On Using Home Networks and Cloud Computing for a Future Internet of Things

Heiko Niedermayer, Ralph Holz, Marc-Oliver Pahl, and Georg Carle

Technische Universität München, Network Architectures and Services,
Boltzmannstrasse 3, 85748 Garching b. München, Germany
lastname@net.in.tum.de
http://www.net.in.tum.de

Abstract. In this position paper we state four requirements for a Future Internet and sketch our initial concept. The requirements: (1) more comfort, (2) integration of home networks, (3) resources like service clouds in the network, and (4) access anywhere on any machine. Future Internet needs future quality and future comfort. There need to be new possiblities for everyone. Our focus is on higher layers and related to the many overlay proposals. We consider them to run on top of a basic Future Internet core. A new user experience means to include all user devices. Home networks and services should be a fundamental part of the Future Internet. Home networks extend access and allow interaction with the environment. Cloud Computing can provide reliable resources beyond local boundaries. For access anywhere, we also need secure storage for data and profiles in the network, in particular for access with non-personal devices (Internet terminal, ticket machine, ...).

1 Introduction

Many problems of today's Internet are not located within exactly one of its layers, and are not limited to packet forwarding or routing on network layer. While the latter are indeed issues in their own right, many other issues are not primarily a question of lower network layers.

We see research into a Future Internet as tackling a two-fold problem: one on lower layers, and one on higher layers. While the lower layers provide a basic core service, we consider the higher layers to actually bring the Future Internet to the users, with more quality and comfort than in today's Internet.

The higher layer may be called the identifier and services layer. Among its purposes is identifier-to-identifier connectivity. Others might be mobility support, multicast and other services, management, end-to-end security, SPAM and SPIT prevention, user interaction, and consumer empowerment (which is EU goal according to [1]).

This is especially true as in recent years a new form of ubiquitous computing is emerging. In home networks, electronic devices are enabled to communicate over the network. Home networks, in turn, are connected to the Internet with many

T. Zseby, R. Savola, and M. Pistore (Eds.): FIS 2009, LNCS 6152, pp. 70–80, 2010.

new devices now ready to join the Internet. In the AutHoNe project[1] we looked at autonomic functionality for home networks that will make their usage easy enough for home users. Research on Future Networking needs to address this new form of ubiquitious computing. Security concepts like identity, authentication, and trust must be adapted to support this type of new communication.

Our contribution is to propose a different kind of requirement set for the Future Internet that is more centered around the user. Furthermore, we sketch a potential architecture that might meet these requirements. This also includes the to our knowledge new idea to combine Cloud Computing and Peer-to-Peer overlay networks to a new hybrid form of overlay network. This could empower users to build their own more user-centric networks more easily.

In the following sections, we briefly present Related Work in Section 2 and discuss our views on requirements for a Future Internet in Section 3. Our architectural proposal is introduced in Section 4. Then we discuss some parts of our security concept in Section 5 and finally present Cloud Computing as migration strategy in Section 6.

2 Related Work

Future Internet is an extremely broad subject with many different proposals from evolutionary approaches that gradually extend the current Internet to revolutionary approaches that want to completely renew the Internet on all layers. Many Future Internet proposals follow an approach based on overlay networks [3]. SpoVNet[4] is an example for an overlay concept for the Future Internet. Its base is called Ariba and provides self-organizing end-to-end connectivity in heterogenous networks[5] with different connectivity domains. But Future Internet is not only higher layers and overlay. A typical problem in lower layers is the scalability of routing tables e.g. tackled by Hanka et al.[6]. Locator/Identifier split is also a common and consensus among most proposals, see e.g. [7] for a corresponding survey. The Host Identity Protocol[8] is an attempt to standardize Locator/Identifier split for the current Internet at the IETF (RFCs 5201-5207). Proposals for Content-based Networking like PSIRP[9,10] push methods common on application layer like publish-subscribe or data distribution down to lower layers in the network architecture.

Cloud Computing moves the data and computation from machines of users or companies to resources in the network, more precisely virtual machines in data centers of larger IT companies. Users may run their servers in the cloud and need not bother about hardware and other technical issues anymore. Cloud Computing is not only a conceptual idea: already today many companies offer cloud services that run on the computing grid in their respective data centers. Examples are Google's App Engine [11] and Amazon's Elastic Compute Cloud (EC2)[12].

[1] Parts of the presented work is part of the AutHoNe project which is partly funded by German Federal Ministry of Education and Research under grant agreement no. 01BN070[2-5]. The project is being carried out as part of the CELTIC initiative within the EUREKA framework [2].

3 Requirements for a Future Internet

In this section we present a series of high-level requirements that we try to meet and that are different from common listings of Future Internet topics like mobility, efficient routing, or multicast support. Of course, some of these rather technical requirements are also to be met, not as a primary concern, but to the degree necessary to fulfill our goals. We also believe that our requirements cannot be met without appropriate security and privacy mechanisms.

The Future Internet initiatives receive a lot of money from public bodies all over the world. The taxpayer finances the research. We therefore have to justify Future Internet before the public. Our conclusion is that end-users need to see and experience a difference between old and new Internet. Besides connectivity, there should be comfort functions or services provided by the network. This is requirement (1): more comfort.

To achieve an increase in comfort for the end-user, they need support for the networking in their home and for the devices they use. A user and all her devices need to have a common identity, a home where they belong to. Given a trend to small devices, their interconnection has to become more straightforward with little configuration. Real Plug and Play as well as understandable interaction paradigms have to be provided. That brings us to requirement (2): home networks have to be integrated.

The existence of comfort services in the network implies that there are resources in the network that can be used to realize these services. Furthermore, higher-layer Peer-to-Peer protocols as we consider here do not necessarily have resources in the network at their disposal. The common solution to connect end-hosts as Peer-to-Peer overlay does not seem fully satisfying for a Future Internet with also commercial use-cases. Commerce is another issue that is currently either solved on application layer or at the Internet Service provider on link layer. Both issues together on one layer are related to Cloud Computing. We therefore have requirement (3): resources like service clouds in the network.

Another aspect of comfort is to have access anywhere. A Future Internet should help users to be connected with whatever they currently have at their disposal. Simply put, user should be able to take their home with them anywhere they want. Currently, only cellular networks provide a similar mobility. A Future Internet concept should have concepts for integrated roaming in its design. Furthermore, access on third-party machines may be necessary, e.g. like today in Internet cafes or like tomorrow on various vending machines. Security for such cases goes beyond the usage of appropriate protocols or applications. The network should help the user and ensure that as little as possible information is leaked. Naturally, full security cannot be provided when using untrusted machines. To conclude this paragraph, we state our requirement (4): access anywhere on any machine.

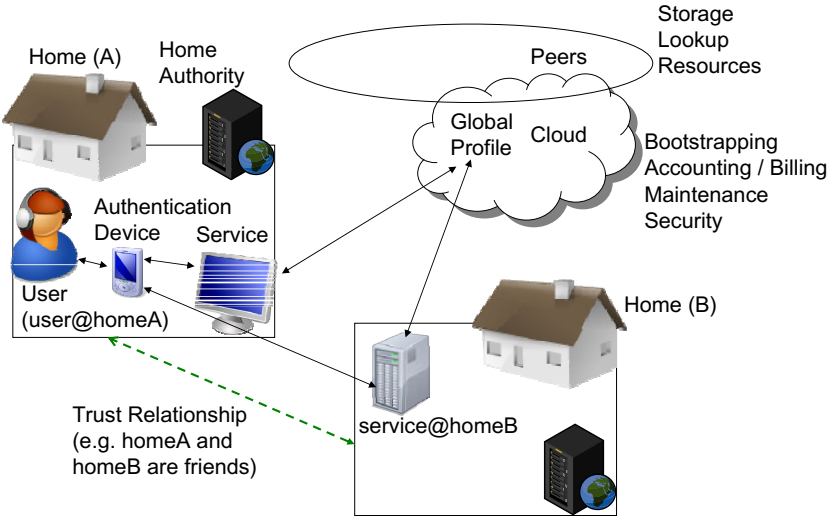

Fig. 1. Overview of components in the architecture

4 An Architecture for an Ubiquitous Future Internet

We center our architecture around the user and base it onto a core Future Internet layer that provides simple locator-to-locator communication. Figure 1 gives a simplified overview of components in the architecture.

Notion of Identity. A *user* has one or more *identities* and certain representations of it. This may be on smartcards or within a PDA or other preferably mobile device. The combination of device and user login onto the device provides two factors for authentication. The knowledge of the password on the device and the private key of the device. For each identity there is a single worldwide profile with basic information about the identity, configuration, and policies. A user has also data which is stored locally or in the network. This is transparent for the application.

Homes. A *home* is a set of devices belonging to an authority, usually represented by a local network and social relations of the users (e.g. families, department). A home is therefore not restricted to a single user. We also think that homes may overlap and contain other homes (e.g. parents, son, daughter). Devices of a home may roam and, thus, be temporarily a part of another home. Each home has a *Home Authority* and devices of a home control their resources in accordance. Self-management of the home includes a distributed knowledge plane that is coordinated by the authority. The knowledge plane provides an abstraction over the heterogeneity of specific devices and allows standardized communication and control for typical networked components as well as home appliance devices and other 'things' with network interface in a Future Internet of Things.

Basic communication within a home is established with zero-configuration protocols. The AutHoNe project[13,14] develops the desired home network concept on top of today's Internet. It shows how such network structures can be established and how a network of homes may look like. In AutHoNe, we consider home networks that consist of home appliance devices as well as devices for communication and computation (e.g. PCs, Smart Phones, and TVs). The network also has a sensor and an actuator domain that interacts with the local physical environment. The abstractions of the AutHoNe project allow a transition to other underlying networks that we might need in a Future Internet.

There is a special relation between *user and home*. A user can be administrator of a home and therefore operate on its behalf. Authentication context and access rights are partially learned from user interaction. Users may brand other entities with certain attributes, e.g. as being part of the home, being a friend, or being a guest. The branding can be done via the authentication device if the entity is already known to the home. A new attribute or friendly name is assigned by the user with the help of software on the authentication device. This will be stored at the Home Authority and the device may receive a certificate if new rights are granted with the operation. To become a new device of a home the branding process involves a device of the owner and the new device. The devices communicate via near-field technology or inside the home's local network. The branding operation is similar to the idea of Zfone[15]. A simple user input with a small code on both devices authorizes the operation and defeats man-in-the-middle attackers and helps to avoid misunderstandings[2] in the local network. Depending on the rank of the user in her corresponding home, the access rights of the other entity will be adapted accordingly (e.g. to the role friend of a user). With these rights, users can use services of the home, access other homes, and monitor and control the home environment.

There is also a special relation between a *user and devices*. Users need to brand their devices to them and to their home. The branding process needs to be a standardized simple interaction of the device, the user, and her device for authentication (smartcard, PDA, ...). A device rates its users according to their credentials. In some cases, users may access the Future Internet on devices foreign to them and their home. Ticket machines and Internet cafes are common scenarios where one might use other devices. When the user uses the device her profile will be downloaded from the network for personalized access and data. The trust into the foreign device determines how much information is exposed. On a friend's machine we will have the basic profile being available. On an unknown machine only a reduced profile will be transfered and a hostile machine will not receive any profile information that it can decode. The assessment and cryptographic operations can be outsourced to the authentication device or her Home Authority and may not necessarily be done by the user.

Cloud and Peer-to-Peer Services. The storage of data and context-dependent access can be realized with a combination of *Cloud and Peer-to-Peer services*. Cloud Computing provides bootstrapping and a security anchor if the authority

[2] Like both entities speak with the wrong other entities.

is not available. An appropriate use of different keys enables this context-aware concept. Untrusted devices only receive the basic cleartext information. The computing cloud realizes global services in the network. In today's Internet DNS root servers serve a similar purpose, yet for a fixed purpose and not being able to be used for other services. Reliability is another argument for the use of cloud resources. For scalability we propose to combine the cloud services with Peer-to-Peer concepts and outsource tasks to peers.

Communication in and between Homes. With respect to *communication* we propose to use zero-configuration protocols for the local network communication. This provides the access to the local authority, the knowledge plane, and local services. Interaction with the Home Authority will give the device a routable locator. Devices in a foreign home may now update their locator in their distant home networks to support their mobility. Foreign devices will not get the same access as home devices. In particular, in many networks they may not get access or only a reduced access unless they belong to a user known to the home or even have been registered as a friend.

To rate devices as guests, friends, etc. we need to establish *trust between homes*. There are some ways to build it. Users may not only brand their devices, but a similar mechanism is used to assign certain attributes like friend to other homes, devices, or users. To some degree we expect that privileges of different roles may be learned from user interaction and feedback. Section 5 provides more details. Additionally, clouds may serve as a trust anchor for other homes and users. In that case, a connection to the cloud can be allowed and the device may access services from there. This is a rather user-centric approach. The resulting 'network of homes' reflects the underlying social structure. Given the trend towards more mobility, we expect that portable devices will connect to foreign networks when they are away from their home network. Following social graphs, trust relations can scale and will be available most of the time for roaming users. The above mentioned extention to access to trusted cloud operators bridges this gap even further.

Access, of course, is useless when one cannot find nodes, users, homes, or services. We therefore need a *lookup* concept. Within a home, we consider naming as a hierarchical scheme with the Home Authority as local root. Between the homes we prefer a flat address-space and in AutHoNe we currently use Pastry[16] for the lookup. For a Future Internet one may use a specialized Peer-to-Peer system instead and integrate the clouds for further improvement.

In our scheme we have *addresses* for users, nodes, homes, and services as well. Semantically, they are their identities and the identity is used for the lookup of their locator. Each entity in the network – physical as well as virtual entity – needs an identity. We currently consider as entities users, devices, homes, and networks. Homes are subnetworks in the identifier space, and networks are subnetworks in the locator space. Locators exist for homes as well as for devices in combination with a home (device@this-home).

Self-certifying identifiers are a solution to reliably authenticate without central authorities. The drawback is that real-world identities can only be learned from contact and not be proven initially. The Peer Domain Protocol (PDP) suite [17] is a protocol suite that can be adapted to learn secure identifiers from interactions and store this information. Section 5 gives more details.

5 Trust Establishment between Homes

As homes form a self-organised network, we believe a PKI with Certification Authorities (CAs) in which all home networks participate to be unlikely: it is implausible that all homes should agree on one CA. Cross-certification with many CAs is also very difficult to achieve [18]. We also assume a very dynamic environment: new devices become part of a home, keys change, users lose their keying material and need to establish it again etc. A more or less static PKI thus seems out of the question.

It is also a stated goal of our approach to empower the user. We thus allow homes to choose themselves with which other homes they build up security contexts. Concerning key exchange and authentication, this is similar to a Web of Trust where users cross-certify each other. However, we have the advantage that our homes have a natural domain structure: in each home, there is one central entity (Home Authority) that can be used as a local trust anchor.

We have previously developed a protocol for the cross-domain authentication of entities [17], PDP, which we can also employ here. It is a four-party protocol where domain servers act as intermediaries and participate in the authentication process. Where two domain servers have a pre-existing security association, e.g. because they have securely exchanged keys, their clients can authenticate to each other securely.

The more interesting property of PDP, however, is that it allows to carry additional information between server and client. Home authorities can store information about previous contacts with other homes and supply this information to the (human) user. The user can then make a better informed decision on how far to trust a certain contact, and how far to trust the authentication (or more precisely, the binding of a home's key to a certain identity that is being claimed). This allows users to gradually build up relationships between each other.

This is best explained by example. Consider two users John and Fred who know each other only from brief contacts in personal life, but now they communicate via the network without physical contact. They establish a first contact between devices of their homes. PDP would return an inconclusive authentication result as no keys have been exchanged between the homes. But during their communication, John and Fred may become more certain that the claimed binding between the other's key and identity is, in fact, correct. They may, for example, use VoIP or instant messenging and speak about something during their last meeting. This may be sufficient for, e.g., John to brand the other entity as 'Fred' and store this in his home authority's storage. He may also give him the role 'Guest' in this way. The Home Authority will store this together with

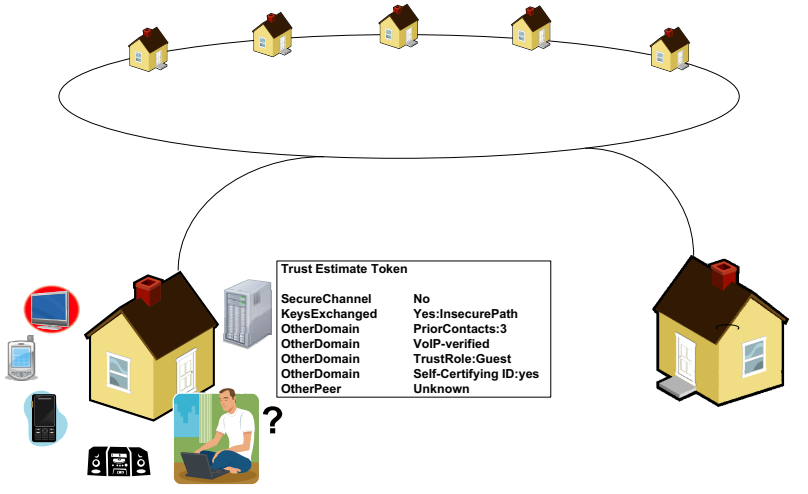

Fig. 2. Trust Estimate Token

the keys of the other server and client. The next time a contact between (different) devices from the two homes occurs, the Home Authorities will display this information in a *Trust Estimate Token* to the user. The devices, however, can be authenticated due to the previous contact with *another* device. The authentication is trust-rated in this case, at 'Guest' level. As John and Fred continue to communicate, they will become more sure of the other's identity over the network and may subsequently raise the trust level. If they, finally, wish to make an entry as 'trusted as a friend' or the like, with according access rights to the home's infrastructure (e.g. allowing to send and store videos or music), they can also exchange further keys during a contact in real life. Figure 2 shows this scenario after a few exchanges.

Note that this approach is also useful if John and Fred do not know each other at all at first. In this case, the process of establishing trust between them would probably be ruled by very prohibitive access rights at first. Also note that the feedback does not necessarily have to be evaluated by a human user. For some applications, e.g. simple instant messenging, policies may allow certain actions based on the current trust level. The idea of AutHoNe and its Knowledge Plane is to learn the level of trust and to adapt the evaluation automatically to the requirements of its applications and services.

6 Cloud Computing as Deployment and Migration Strategy

Cloud Computing is a concept where users access and use network and computing resources instantly in return for money, e.g. provided by Amazon Webservices (AWS). We do not restrict Cloud Computing to the model of today's IaaS

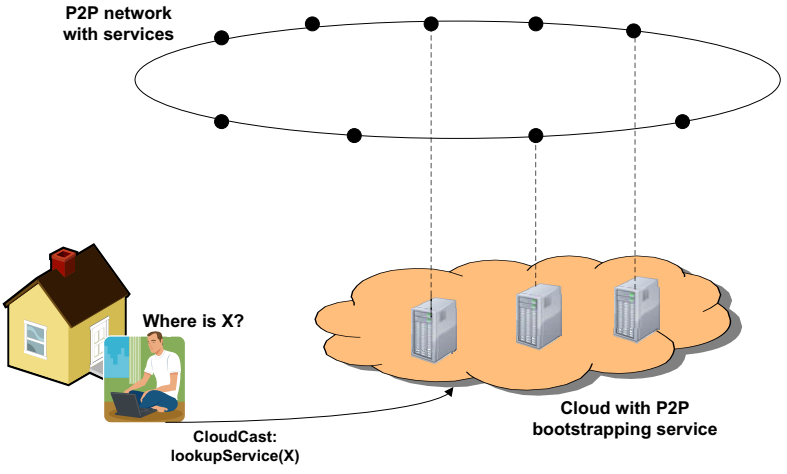

Fig. 3. Cloud Computing can help with the bootstrapping and service provisioning

(Infrastructure as a Service) providers. The major benefit of Cloud Computing in our understanding is that it provides reliable computing and storage resources in the network. These resources provide an additional anchor point for security due to the responsibilities of the infrastructure provider and its accounting. Cloud Computing also provides new means to sponsor a service[3] or to generate revenue from a service[4]. There are also a variety of resources available in home networks that can be useful, and thus this idea of computing clouds can be extended to combine cloud resources with resources in homes (peers). Clouds cooperate with peers and as a consequence will form new kinds of Peer-to-Peer networks.

The resources of a cloud can not only be used by applications, but they can be used by the network and its services. In our vision, users can access both resources in their home as well as resources in the network provided by provider clouds and other peers in the network.

Cloud Computing can also be seen as a deployment and migration strategy for services and higher layer networking approaches. The computing cloud makes it easier to deploy new services on the Internet. It provides an initial set of nodes to bootstrap overlays and their services as shown in Figure 3. An additional benefit is that its nodes reside closer to the Internet core and thus boost network and service performance. This is an advantage over common Peer-to-Peer proposals. It might be possible to introduce anycast messages to a near-by cloud (cloudcast) that can resolve service requests to yet unknown services as well as coordinate access to data stored in the network. For this so-called cloudcast, each home is in contact with either machines in the cloud or peers in the Peer-to-Peer network. These corresponding nodes operate as default gateways and process and forward the requests accordingly. For migration, old and new services may run in parallel.

[3] e.g. sponsor a virtual machine in the cloud.
[4] e.g. some of the payment goes to the software or service provider.

Once the old service is shut down, the cloud may still provide its interface and transcode the messages into messages of the appropriate new service or even a new Internet.

Providing centralized network resources with a combination of Cloud Computing and Peer-to-Peer networks is similar to but more flexible than today's fixed relations. DNS root servers are one example for fixed resources in the network. A cloud can provide them for yet undefined services and allows an adaptation to the ever changing use of the Internet[5].

7 Conclusions

In this article, we have extended the focus of Future Internet from the core to the peripherals. This is in consensus with many overlay and virtualization approaches for Future Internet. We centered our proposal around the user. Home networks are the environment in which the user acts and cloud computing provides necessary resources for tasks beyond the scope of local networks. While we have sketched potential solutions for some aspects, there is no complete architecture yet. Many open questions are related to this vision which we believe future research has to tackle.

References

1. Lemke, M.: The EU Future Internet Research and Experimentation (FIRE) Activities. In: 8th Würzburg Workshop on IP (EuroView 2008) (July 2008)
2. AutHoNe-DE Consortium: AutHoNe-DE Project - Home Page (2009), `http://www.authone.de`
3. Cheng, L., Galis, A., Mathieu, B., Jean, K., Ocampo, R., Mamatas, L., