of PSIRP.

3 Migration Through Service Development

In the previous sections, we outlined ways in which migration could proceed according to either an Overlay or Underlay solution. The Overlay solution is envisaged as a ‘shim layer’ or IP Overlay in which information-centric identifiers are used but are mapped to an underlying IP network. The Underlay solution is envisaged as leading to a ‘native’ Information Cloud, in which the full PSIRP architecture is implemented on the network, either within a Virtual Private environment, or within the internet at large. We now consider the vertical applications which were identified in D4.6 as key drivers for the Information Cloud and explore how they will motivate the need for Overlay or Underlay solutions, and therefore how the transformation of the network is likely to proceed.

As concluded in D4.6, for the network operator to get a return on investment in network transformation to create an information-centric Internet, application sector investment is critical. If the network operator provided functionality in the network, without market demand or market understanding or application enablers, the investment would be wasted. In considering migration, we assume that it will occur according to a succession of specific and relatively short-term market drivers; although it is expected that progress made in meeting the needs of early adopter applications will benefit later applications. In other words, we do not expect that a planned systematic migration will be carried out by a network operator, but rather each step of network transformation will have to make commercial sense. On that basis, we have drawn from the application sector drivers of D4.6, a logical progression of migration, which serves as an example of what could happen to achieve an Information Cloud. In so doing, we assume that there would be close collaboration between the network operator and its customers and their application needs, with the ideal being that the network operator would also write many of the solution enablers, components, interfaces etc., as this would be the most effective way to catalyse development and re-use of information-centric technologies.

In the following sections we reprise the main application opportunity areas for the Government, Collaborative ICT and Content-Centric sectors. For each application area, we discuss the most plausible type of information-centric solution and then summarise how each sector contributes to the migration story.

3.1 Government Sector Applications

We first look at the Government Sector opportunities in terms of migration towards an information-centric internet, or Information Cloud. (Note that this is largely from the UK perspective. Other countries may already have more advanced government systems in place, such as Austria’s eGovernment applications [GOE 2008].)

3.1.1 Removal of Departmental Silos

We consider the kind of steps that would be needed to go from a set of functionally and physically segregated monolithic departmental systems (data silos), to an information-centric (citizen-centric) infrastructure, allowing efficient data storage, access and manipulation.

In going from a situation where departmental systems are logically isolated and locked-down in security terms, it would be important to start by developing appropriate security techniques to enable only partial access to data associated with a given citizen.

Data would then need to be re-partitioned, associating all data held on a citizen, with that citizen. It would be necessary to allow a number of different identifiers to be used, i.e. for the system to translate between identifiers that were originally department-specific, such as health Ids. Once data was logically citizen-specific, it is likely that data stores would still be physically separate, for historical reasons. If so, then it would make sense, for efficiency of access, to initially partition the citizens according to an alphabetical list of names, between these data stores, for example. At this point, a simple version of Rendezvous could be implemented to handle government-specific search terms and publish/subscribe events. No network engineering would have been needed.

Mapping on to IP addresses would take place within a virtual private network. This is clearly an Overlay or ‘shim layer’ solution. A further migration could then take place to enable search terms to be handled across the whole government network, possibly via peer-to-peer methods, resulting in the integration of previously separate silos, which will ultimately be seamless.

3.1.2 Accountability

Maintaining accountability is especially important for citizen data. Simplistic approaches that grant access to all areas of a database to civil servants (government employees) are not appropriate for citizen data. In general, data should be visible only according to explicit lawful relevance, and should not be retained beyond that use, and should not be copied or forwarded. These requirements provide strong drivers for many of the ideas coming out of PSIRP. These can be summarised as the tagging associated with access to scopes; and Packet Level Authentication.

With Packet Level Authentication (PLA) every node in the network is capable of checking the authenticity and integrity of packets. If we combine PLA with tagging of scope Ids, which could include time to live constraints, or application constraints or readability constraints e.g. only on certain MAC addresses, we begin to see how accountability of data access can be achieved. We can enforce authentication of a person as they access a piece of information, leave an audit trail of where that data has gone, and prevent copying or re-use of that information, and hence prevent data being mislaid or falling into the wrong hands.

This functionality would be hugely advantageous. In terms of migration, the deployment of Packet Level Authentication would require significant investment and would provide important new functionality. However, it is not clear that a fundamental re-engineering of the underlying network would be necessary. Mapping to scopes could still take place as an overlay on to an IP network.

Furthermore, there is a demand to make users accountable for their Internet traffic. Currently, the IP address data retention mechanism tries to achieve such accountability. For example, in European Union operators must store necessary data (such as user's real identity, IP address allocated to the user, time frame when the IP address has been used, etc.) for 6 - 24 months [EU2006]. Such approach has several downsides in terms of security, privacy and flexibility. IP address data retention is not based on strong security measures and IP addresses can be spoofed. Storing user's personal details in multiple places for months or years introduces a significant privacy risk, especially since not all operators are able to store sensitive data correctly. Finally, this consumes operator's resources and increases barriers of entry, preventing lightweight operators from functioning.

The above mentioned method can also be used to implement the Internet-wide user authentication and roaming. In this case the trusted third parties (TTP) would form trust relationships between each other, and also would certify the network access providers. For example, if Google and MasterCard would make a co-operation agreement, then the user possessing Google's TTP certificate would be able to utilize a network certified by MasterCard on the other side of the world. Just as the roaming and billing is handled through the SIM card in current mobile networks, these tasks would be handled with the PLA header information and TTP certificate. The privacy would be further improved since the network access provider would not necessary learn a user's real identity at all. Such a scheme would significantly increase the competition and flexibility within the network.

In IP networks the above mentioned scheme requires that all traffic is secured by PLA. In a case of PSIRP, it may enough to just add the PLA header to the control traffic (bootstrapping, publish and subscribe messages) to provide adequate accountability. Therefore, the amount of necessary PLA-enabled hardware would be significantly lower.

3.1.3 Citizen-Centricity

We envisage that data has become logically citizen-centric, and has been partitioned into relevant scopes, such as health, work & pensions, address, passport, car registration etc., which are only accessible in appropriate ways. We then take this to the next step and consider the citizen having access to his or her own data, and even being able to input new values into that data directly, e.g. updating their address. This would be an ideological change that would make society arguably more democratic. It could also help to make government services fairer and more consistent; and ought to offer efficiency savings as well. Again, we would need the accountability safeguards described above. We would also need strong authentication mechanisms to ensure the identity of the citizen. This could be achieved via medical records and liaison with local health professionals, or even biometric readers.

In terms of migration, this third application is an extension of the first two and would not require significant new development. However, the benefits could be far-reaching, especially in the areas of social care, for example, enabling people to interface with professionals in a more equally balanced way, returning much more control to the individual.

3.1.4 Migration Via Developing Government Solutions

The Government sector is a logical early adopter of a simplified information-centric network. In many countries there are poorly structured databases, with inadequate accountability of data access, and limited flexibility. If convinced of the benefits of a new system, however, a significant contract would be placed that would enable technology development to take place, experiences to be gained, and re-usable components to be developed.

According to our analysis, this development would drive a simple form of Rendezvous within a constrained-scale deployment and a virtual private environment, in the first instance. There would be an exploration of the power of tagging scopes and publish/subscribe search mechanisms. There would also be a managed migration from physically distinct databases to a peer to peer merging of systems. More significantly, the applications would drive the development of innovative security techniques to enable the Information Cloud paradigm, by providing scope-limited access to information which is logically associated with an individual, both from a privileged position (government employee) and a democratic position (seeing one’s own data.)

To achieve accountability, Packet Level Authentication would almost certainly have to be implemented. This, in itself, would go some way towards creating a ‘native’ Information Cloud solution, but without having to significantly re-engineer the network nodes. A direct benefit of achieving this accountability of data access is likely to be in meeting the needs of Media Rights Management. If the use of a piece of content became inextricably linked to an individual or specific device, it would become much harder for people to engage in free peer to peer sharing sites. There would be a strong incentive for content owners to participate in any new technology that made it obvious that content had or had not been issued from its legitimate publisher. So, in providing solutions for Government, another significant application area would be activated.

3.2 Collaborative Business ICT

We now look at applications within the sector of Collaborative Business ICT, and how they contribute to the migration towards an information-centric internet.

3.2.1 Multi-Service Retail

A core driver to enable a company to be able to provide new service offerings to its customer base, in a cost-effective way, is the ability to make its business customer-centric. In many ways, this is an extension of the case described above in which the Government goes from being built around government departments, to a position of having a citizen-centric information system. The multi-service business has as its goal the ability to create new service offerings merely by writing a new application for the same single-truth customer data.

In order to achieve this, as in the Government case, it is critical that there are appropriate security safeguards to separate access to data associated with regulated industries, e.g. financial. The use of authenticated access to scopes would form an important part of this. The accountability of access to that data will need to be almost or exactly as rigorous as for the Government case. There are also clear parallels with the Government sector, in the need to enable customers to interface with a business system, to update personal details, and to request to see the information that is held on them. However, this application would require further development of the Rendezvous function, with a greater range of search terms and ontologies needing to be handled. It would also require the handling of different modes of operation. As well as being able to do basic data querying, a business would need to use an information system for a 'single version of the truth', real-time business intelligence, transactional applications and federated views across multiple lines of business. In these goals, there is huge overlap with emerging offerings in the Information As a Service space, and – as such – it would be important for developers of information-centric solutions to understand and perhaps partner with vendors in this space.

Initially we would envisage a business operating within a single-tenanted solution, with a large business owning its own virtual private network. However, it is also logical to expect that, over time, multi-tenanted solutions would develop, to enable de-risking of new service offerings (via incremental growth) and to enable smaller players to compete. This would require a more mature technology environment, offering greater safeguards to businesses to keep their records and transactions safe from their competitors. However, these are issues at the forefront of Cloud Computing, with which Information Cloud is likely to inter-operate, so we can envisage convergent solutions being developed. Note that we are still able to meet these requirements with a 'shim layer' or overlay solution to IP networks.

3.2.2 Dynamic Supply Chains

Dynamic supply chains represent a more sophisticated form of customer-centric information network, in which the network becomes, at another level, multi-business-centric, where we mean that multiple loosely-coupled businesses are working together in a supply chain. Ideally we want all businesses to be able to participate in supply networks with any other business, using common interfaces and common information models. Here, the driver is to be able to 'swap in' and 'swap out' suppliers so easily that the new supply chain configuration can operate for short periods of time, to meet local or transient supply needs, without protracted negotiation or system integration being needed. We envisage a step-change from current forms of supply chain set-up, even those that use ebXML registers. The opportunity that the PSIRP architecture offers is in the ability for rapid set-up of specific, constrained interactions amongst different (legal) entities. Here, we are leveraging publish/subscribe mechanisms to flag gaps in the supply chain and identify new suppliers, but – more importantly – we are leveraging the ability for businesses to be able to inter-work on a 'need to know' basis only. By this, we mean that other businesses are granted temporary access to their stock/customer/financial systems within clearly defined limits, according to the relevant scope of the current supply chain, only.

This application shows the importance of having an information-centric model of the business, at the level of its customers, its stock, and its operational systems. Business processes as well as physical entities become information items. Achieving this would require greater sophistication of Rendezvous points and pub/sub systems, than previously seen. Additionally, we would need third party secure authentication methods to validate the claims of business to meet supply demand, at appropriate quality. We can envisage trust networks developing, which would also operate on an information-centric model. We see the exploitation of a key aspect of the PSIRP architecture, which is the separation of governance from information. As such, it becomes more difficult to meet all these requirements from a ‘shim layer’ solution. We start to see real benefit from a deeper relationship between the physical network and the information network, particularly when we need to optimise performance. Multi-tenanted solutions would be necessary. A logical development of dynamic supply chains is being able to carry out dynamic service composition. Also, if we are considering real-time, seamless inter-operation amongst businesses over a shared or public underlying network, there is increased motivation to provide dynamic bandwidth allocation, in order to provide end-to-end business QoS. So, although an overlay solution could work in principle, the balance is beginning to tip towards an underlay or native Information Cloud solution.

3.2.3 New Markets

The application headed by ‘New Markets’ is the next stage in the evolution of business in terms of information-centricity. Here we envisage software components, customer data and operational data being information elements. Native Information Cloud gives us the opportunity to enhance all operational systems by its use of inherent auditing. Here we are thinking of the Packet Level Authentication accountability mechanism being applied to business processes, as well as information queries. Separate auditing functions are then not needed. The auditing of applications means that flows of service attributes: the performance of service components, their usage and billing can be captured without needing additional system layers. This can be used for managing stock and cash-flow levels as well as providing a clear understanding of the current pressure points of a system, in order to optimise its dimensioning.

The separation of operational functions from customer functions is also not needed. A business could become item-centric, providing the means to know the complete end to end lifecycle of a manufactured item, and the complete lifecycle of an item of food: ‘from farm to fork.’ However, this would simply be a different way of ‘cutting the information cake’. In other words, the business would be simultaneously finance-centric, customer-centric, process-centric, stock-centric etc.

In moving into new markets, e.g. in a foreign country, new service components would be needed to meet new regulatory and market pressures. However, services would be written as re-usable software components. The PSIRP architecture gives us ways to express policies to constrain the choice of component elements and the way that they are put together, to optimise service operations. This could vary and dynamically re-configure in real time. In applying an Information Cloud approach to software, as well as data, we begin to see how software components, system storage and ultimately the network capacity itself could adapt and self-organise to optimise commercial outputs. This level of optimisation and integration of systems would require a native Information Cloud solution, to be fully realised. Although a single-tenanted solution could work, within a data centre architecture, we would almost certainly see convergence with multi-tenanted Cloud Computing in order for all the functionality to be realised.

3.2.4 Migration Via Developing Collaborative Business ICT Solutions

In looking at Collaborative ICT solutions, we see a progressive evolution towards a native Information Cloud solution. We move from the multi-service retail case, in which an overlay solution is highly plausible and we mainly see increased sophistication in the Rendezvous and publish/subscribe mechanism; through dynamic supply chains, in which the need for integrated underlying networking at the information level begins to be attractive; through to the application referred to as ‘new markets’ in which a much deeper level of integration at the information, network and system level becomes persuasive. Here, we are not referring to permanent hard-wired integration, but the ability to have dynamic loosely-coupled relationships across a number of different logical and system planes. Although mapping to an underlying IP network would still be feasible in many cases, in order for operational performance criteria to be met, we begin to see real advantage in having content-centric routing, the ability to prioritise traffic according to business need, and the ability to rapidly re-configure trust and security relationships, leading to a new networking paradigm.

3.3 Content-Centric

Here we consider the part that content-centric applications have to play in the migration towards an information-centric internet.

3.3.1 Empowerment

The term 'empowerment' is used to label a number of content-centric drivers that enable new mass-market applications.

One application of information-centric network, in which individual pieces of content can have a rich level of tagging, is in providing much finer grain descriptors of content, and hence reach appropriate audiences. This would result in the ability to have much finer grain safeguards and censorship rules within institutions. This does not directly drive for content-centric networks. However, tagging that was linked to routing would be more powerful than tagging alone, in that it would mean that unwanted content was not only unavailable at the edge of the network, but also would not pass through a network. Though, if PLA were really water-tight, this would not matter.

From the industrial liaison discussions described in D4.6, emerged a driver to enable democratic TV formats, in which members of the public provide content to remote audiences. For this to work as a format, it is critical that content does not occupy bandwidth, without having an audience. This leads to a powerful driver not only to inextricably link subscribers to publishers, but also to do so in a way that optimises network usage. In other words, it provides a strong driver for content-centric routing, whereby content will not move in the network, unless requested; as well as dynamic bandwidth allocation to guarantee viewing QoS. Content with few subscribers will be poorly served. Content that is in demand will proliferate and be easy and fast to access. The driver to control bandwidth in these ways makes a native Information Cloud solution desirable.

Another outcome of industrial liaison was the driver to enhance content discovery, generally, and – more specifically – to empower people to set up their own Rendezvous points. There are several attractive features of a Rendezvous Point (RP), compared with a website: content can be drawn to the RP via trust networks, meaning that entities that are the most discerning, whether in terms of spotting trends, guaranteeing authenticity or having specific values etc., will attract the best content and the best click-through revenues. An RP should also be inherently easier to re-configure, in terms of the company's main message, brand and purpose, making it a more dynamic and cost-effective solution than having to maintain and occasionally completely re-build a website.

The widespread use of Rendezvous Points would not itself be a driver for a native Information Cloud solution, but as content provision relies on bandwidth, the ability to optimise the way that bandwidth serves content distribution, as described above, would greatly enhance the performance of Rendezvous Points, where media are concerned.

3.3.2 Truth

This section considers content-centric drivers that are related to content authentication.

The Rendezvous Point would be the new paradigm for breaking news, acting as a mechanism intermediate between unregulated publication on YouTube and publication via a news agency, which may be subject to political bias or simply difficult to access. This would also be an effective mechanism for 'whistle blowers', where their authenticity can be validated, whilst maintaining their wider anonymity. This application does not require a native Information Cloud solution, but – as it may evolve after Collaborative ICT solutions – it may be a late-adopter spin-off beneficiary of native Information Cloud technology.

Another aspect of authentication concerns the prevention of content manipulation by tagging key attributes of an image, e.g. the face. Industrial liaison discussions highlighted the importance of also having control of the context in which content is shown, so that its original message is not subverted. The first of these drivers can be met already using MP4 formats. However, the second driver makes the content audit trail uppermost, which is where the widespread use of authentication techniques built around tagging and PLA become important, which might be most effectively implemented using a 'native' Information Cloud.

3.3.3 New Realities

In order to create immersive new reality applications, there are a number of requirements of relevance to information-centricity, as discussed below.

Information tagging would enable meta-data to be provided to end-users and their devices, and this is expected to be critical to rendering 3-D models, created by user actions, to provide real/virtual multi-media environments. This is not in itself a driver for a native Information Cloud solution, but as soon as we consider performance, the use of an IP overlay solution looks less viable. A native Information Cloud solution, with content-centric routing, should mean that gaming software components and player data would proliferate and move towards the players to optimise the gaming experience.

Another key factor, which applies to all New Reality applications, would be the ability to dynamically reserve high bandwidth links, for the duration of the virtual reality relationship, and to be charged accordingly. Tagging of information (software, data etc) via the Rendezvous identifier, could be used to do rapid prioritisation of traffic, without needing to do deep packet inspection or to know who the traffic is coming from. The subscriber to publisher relationship would independently provide governance and payment. The same considerations would apply to other co-operative New Realities apps, such as augmented reality for product design or teaching specialised skills, such as surgery.

In summary, for reasons of performance, New Realities applications are likely to need a native Information Cloud solution.

3.3.4 Migration Via Content-Centric

The main drivers for a native Information Cloud solution in the content-centric sector are around performance: both content-centric routing (with implied self-organisation of software and data) and dynamic bandwidth allocation for guaranteed end to end QoS, are key to this. Even with the roll-out of high bandwidth broadband, it seems inevitable that bandwidth optimisation will always be needed.

A few content-centric applications, for example in the medical sphere, could be delivered via dedicated virtual private network solutions. However, what is primarily envisaged in the content-centric sector is mass-participation, whereby, for a fee, or in exchange for viewing advertising, the general populace can interact in new multi-media. As such, however, there is no obvious coherent driver for roll out of native or even shim layer Information Cloud. Therefore, it is expected that the majority of these applications will only become mainstream after the early adopter communities of Government and Collaborative Business have driven technology development and network transformation. However, business models for future technologies are notoriously difficult to predict.

4 Conclusion

In this document we have attempted to build on the conclusions and insights drawn from D4.6 to consider in more detail the likely path of migration towards an information-centric Internet. In so doing, we have tried to identify what are the key dependencies for each application opportunity identified, so that we can clearly distinguish what drivers could be satisfied by simple modification of existing networks, and where significant re-engineering would be necessary. We have also tried to put the applications on a logical time-line, in order to be able to see which applications make sense to consider first. We have also stressed the need to identify commercial drivers for each application, rather than assuming that the sum of drivers would enable a coherent migration towards the most highly transformed network.

References

[COR] OMG, “The Common Object Request Broker Architecture”, at http://www.omg.org, 2009

[Est2009] C. Esteve, P. Jokela, P. Nikander, M. Särelä, and J. Ylitalo, "Self-routing Denial-of-Service Resistant Capabilities using In-packet Bloom Filters", Proceedings of European Conference on Computer Network Defence (EC2ND), 2009.

[EU2006] European Parliament, Directive 2006/24/EC.

[GOE2008] I Goetzl, T Greeching and G Fisher, “eGovernment in Austria and Vienna: Progress by vertical co-operation” in A Shark and S Torporkoff, Beyond eGovernment and e_Democracy: A Global Perspective, Booksurge Publishing, 2008.

[D2.1] A. Karila (ed], PSIRP Deliverable D2.1: State-of-the-Art Report and Technical Requirements, 2008

[D2.3] Ain M. (ed), PSIRP Deliverable D2.3: Architecture Definition, Component Descriptions, and Requirements, 2008

[D2.4] M. Ain (ed.), PSIRP Deliverable D2.4: Update on the Architecture and Report on Security Analysis, 2009

[D3.5] P. Jokela (ed.), PSIRP Deliverable D3.5: Final description and evaluation implemented components: including open source deliveries, April 2010

[D4.5] J. Riihijärvi (ed), PSIRP Deliverable D4.5: Final Architecture Validation and Performance Evaluation Report, 2010

[D4.6] Final Evaluation Report on deployment incentives and Business Models, 2010 (to be finalized).

[D5.3] H. Flinck (ed.), PSIRP Deliverable 5.3: Dissemination Report, 2008

[D5.4] H. Flinck (ed.), PSIRP Deliverable 5.4: Dissemination and Exploitation Report, 2009

[IT] L. Maxwell, It’s ours; why we, not government, must own our data, Centre for Policy Studies, http://www.cps.org.uk/cps_catalog/it's%20ours.pdf, 2009

[Jok2009] P. Jokela, A. Zahemszky, C. Esteve, S. Arianfar, and P. Nikander. "LIPSIN: Line speed publish/subscribe inter-networking", In Proceedings of ACM SIGCOMM’09, Barcelona, Spain, Aug. 2009.

[Lag2010] D. Lagutin, and S. Tarkoma, “Cryptographic Signatures on the Network Layer – an Alternative to the ISP Data Retention,” In Proceedings of the IEEE Symposium on Computers and Communications (ISCC'10), Riccione, Italy, June 2010.

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

[Zah2009] A. Zahemszky, S. Arianfar. "Fast reroute for stateless multicast". In Proceedings of International Workshop on Reliable Networks Design and Modelling (RNDM), St. Petersburg, Russia, Oct. 2010.

[Zah2010] A. Zahemszky, P. Jokela, M, Särelä, S. Ruponen, J. Kempf, and P. Nikander, "MPSS: Multiprotocol Stateless Switching", in proceedings of 13th IEEE Global Internet Symposium 2010, San Diego, CA, USA, March 19, 2010.