Towards a Secure Rendezvous Network for Future Publish/Subscribe Architectures

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

Athens University of Economics and Business, Athens, Greece,
{fotiou,marias,polyzos}@aueb.gr,
WWW home page: http://mm.aueb.gr

Abstract. Publish/Subscribe is often regarded as a promising paradigm for Future Internet architectures. Its information oriented nature and its particular security features have stimulated current research efforts which aim at applying publish/subscribe principles to a clean-slate Internet architecture. One of the core components of publish/subscribe architectures is the rendezvous network. Any security failure that a rendezvous network may face will probably jeopardize the operation of the whole (inter-)network. In this paper we highlight security requirements and potential security issues of rendezvous networks and we present security solutions that can be applied in order to shield them.

Keywords: Future Internet, Publish/Subscribe, Security

1 Introduction

Publish/subscribe is regarded as a promising paradigm for the development of future Internet applications and even as a candidate for a clean slate future Internet architecture. Publish/subscribe is currently under investigation in a variety of research efforts-such as the PSIRP [1] and CCNx [2] projects. Its information centrism manifests a significant shift from the traditional endpoint-centric Internet paradigm. In addition, the interest-based decision on accepting data or actively looking for information at the subscriber (recipient) side, shifts the power towards the receiver and thus restoring some balance, compared to the commonly used send-receive paradigm that empowers the sender.

Publish/subscribe architectures are built with three basic components: publishers, subscribers, and an event notification infrastructure also known as the *rendezvous network* [3]. Publishers are information providers who advertise their information by virtue of publications. Subscribers on the other hand are information consumers, who express their interest on specific pieces of information by issuing subscriptions. The rendezvous network is responsible for matching subscriptions with publications and for the initiation of the data transfer from the information publisher to the information subscriber(s) when both are present and ready. Publication and subscription operations are decoupled in time and space, and the fact that the publication-subscription matching takes place at an independent point in the network allows for efficient deployment of multicast,

e.g. [4], mobility, as described in [5], as well as multihoming and indirection, e.g. [6] [7].

The rendezvous network is probably the most critical part of a publish/subscribe architecture. Apart from matching subscriptions with publications, the rendezvous network is the place where access control, policy implementation, and publication scoping can take place [8]. The rendezvous network may be composed of separate, dedicated nodes, such as in [1], or it may be part of the core network, as it happens in [2]. In the former case the rendezvous nodes are usually organized in one-dimensional DHTs–such as Chord [9] and Pastry [10]–when the subscription-publication matching takes place using keys, or in multi-dimensional DHTs–such as CAN [11]–when it comes to semantic-based subscription-publication matching.

Being the core component of publish/subscribe architectures, the rendezvous network is expected to be the target of various security attacks. The purpose of this paper is to describe possible threats against the rendezvous infrastructure by taking into consideration previous experience with DNS and DHTs, to identify the role of the rendezvous network in the already defined security weaknesses of publish/subscribe architectures, as well as to present a set of security solutions that can be applied in securing the rendezvous network. The remainder of this paper is organized as follows. Section 2 presents the rendezvous network's security requirements, while Section 3 outlines some potential security threats. Section 4 overviews security solutions that can be used in order to secure a rendezvous network and finally Section 5 presents our conclusions with discussion and ideas for future work.

2 Security Requirements

General security requirements for publish/subscribe architectures have been defined by Wang et al. [12] as well as by Lagutin et al. [13]. The role of the rendezvous network in order to achieve these requirements is of critical importance.

Publish/subscribe architectures are required to provide publication confidentiality, i.e., publications should not be revealed to unauthorized subscribers, as well as subscription privacy, i.e., users' subscriptions should be kept secret. Rendezvous networks have to provide *access control* mechanisms in order to achieve publication confidentiality, as well as encryption and trust mechanisms in order to provide subscription privacy.

