Publish-Subscribe Internetworking Security Aspects

Nikos Fotiou, Giannis F. Marias and George C. Polyzos

Abstract Publish-Subscribe is a paradigm that is recently receiving increasing attention by the research community, mainly due to its information oriented nature. Although the publish-subscribe paradigm yields significant security advantages over the traditional send-receive one, various security and privacy challenges are raised when it comes to the design of an internetworking architecture that is solely based on this paradigm, such as the Publish Subscribe Internet (Ψ) architecture. Ψ is the main outcome of the Publish-Subscribe Internet Routing Paradigm (PSIRP) project, which was launched with the ambition to develop and evaluate a clean-slate architecture for the future Internet based on the publish-subscribe paradigm. Availability, security, privacy and mobility support are considered as core properties for this new form of internetworking, instead of being provided as add-ons, as in the current Internet. This paper discusses the security and privacy properties of and challenges for publish-subscribe internetworking architectures and specific techniques and solutions developed in PSIRP for Ψ.

1 Introduction

The Publish-Subscribe paradigm has been in the spotlight of recent research efforts. Its information oriented nature, the decoupling it offers between information providers and information consumers as well as its location-identity split properties, have inspired a variety of–mainly overlay–architectures that focus on multicast [6], mobility [15], indirection [29] as well as on caching [16].

Publish-Subscribe architectures are composed of three main components; publishers, subscribers and a network of brokers [8]. Publishers are information providers that 'publish' information (advertisements). Subscribers on the other hand are information consumers that express their interest in specific pieces of informa-

Athens University of Economics and Business, Mobile Multimedia Laboratory, Patision 76, Athens 104 34, Greece, e-mail: {fotiou,marias,polyzos}@aueb.gr

tion by issuing subscriptions. Brokers are responsible for matching publications with subscriptions and initiate the (information) forwarding process from information providers towards information consumers. The broker, responsible for the publication-subscription matching, is often referenced to as the rendezvous point and, therefore, the network of brokers is usually referred to as the rendezvous network. Publication and subscription operations are decoupled in time and space allowing for the support of mobility as well as anonymization mechanisms. Moreover a publication can be provided by multiple nodes and similar subscriptions can be aggregated, creating opportunities for multicasting and multihoming. Inherently, the publish-subscribe paradigm has many security advantages compared to the commonly used end-to-end, send-receive oriented paradigm.

PSIRP (Publish-Subscribe Internet Routing Paradigm),[1] an EU FP7 funded research effort, has designed, implemented in prototypes, and initially evaluated a clean-slate, information oriented future Internet architecture; we call it the *Publish-Subscribe Internet* (PSI) architecture, Ψ for short. This architecture aims at overcoming most limitations of the current Internet and at emphasizing the role of information as the main building block of the (future) Internet. This new architecture is based on a paradigm completely different from the current one. Ψ is based on pure, through-the-stack application of the Publish-Subscribe paradigm. Moreover by abiding to the Trust-to-Trust (T2T) principle [4], i.e., all functions take place only in trusted points, the Ψ architecture considers security as a building block of its architecture rather than as an 'add-on'. Ψ harvests the security advantages the publish-subscribe paradigm offers, whilst Ψ-specific security mechanisms are also incorporated.

The purpose of this paper is twofold: to give an overview of the security features of and challenges for the publish-subscribe paradigm, as well as to show the additional techniques and mechanisms developed in PSIRP in order to secure the Ψ architecture. The remainder of this paper is organized as follows. Section 2 highlights some of the security and privacy challenges that exist in publish-subscribe architectures. Section 3 presents the security advantages of the publish-subscribe paradigm. Section 4 overviews the Ψ architecture and its specific security solutions. Section 5 investigates how other, related architectures handle security requirements. Finally, our conclusions as well as ideas for future work are presented in Section 6.

2 Security and Privacy Challenges in Publish-Subscribe Architectures

As previously mentioned, in the publish-subscribe model, producers publish event notifications to announce information availability and consumers subscribe to specific information items to explicitly declare their interest. Matching is achieved through the rendezvous network, which is envisioned as a distributed service that

[1] http://www.psirp.org

spans over a large number of providers and administrative domains. In the case where one or more matches are provided by the rendezvous service, then a particular sub-graph over the network topology is determined and activated to support a multihomed and multicasted communication service that transports information elements from publisher(s) to subscriber(s).

Security issues and requirements that arise in a global-scale publish-subscribe system have already been extensively addressed. Wang et al. [31] as well as Lagutin et al. [21] have specified security requirements for a publish-subscribe architecture, whereas Wun et al. [33] have identified and classified possible DoS attacks in content-based publish-subscribe systems. Various mechanisms have been developed in order to secure publish-subscribe systems–such as Eventguard [28]–and most of them base their operation on traditional security mechanisms, adapted to the concept of the publish-subscribe paradigm. In this paper we are focusing on security, trust, and privacy requirements focusing on a different level of abstraction and trying to enrich the existing work with recent results for the publish-subscribe paradigm.

In the information level, integrity, authenticity and validity of information are required. Integrity protection methods will ensure that any violation or fabrication of information elements' content will be detectable. Authenticity means that the information that is received by the subscriber is identical with the subscriber's initial request, and it is not forged. Validity means that the information items announced by the publisher, matched with the subscriber's request, and then forwarded to the subscriber are identical. Detecting integrity violation is a task that mainly is based on public key certificates and signatures, and, thus, it requires trusted third parties or bilateral trust (e.g., symmetric secrets, or HMAC key-based approaches). On the other hand, publication and subscription operations might be decoupled in time. Thus, subscribers might never recognize the publishers' identities, or even their certificates. Thus, information integrity verification should be assisted by the rendezvous-network. In order to avoid bottlenecks due to processing or signing every information element, rendezvous nodes might produce sequences of integrity evidences, such as TESLA seeds if a TESLA approach [25] has been adopted between publication end-points and consumers. Verifying authenticity and validity of the information requires a different, reaction-oriented approach, which is based on subscriber's evaluation on the received information. Such an approach will rank published information elements, and recommend the accurate ones, avoiding DoS attacks [10] or spamming [11].

At the application layer, a main security challenge is the design of a mechanism that grants to subscribers the appropriate access privileges to publication announcements. This is akin to making confidential the existence, and not the content, of publications. Assuming that publishers are always privileged to submit events and announcements, the rendezvous network should enforce an access framework that makes the notifications reachable to preferred subsets of subscribers. For application-level access control, such subsets are formed using scopes [9], role-based access control [3] [26], as well as identification and authentication schemes [24]. On the other hand, publication content confidentiality is achieved mainly through encryption. Finally, when a forwarding topology will be deployed to

transport information to subscribers, then there is a potentially strong anonymity requirement to unlink the information and the publisher and subscribers among themselves and from the networking attachment and relay points.

From the subscriber's privacy point of view, a central objective is to unlink his identity from his subscription interests, e.g., by supporting anonymous subscriptions. Subscription privacy might rely on an anonymity framework related to trusted proxies (anonymizers) that receive and process the original request, change its time reference, hide the subscriber's identity and obfuscate his network attachment point. This approach might introduce significant delays, but fulfills the demand for strong anonymity support at the network layer. Additionally, such a system should be designed and deployed appropriately to avoid attacks that have been reported on mix-based privacy enhancement approaches, such as traffic analysis, blending and trickle attacks [32].

3 Publish-Subscribe Security Features

The publish-subscribe paradigm can be seen as a remedy to the imbalance of power between senders and receivers in the traditional send-receive paradigm. With the original Internet architecture, the network will make a best effort attempt to deliver whatever any sender sends, irrespective of the interest of and no matter the cost for the receiver and the network(s). This imbalance is often accused for the increasing number of (Distributed) Denial of Service (DDoS) attacks, as well as for the emergence of spamming. In publish-subscribe systems there is no information flow as long as the receiver has not expressed interest on a particular piece of information, i.e, the receiver in a publish-subscribe architecture is able to instruct the network which pieces of information shall be delivered to it. Moreover, and even though the model is so powerful so that there can be subscriptions before the corresponding publications have been published, no information is requested from a publisher, unless the publisher has explicitly denoted the availability of that information, i.e., not before the publisher has issued a publication message (for this particular piece of information).

Publication and subscription operations are decoupled in time and space, i.e., they do not have to be synchronized neither do they block each other. Moreover publishers and subscribers do not communicate directly and they can hide their identity as–in general–subscribers are only interested for the information itself rather than on who provides it, and publishers–usually–disseminate publications using multicast so they cannot (and usually should not) be fully aware of the publication's recipients. Therefore, anonymity can be easily achieved in publish-subscribe architectures. Moreover having a point in the network where subscription and publications are matched, effective deployment of access control mechanisms is enabled.

Publish-Subscribe architectures offer great availability. The rendezvous network of a publish-subscribe architecture is usually implemented using a DHT. DHTs provide significant load balancing–usually at the cost of some communication stretch.

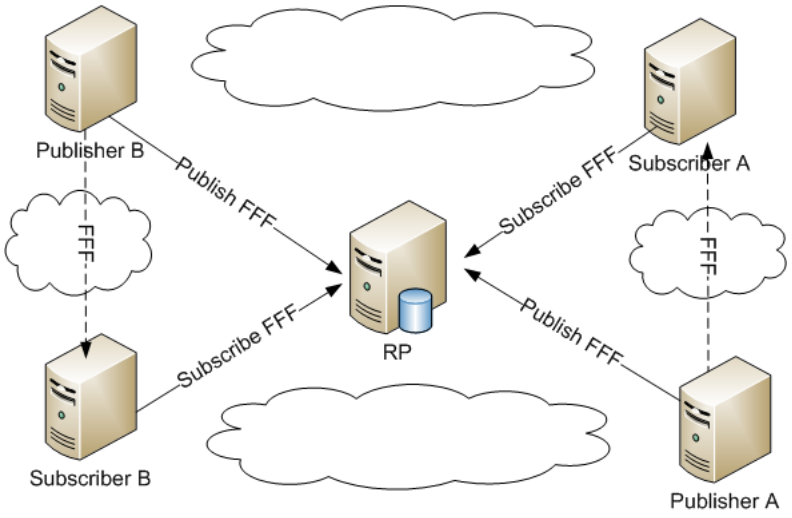

Fig. 1 Example of multihoming in a publish-subscribe architecture

Moreover in a publish-subscribe architecture multihoming can be easily achieved, as multiple publishers may advertise the same publication to a Rendezvous Point (RP), therefore a RP has a number of options with which it can satisfy a subscription. Figure 1 shows an example of multihoming in a publish-subscribe architecture. Publishers A and B, both publish publication FFF. Subscribers A and B subscribe to this publication. For each subscription message the RP knows two publishers that can provide the publication matched, therefore for each subscription message it could choose the publisher that is closer (in any sense) to the respective subscriber, e.g. here, it chooses publisher A to serve subscriber A and publisher B to serve subscriber B.

Publish-subscribe architectures allow for subscription aggregation and they create opportunities for multicast to be useful, therefore in these architectures resource sharing can be achieved, leading to greater availability. In Figure 2 both subscribers A and B subscribe to publication FFF. The subscription messages are aggregated within the networks, when following the same path towards the RP. Moreover publisher A forwards a single data flow, which is copied (bifurcated) in an appropriate place in the network in order to serve both subscribers.

4 The Ψ Architecture

The core element of the Ψ architecture is information; information is everything and everything is information [30]. In Ψ every piece of information is identified by a unique, flat, self-certified identifier, known as the *rendezvous identifier* (RId).

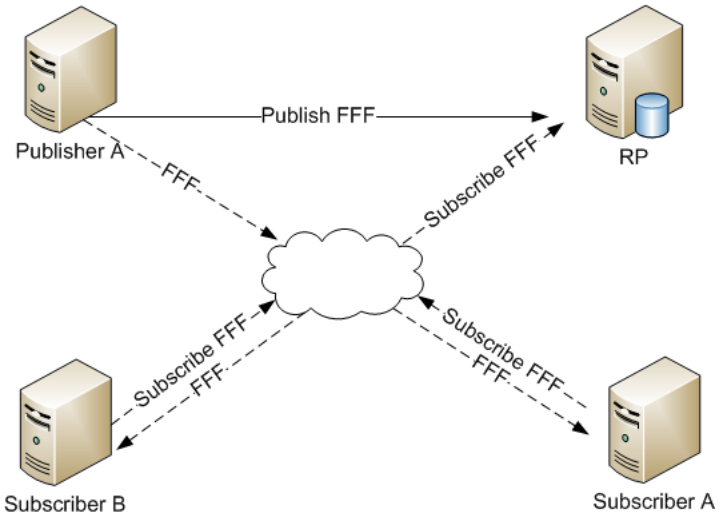

Fig. 2 Resource sharing in a publish-subscribe architecture using subscriptions aggregation and multicast

Information is organized in *scopes*. Scopes are physical or logical structures that facilitate the finding as well as access control over a piece or collection of information. A physical scope can be for example a corporate network, whereas a logical scope can be a group of friends in a social network. Scopes can be included within each other, creating a flexible structure. Scopes are identified by a flat identifier known as the *scope identifier* (SId). Each SId is managed by a rendezvous point (RP) which can be a single *rendezvous node* or a large *rendezvous network*.

The publication operation in Ψ involves 3 steps [12]; initially the SId of the publication scope is identified, then the RId of the publication is created and, finally, the publication is published in i.e. the publication message, including the Rid and Sid, is sent to the RP responsible for handling this SId. The publication message may also contain *metadata*–such as size of the data, encoding and other general information about this publication. Figure 3 shows the publication operation in a Ψ network with three scopes; the scope MyUniversity and its sub-scope MyLab and the scope MyFamily. As it can be seen in this figure, a publisher issues a publication to the scope MyFamily. The publication message should contain a scope-unique publication identifier (RId), the MyFamily scope identifier (SId) as well as metadata that describe this publication. The publication message reaches the rendezvous node RN B, which is part of the MyFamily rendezvous network.

The subscription operation involves the identification of the SId and RId of a publication–which can be done, for instance, with the help of a search engine–and the sending of a subscription message. Initially the subscription message will be forwarded to the appropriate scope as all the other scopes are not aware of the publication in question. When the subscription reaches the appropriate scope it will be

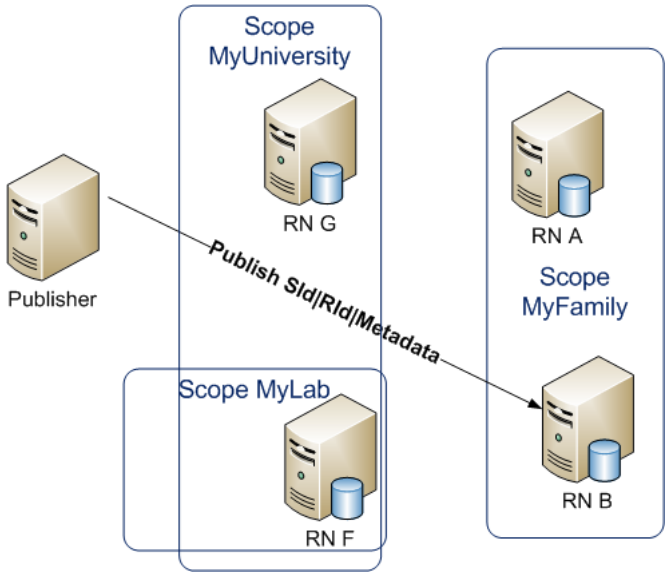

Fig. 3 Publication in a Ψ network

forwarded to the publication RP. The network is responsible for routing publication and subscription messages towards the RP as well as for forwarding publications from publishers towards subscribers. Figure 4 shows the subscription operation. A subscriber subscribes to an already published publication. When the subscription message reaches the appropriate RP, and as long as there is a publication that matches this subscription message, the RP creates a forwarding path, from the publisher towards the subscriber, and instructs the publisher to send the publication using a specifically created identifier (FId) for this path. A forwarding path is realized through zFilters [14], a Bloom filter based structure that contains the link identifiers that a data packet must traverse in order to reach its destination(s). Ψ uses a slow path for signaling, i.e., publication and subscription messages, and a fast path for data forwarding. Moreover multicast is the preferred delivery method.

4.1 Ψ-Specific Security Mechanisms

Security in Ψ plays an important role and trust is at the center of a Ψ declared principle. Security mechanisms are considered at all levels of the architecture. Information in Ψ is transmitted in encrypted packets using the Packet Level Authentication (PLA) technique [19]. PLA is a novel mechanism, applied in Ψ, for protecting the network based on the assumption that per packet public key cryptographic opera-

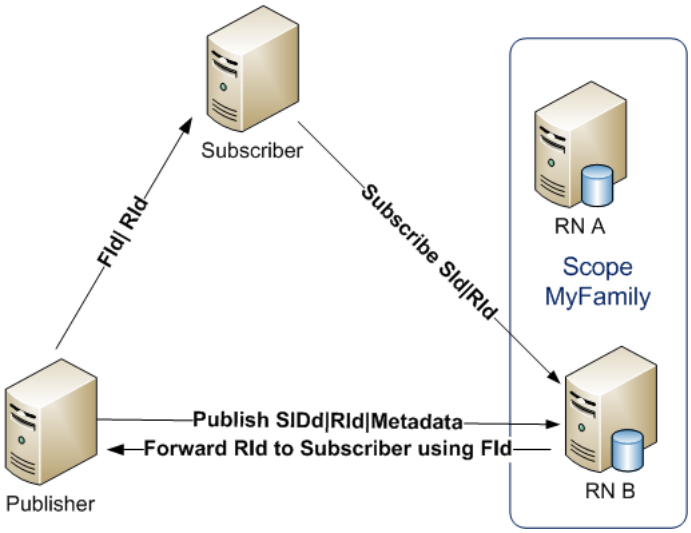

Fig. 4 Ψ subscription: initially (e.g.) the publisher issues a publication, then a subscriber, subscribes to this publication and the rendezvous point instructs the publisher to forward this publication to the subscriber

tions are possible at wire speed in high speed networks with the help of new cryptographic algorithms and advances in semiconductor technology. Moreover when applied in wireless environments PLA has been proven to offer significant energy efficiency [20].

As already described Ψ's forwarding mechanism is based on the formation of a Bloom filter–called zFilter–that describes the path that a data packet should follow [14]. The computation of the zFilter is based on the identifiers of the links that compose the data path. These identifiers are dynamically generated every time a zFilter is created, making this way almost impossible for an attacker to create crafted zFilters or link identifiers that will lead to DoS attacks or to information leakage. Forwarding using zFilters is achieved at line speed, leading to excellent performance and scalability. Network attachment in Ψ [17] assures proper user authentication protecting both users from improper configuration as well as the network from (D)DoS attacks that can be caused by malicious users who repeatedly try to attach themselves to a Ψ network.

At the higher layers of the architecture, existing security mechanisms can be used. Nikander et al. [23] studied the application of existing work on cryptographic protocol analysis in a pure publish-subscribe architecture and found out that, even if networking protocols are revised very drastically, current cryptographic protocol analysis can be applied to a certain extent, with only minor modifications, mostly on the notation side. Moreover, novel trust mechanisms should be considered applied to information ranking [11] rather than ranking end-users.

Ψ security is going to be primarily based on the notion of scopes. Although not yet fully designed and implemented, scopes are expected to control information dis-

semination as well as to play a significant role in applying access control policies, as well as accounting mechanisms. Scopes are expected to be Ψ's information firewalls.

5 Security Aspects of Comparable Internetworking Architectures

CCNx [7] (Content-Centric Networking, now termed Named Data Networking: NDN) is an ongoing research project that investigates the potential of an information-oriented Internet architecture. In contrast to Ψ, CCNx proposes an architecture organized using hierarchical naming [13]. Moreover CCNx uses a broadcast-based mechanism for information location, rather than a rendezvous driven one. CCNx does not rely on flat self-certified identities, it rather uses a scheme that assures the relationship between publications and their identities and it provides validity, provenance, and relevance [27]. In this scheme every publisher is allowed to generate a user-friendly tag label for their publication, which in a next step is incorporated into the body of the publication as a digital signature. This digital signature is generated by applying the publisher's public key over the publication's data and the publication label. When a subscriber receives the publication, and provided that the publisher is reliable, he is able to verify that the publication he received matches its label. On the other hand in case of a malicious publisher that uses forged labels, this publisher can be held accountable for his behavior, as its public key has been used in order to generate the publication's digital signature.

The Data-Oriented Network Architecture (DONA) [18] and Routing on Flat Labels (ROFL) [5] are two pioneering architectures that introduced flat identifiers. DONA aims at replacing DNS with flat self-identifying labels that will enable data location and retrieval. In contrast to Ψ, DONA uses the same path, for information location and forwarding. DONA's main security mechanism is its self-certified naming. DONA names are organized around principals and they are of the form P:L, where P is the cryptographic hash of the principal's public key and L is a label chosen by the principal, who ensures that these names are unique. Every publication is accompanied by a metadata file that includes the principal's public key as well as her digital signature over the publication data. Users in DONA are expected to learn a publications' name using external, reliable mechanisms. In order to defend against DoS attacks, DONA relies on IP-level mechanisms, as well as on the limits that providers will pose on users' publications and subscriptions. Finally DONA assumes the existence of third trusted parties for public key status retrieval and revocation.

ROFL creates an internetworking architecture in which routing takes place solely based on the data–flat–identifiers. In ROFL there is no information hierarchy, as there is in Ψ (with the usage of scopes) and DONA. ROFL security is also based on self certified identities. In ROFL, in every network node, i.e., router or host, a unique ID is assigned, which is tied to a public-private key pair. This key pair is used to

sign-verify every packet that traverses the system. ROFL secures its routing infrastructure by using the so-called *filtering* and *capabilities* techniques. With *filtering*, every host can control its reachability and therefore filter out malicious hosts. With *capabilities* the architecture is able to perform fine-grained access control. Whenever a (legitimate) host requests the creation of a network path, a *capability* token is provided, which proves that the host has the proper access control credentials for this path. *Capability* is a cryptographic token designating that a particular source (with its own unique object identifier) is allowed to contact the destination.

The Internet Indirection Infrastructure (i3) [29] and the Host Identity Protocol (HIP) [2] are two rendezvous-based overlay solutions that aim at supporting mobility, multicast and multihoming. Ψ's rendezvous and topology processes use similar concepts, at all levels of the architecture.

i3 implements an IP overlay network that replaces the point-to-point communication model with a rendezvous-based paradigm. In i3 *sources* (akin to Ψ publishers) send packets to a logical identifier, whereas *receivers* (akin to Ψ subscribers) express interest in packets by inserting a trigger into the network. A distributed lookup service is responsible for matching triggers with packets and an overlay network of i3 nodes is responsible for forwarding packets. An i3's extension, known as the *Secure-i3* [1], further enhances the security of the proposed architecture by allowing hosts to hide their IP address as well as to defend against DoS attacks without introducing new vulnerabilities. IP address hiding is accomplished with the usage of the so-called *private IDs*; when an end-host issues a new trigger, instead of using its real IP address, it uses the public ID of an i3 (reliable) node that acts as the end-host's representative. The public ID of this i3 node is the private ID of the end-host. Even if the representative node removes its public ID it will not affect the already established end-host's connections. Every node in i3 may have multiple public IDs. In case of DoS attacks a node may remove all of its public IDs to eliminate the attack, or remove some of them in order to mitigate the attack. Moreover, puzzles can be used as a countermeasure against DoS attacks; before a suspicious host is allowed to send a new packet, it is requested to solve a cryptographic puzzle. Finally, hosts in i3 can manipulate the path that a packet should follow in order to reach them, this way they are able to circumvent parts of the network that are under attack.

HIP introduces a new layer that decouples host identity from location identity in the internetwork stack, between the IP layer and the transport layer. When HIP is used, the applications no longer connect to IP addresses, but to separate *Host Identifiers*. A Host Identifier is cryptographic hash of the host's public key, which in turn, is used for securing communication between hosts. The resolution from a Host Identifier to an IP address can be achieved either by using a DNS-like mechanisms or a DHT. *Host Identity Indirection Infrastructure* (Hi3) [22] is the secured version of the HIP protocol, which utilizes Secure-i3's rendezvous principles. Secure-i3 is used in order to perform Host Identifier to IP address resolution, whereas IPSec is used for the rest of the communication between hosts.

6 Conclusion and Future work

The Publish-Subscribe paradigm achieves a significant shift from the current end-host driven internetworking towards an information oriented Internet architecture. This paradigm offers significant security advantages, including greater availability and enhanced privacy. The opportunities for multicast, mobility support and caching, as well as, the decoupling it offers between the communicating parties, make the publish-subscribe paradigm a strong candidate for a future internetworking architecture. Nevertheless various security and privacy challenges remain and further research is needed in order to identify and tackle them. Towards this direction the PSIRP project has created the so-called Ψ architecture; a clean slate Internet architecture that is based on the publish-subscribe paradigm. The Ψ architecture demonstrates the significant capabilities of this paradigm and through the development of Ψ-specific security mechanisms shows the road towards a secure future internetworking architecture.

The research in this field is a very active ongoing effort. Various research projects around the world investigate the potential of new internetworking architectures based on the publish-subscribe paradigm–or other similar ones. Security remains in the spotlight of all these research efforts. As far as the Ψ architecture is concerned, its research and development continues during the EU FP7 PURSUIT[2] project, which plans to further explore security, privacy and trust issues of this architecture, as well as to create novel mechanisms and evaluate them including aspects of them experimentally over the newly established Ψ tesbed which spans Europe.

Acknowledgements The work reported in this paper was supported in part by the FP7 ICT project PSIRP, under contract ICT-2007- 216173.

References

1. Adkins, D., Lakshminarayanan, K., Perrig, A., Stoica, I.: Towards a more functional and secure network infrastructure. Tech. Rep. UCB/CSD-03-1242, EECS Department, University of California, Berkeley (2003). URL http://www.eecs.berkeley.edu/Pubs/TechRpts/2003/6241.html
2. Al-Shraideh, F.: Host identity protocol. In: Networking, International Conference on Systems and International Conference on Mobile Communications and Learning Technologies, 2006. ICN/ICONS/MCL 2006. International Conference on, pp. 203 – 203 (2006)
3. Belokosztolszki, A., Eyers, D.M., Pietzuch, P.R., Bacon, J., Moody, K.: Role-based access control for publish/subscribe middleware architectures. In: DEBS '03: Proceedings of the 2nd international workshop on Distributed event-based systems, pp. 1–8. ACM, New York, NY, USA (2003)
4. Blumenthal, M.S., Clark, D.D.: Rethinking the design of the internet: the end-to-end arguments vs. the brave new world. ACM Trans. Internet Technol. **1**(1), 70–109 (2001)
5. Caesar, M., Condie, T., Kannan, J., Lakshminarayanan, K., Stoica, I.: Rofl: routing on flat labels. In: SIGCOMM '06: Proceedings of the 2006 conference on Applications, technologies,

[2] http://www.fp7-pursuit.eu/

architectures, and protocols for computer communications, pp. 363–374. ACM, New York, NY, USA (2006)
6. Castro, M., Druschel, P., Kermarrec, A.M., Rowstron, A.: Scribe: a large-scale and decentralized application-level multicast infrastructure. Selected Areas in Communications, IEEE Journal on **20**(8), 1489 – 1499 (2002)
7. CCNx: Web site (2010). `http://www.ccnx.org`
8. Eugster, P.T., Felber, P.A., Guerraoui, R., Kermarrec, A.M.: The many faces of publish/subscribe. ACM Comput. Surv. **35**(2), 114–131 (2003)
9. Fiege, L., Zeidler, A., Buchmann, A., Kilian-Kehr, R., M
”uhl, G.: Security aspects in publish/subscribe systems. In: In Proc. of Third Intl. Workshop on Distributed Event-based Systems (DEBS04 (2004)
10. Fotiou, N., Marias, G., Polyzos, G.: Fighting Spam in Publish/Subscribe Networks Using Information Ranking. In: Proceedings of the 6th Euro-NF Conference on Next Generation Internet Networks (NGI). Paris, France (2010)
11. Fotiou, N., Marias, G., Polyzos, G.: Information Ranking in Content-Centric Networks. In: Proceedings of the Future Network and MobileSummit 2010. Florence, Italy (2010)
12. Fotiou, N., Polyzos, G., Trossen, D.: Illustrating a Publish-Subscribe Internet Architecture. In: In Proceedings of the 2nd Euro-NF Workshop on Future Internet Architectures. Santander, Spain (2009)
13. Jacobson, V., Smetters, D.K., Thornton, J.D., Plass, M.F., Briggs, N.H., Braynard, R.L.: Networking named content. In: CoNEXT ’09: Proceedings of the 5th international conference on Emerging networking experiments and technologies, pp. 1–12. ACM, New York, NY, USA (2009)
14. Jokela, P., Zahemszky, A., Esteve Rothenberg, C., Arianfar, S., Nikander, P.: Lipsin: line speed publish/subscribe inter-networking. In: SIGCOMM ’09: Proceedings of the ACM SIGCOMM 2009 conference on Data communication, pp. 195–206. ACM, New York, NY, USA (2009)
15. Katsaros, K., Fotiou, N., Polyzos, G., Xylomenos, G.: Overlay Multicast Assisted Mobility for Future Publish/Subscribe Networks. In: Proceedings of the ICT Mobile Summit. Santander, Spain (2009)
16. Katsaros, K., Xylomenos, G., Polyzos, G.C.: A hybrid overlay multicast and caching scheme for information-centric networking. In: Proceedings of the 13th IEEE Global Internet Symposium. San Diego, CA, USA (2010)
17. Kjallman, J.: Attachment to a Native Publish/Subscribe Network. In: ICC Workshop on the Network of the Future. Dresden, Germany (2009)
18. Koponen, T., Chawla, M., Chun, B.G., Ermolinskiy, A., Kim, K.H., Shenker, S., Stoica, I.: A data-oriented (and beyond) network architecture. In: SIGCOMM ’07: Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications, pp. 181–192. ACM, New York, NY, USA (2007)
19. Lagutin, D.: Redesigning internet-the packet level authentication architecture. Licentiates Thesis in Computer Science, Helsinki University of Technology, Espoo, Finland (2008)
20. Lagutin, D., Tarkoma, S.: Public Key Signatures and Lightweight Security Solutions in a Wireless Environment. Smart Spaces and Next Generation Wired/Wireless Networking **5764**, 253–265 (2009)
21. Lagutin, D., Visala, K., Zahemszky, A., Burbridge, T., Marias, G.: Roles and Security in a Publish/Subscribe Network Architecture. In: Proceedings of the 2010 IEEE Symposium on Computers and Communications (2010)
22. Nikander, P., Arkko, J., Ohlman, B.: Host identity indirection infrastructure (Hi3). In: Proc Second Swedish National Computer Networking Workshop SNCNW. Karlstad, Sweden (2004)
23. Nikander, P., Marias, G.: Towards Understanding Pure Publish/Subscribe Cryptographic Protocols. In: Sixteenth International Workshop on Security Protocols. Cambridge, England (2008)
24. Pallickara, S., Pierce, M., Gadgil, H., Fox, G., Yan, Y., Huang, Y.: A framework for secure end-to-end delivery of messages in publish/subscribe systems. pp. 215 –222 (2006)

25. Perrig, A., Canetti, R., Tygar, J., Song, D.: The TESLA broadcast authentication protocol. RSA CryptoBytes **5**(2), 2–13 (2002)
26. Pesonen, L.I.W., Eyers, D.M., Bacon, J.: Encryption-enforced access control in dynamic multi-domain publish/subscribe networks. In: DEBS '07: Proceedings of the 2007 inaugural international conference on Distributed event-based systems, pp. 104–115. ACM, New York, NY, USA (2007)
27. Smetters, D., Jacobson, V.: Securing Network Content. Tech. rep., PARC (2009)
28. Srivatsa, M., Liu, L.: Securing publish-subscribe overlay services with eventguard. In: CCS '05: Proceedings of the 12th ACM conference on Computer and communications security, pp. 289–298. ACM, New York, NY, USA (2005)
29. Stoica, I., Adkins, D., Zhuang, S., Shenker, S., Surana, S.: Internet indirection infrastructure. IEEE/ACM Trans. Netw. **12**(2), 205–218 (2004)
30. Tarkoma, S., ed.: PSIRP deliverable 2.3, architecture definition, component descriptions, and requirements (d2.3) (2010). `http://www.psirp.org/`
31. Wang, C., Carzaniga, A., Evans, D., Wolf, A.: Security issues and requirements for Internet-scale publish-subscribe systems. In: Proceedings of the 35th Annual Hawaii International Conference on System Sciences, pp. 3940–3947 (2002)
32. Wright, J., Stepney, S., Clark, J., Jacob, J.: Formalizing anonymity: A review. REPORT-UNIVERSITY OF YORK DEPARTMENT OF COMPUTER SCIENCE YCS **389** (2005)
33. Wun, A., Cheung, A., Jacobsen, H.A.: A taxonomy for denial of service attacks in content-based publish/subscribe systems. In: DEBS '07: Proceedings of the 2007 inaugural international conference on Distributed event-based systems, pp. 116–127. ACM, New York, NY, USA (2007)