Integrity is also a security requirement for publish/subscribe architectures. Integrity can be classified into three categories: Information integrity, which concerns the exchanged messages, subscription integrity, which concerns users' subscriptions, and finally service integrity, which concerns the protection of the service from malicious faults, such as from malicious nodes inserting bogus messages. The rendezvous network should provide supporting mechanisms that will enable integrity in publish/subscribe systems. Such mechanisms include encryption, digital signatures, as well as access control.

Publish/Subscribe architectures are also required to have high availability. Being the most critical component of a publish/subscribe architecture, the rendezvous network should be fault tolerant and it should also support fair distribution of the load to all rendezvous nodes. Moreover the rendezvous network should be robust and particularly resistant to Denial of Service attacks.

Apart from these traditional security requirements, publish/subscribe architectures should provide authentication, anonymity, accountability, as well as scoping. The rendezvous network may–or may not–participate in the authentication mechanism. Nevertheless it should be able to properly authenticate publishers, subscribers, as well as rendezvous nodes. Moreover, rendezvous nodes should be able to identify themselves to publishers as well as to subscribers. On the other hand, the rendezvous network should support anonymity, accountability, and scoping, as it is responsible for managing user subscriptions, which should be kept secret, and for matching subscriptions with publications, therefore it is the proper place for applying accounting mechanisms and for limiting publication dissemination.

3 Security Threats

Most of rendezvous network functionality resembles to DNS, as it resolves subscription requests to publications and publishers. Nevertheless rendezvous networks are mainly implemented using DHTs–or DHT-like architectures–since DHTs have various good properties, including good load distribution. Security lessons learned from both systems should be considered.

From its inception DNS suffered from some severe problems. Poor implementation, the absence of a trust management architecture and the lack of security mechanisms led to DNS spoofing and DNS cache poisoning attacks [14]. Attackers were able to alter DNS messages or even hijack DNS sessions by taking advantage of the lack of cryptography as well as of the small set and non-randomized transaction identities. Even in 2005, after many years of DNS usage, it was found that reliance on poor transitive trust relations can lead, in many cases, to failures–as they occur in different administrative domains than the attacked DNS server [15]. Rendezvous networks should be designed with trust in mind, moreover as we learned from the DNS paradigm, no matter how well-designed a protocol is, its implementation may suffer from security vulnerabilities, therefore a reliable rendezvous network should be fault tolerant. Finally, the distributed nature of rendezvous networks and the fact that a subscription message may traverse various rendezvous nodes before it reaches its final destination, requires particular care towards resilience to (partial, intermediate) failures.

DHTs allow for efficient key lookup and load balancing among peers. Nevertheless unreliable nodes may lead to a series of problems. Sit et al. [16] classify DHT security attacks into three broad categories:

– Routing attacks: these are attacks in which a malicious node tampers with routing, e.g., it forwards lookups to an incorrect node, it issues wrong routing

updates, or misleads other nodes to join a fake network, leading to network partition.
- Storage and retrieval attacks, in which a malicious node denies the existence of a piece of information, or it refuses to serve particular requests.
- Miscellaneous attacks, such as byzantine nodes, nodes that overload the rest of the network with bogus traffic leading to denial of service attacks as well as unstable nodes.

Moreover DHTs have been found to be susceptible to *Sybil* attacks, i.e., attacks in which a malicious node may spoof multiple identities [17], as well as to Eclipse attacks, i.e., attacks in which multiple malicious nodes cooperate and hide a set of a DHT network from a legitimate node [18]. The security threats that DHTs pose make the role of trust in a rendezvous system even more dominant. Trust mechanisms that promote the cooperation between reliable nodes while isolating the unreliable ones should be used. Moreover the attacks on DHTs show us–among other things–the importance of identity management; in a reliable rendezvous network efficient identification and authentication mechanisms should exist. These mechanisms will guarantee the authenticity of rendezvous nodes (and their functionality).

The current Internet is continuously threatened by Denial of Service (DoS) attacks. Wun et al. were the first that identified and classified this kind of attacks in content based publish/subscribe networks [19]. Interestingly enough, all kinds of attacks identified are based on the rendezvous network; the attackers may flood the rendezvous network with publications, causing the creation of an intolerable number of notifications towards subscribers, or if the rendezvous network allows parameterized subscriptions–e.g., all publications that concern Future Internet published between 2007 and 2009–attackers may use complex expressions and consume the resources of the rendezvous nodes. Finally, attackers may issue a large number of subscriptions, saturating rendezvous nodes memory, since state has to be kept for each subscription. It is observed that DoS attacks against rendezvous networks originate from malicious publishers and subscribers, therefore rendezvous networks should be supported by mechanisms that are able to detect, react, and mitigate them. Nevertheless, if we observe the current Internet it can be seen that users can easily hide their identity, or even have multiple identities. As a result a single mechanism that will simply identify and isolate misbehaving users will probably be ineffective. It will be necessary to identify messages, i.e., publications and subscriptions, that are suspect to lead to a DoS situation and manage them accordingly.

Spam is another security threat that is expected to target publish/subscribe architectures. Although publish/subscribe architectures are considered to be less vulnerable to this kind of attack, as subscribers have to explicitly express their interest on a specific piece of publication before they receive it, Tarkoma has shown that spam is feasible in publish/subscribe architectures [20]. Spam in these architectures can be achieved by learning user preferences. A malicious publisher that is able to predict subscription messages can easily craft publications that will match these subscriptions. The role of the rendezvous network in learning

user habits is very important. Therefore, a reliable rendezvous network should guarantee user anonymity and privacy. Moreover a rendezvous network should have mechanisms that will allow the isolation of malicious publications.

4 Security Solutions

Various security solutions have been proposed for securing publish/subscribe architectures, each one satisfying some of the desired rendezvous network security requirements and offering some protection against the described security threats.

EventGuard [22], is a mechanism that aims at providing security for content-based publish/subscribe systems. Its goal is to provide authentication for publications, confidentiality, privacy, and integrity for publications and subscriptions, as well as to assure availability while keeping in mind performance, scalability, and ease of use. It uses six "guards," that secure six critical publish/subscribe operations (subscribe, advertize, publish, unsubscribe, unadvertize, and routing) as well as a meta-service that generates tokens and keys. All operations involve communication with the meta-service before sending any message. The approach that EventGuard uses for achieving subscription privacy is not so effective, as two–encrypted–subscriptions to the same publication will be the same, and by taking into consideration the number of publications that may exist in a specific rendezvous node, brute forcing subscription messages is a realistic threat. Moreover EventGuard involves many message exchanges with the security meta-service and its re-keying functionality is not well described.

Pallickara et al. [23] achieve message confidentiality and integrity by using *key management centers* (KMC). KMCs are rendezvous nodes enhanced with cryptographic mechanisms. Each of KMC can be responsible for multiple topics. However, each topic can be handled by a single KMC. Publication and subscription operations involve message exchanges with the KMCs, which lead to the creation of a symmetric key that is used for encrypting messages. Although this approach appears to be scalable, the fact that each topic is handled by a single KMC makes the system less tolerant to faults.

The aforementioned security solutions, apart from confidentiality and integrity services, provide the means for identifying nodes in a publish/subscribe architecture. Node identification can be the basis of an anti-spam mechanism–as described in [20]– in which malicious users are blacklisted. User blacklisting may also be a first step towards a DoS resistant rendezvous network.

Miklos was one of the first that discussed access control in publish/subscribe architectures [24]. His approach is based on assigning positive access rights according to a policy list that exists in every rendezvous node. This policy list is used in order to examine whether a user has the credentials to perform an operation. A more sophisticated solution was presented by Belokosztolszki et al. [25] that is based on the OASIS role-based access control system [26]. Attribute based encryption (ABE) [21] is another powerful solution that can be used for access control in a rendezvous network. Badet et al. use ABE to create a social

network in which users can safely publish personal data in their profile [27], which in terms of publish/subscribe is the rendezvous network.

Subscription privacy is–to our knowledge–still an open issue in publish/subscribe architectures. This happens because from the one hand subscriptions should be kept secret in order to assure subscriber privacy, but on the other hand rendezvous nodes should be able to handle the encrypted subscription messages in order to compare them with publications. Moreover due to the, not so big, number of publications that may exist in a rendezvous node, it is probably computationally feasible to decrypt the encrypted subscription messages by brute force. New trends in security such as homomorphic encryption [28] or even schemes that are used in distributed databases for private information retrieval, e.g., the scheme described by Ambainis [29], may be the building blocks for a subscription privacy mechanism.

Denial of Service attacks are, probably, the most difficult to be tackled, as their causes, effects, or even the layer of execution, may vary. Although it seems impossible to create a DoS attack-free rendezvous network, authentication mechanisms, challenge-based and CPU-intensive operations, as well as micropayments, may possibly offer a level of protection for the rendezvous network.

5 Conclusions

The rendezvous network is a critical part of a publish/subscribe architecture as it is responsible for its core functions. As a result, the rendezvous network has an important role regarding the satisfaction of publish/subscribe security requirements. Therefore, at the outset, the rendezvous network seems to be the Achilles heel of a publish/subscribe architecture, since a successful security attack targeting it may jeopardize the whole operation of the system.

In this paper we highlighted the role of the rendezvous network in fulfilling the security requirements of publish/subscribe architectures and we overviewed some security mechanisms that can enable the rendezvous network in achieving this target. Future work in this area includes the incorporation of a selection from all these mechanisms into a single scheme, which will lead to the creation of a secure and robust rendezvous network.

Since most of the proposed solutions are seen basically in the context of the application layer, serious investigations into their feasibility, effectiveness, and efficiency for the core of a clean-slate architecture for a general purpose future internet are required. Such investigations are in pur future research plans in the context of the PSIRP [1] and PURSUIT (EU FP7 funded) projects.

Acknowledgment

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

References

1. PSIRP Project Website: http://www.psirp.org (Last accessed April 2010)
2. CCNx Project Website: http://www.ccnx.org (Last accessed April 2010)
3. Eugster, P.T. and Felber, P.A. and Guerraoui, R. and Kermarrec, A.M.: The many faces of publish/subscribe. ACM Computing Surveys (CSUR), vol. 35, no 2, pp. 131, (2003)
4. Castro, M. and Druschel, P. and Kermarrec, A.M. and Rowstron, A.I.T.: SCRIBE: A large-scale and decentralized application-level multicast infrastructure. IEEE Journal on Selected Areas in communications, vol. 20, no 8, pp. 1489–1499, (2002)
5. Huang, Y. and Garcia-Molina, H.: Publish/subscribe in a mobile environment. Springer Wireless Networks, vol. 10, no. 6, pp, 643–652, (2004)
6. Koponen, T. and Chawla, M. and Chun, B.G. and Ermolinskiy, A. and Kim, K.H. and Shenker, S. and Stoica, I.: A data-oriented (and beyond) network architecture. ACM SIGCOMM Computer Communication Review, vol. 37, no 4, pp. 192, (2007)
7. Stoica, I. and Adkins, D. and Ratnasamy, S. and Shenker, S. and Surana, S. and Zhuang, S.: Internet indirection infrastructure. Springer Peer-to-Peer Systems, pp. 192–202, (2008)
8. Fotiou, N. and Polyzos, G.C. and Trossen, D.: Illustrating a Publish-Subscribe Internet Architecture. Future Internet Architectures: New Trends in Service Architectures (2nd Euro-NF Workshop), (2009)
9. Stoica, I. and Morris, R. and Karger, D. and Kaashoek, M.F. and Balakrishnan, H.: Chord: A scalable peer-to-peer lookup service for internet applications. Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications, pp 160, (2001)
10. Rowstron, A. and Druschel, P.: Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. IFIP/ACM International Conference on Distributed Systems Platforms (Middleware), vol. 11, pp. 329–350, (2001)
11. Ratnasamy, S. and Francis, P. and Handley, M. and Karp, R. and Schenker, S.: A scalable content-addressable network. Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications, pp. 172, (2001)
12. Wang, C. and Carzaniga, A. and Evans, D. and Wolf, AL.: Security issues and requirements for Internet-scale publish-subscribe systems. Proceedings of the 35th Annual Hawaii International Conference on System Sciences, pp. 3940–3947,(2002)
13. Lagutin, D. and Visala, K. and Zahemszky, A. and Burbridge, T and Marias, G.F.: Roles and Security in a Publish/Subscribe Network Architecture. To appear in IEEE Symposium on Computers and Communications, (2010)
14. Lioy, A. and Maino, F. and Marian, M. and Mazzocchi, D.: DNS security. Proceedings of the TERENA Networking Conference, (2000)
15. Ramasubramanian, V. and Sirer, E.G.: Perils of transitive trust in the domain name system. Proceedings of the Internet Measurement Conference (IMC),(2005)
16. Sit, E. and Morris, R.: Security considerations for peer-to-peer distributed hash tables. Peer-to-Peer Systems,Springer,pp.261–269,(2002)
17. Douceur, J.: The Sybil attack. Peer-to-Peer Systems, Springer, pp. 251–260, (2002)
18. Singh, A. and Castro, M. and Druschel, P. and Rowstron, A.: Defending against eclipse attacks on overlay networks. Proceedings of the 11th workshop on ACM SIGOPS European workshop, pp. 21,(2004)

19. Wun, A. and Cheung, A. and Jacobsen, H.A.: A taxonomy for denial of service attacks in content-based publish/subscribe systems. Proceedings of the 2007 inaugural international conference on Distributed event-based systems, ACM, pp. 127 ,(2007)
20. Tarkoma, S.: Preventing Spam in Publish/Subscribe. Proceedings of the 26th IEEE International Conference on Distributed Computing Systems Workshops, pp. 21, (2006)
21. Bethencourt, J. and Sahai, A. and Waters, B.: Ciphertext-policy attribute-based encryption. IEEE Symposium on Security and Privacy, 2007. SP'07, pp. 321–334, (2007)
22. Srivatsa, M. and Liu, L.: Securing publish-subscribe overlay services with eventguard. Proceedings of the 12th ACM conference on Computer and communications security, pp. 298, (2005)
23. Pallickara, S. and Pierce, M. and Gadgil, H. and Fox, G. and Yan, Y. and Huang, Y.: A Framework for Secure End-to-End Delivery of Messages in Publish/Subscribe Systems. Proceedings of the 7th IEEE/ACM International Conference on Grid Computing (GRID 2006), pp. 28-29, (2006)
24. Miklos, Z.: Towards an access control mechanism for wide-area publish/subscribe systems. Proceedings. 22nd International Conference on Distributed Computing Systems Workshops, pp. 516–521 (2002)
25. Belokosztolszki, A. and Eyers, D.M. and Pietzuch, P.R. and Bacon, J. and Moody, K.: Role-based access control for publish/subscribe middleware architectures. Proceedings of the 2nd international workshop on Distributed event-based systems, pp. 8, (2003)
26. Bacon, J. and Moody, K. and Yao, W.: A model of OASIS role-based access control and its support for active security. ACM Transactions on Information and System Security (TISSEC), vol. 5, no 4, pp. 492–540, (2002)
27. Baden, R. and Bender, A. and Spring, N. and Bhattacharjee, B. and Starin, D.: Persona: An online social network with user-defined privacy. ACM SIGCOMM Computer Communication Review, vol.39, no 4, pp 135–146, (2009)
28. Gentry, C.: Fully homomorphic encryption using ideal lattices. Proceedings of the 41st annual ACM symposium on Theory of Computing, pp 169–178, (2009)
29. Ambainis, A.: Upper bound on the communication complexity of private information retrieval. Springer Automata, Languages and Programming, pp. 401–407, 1997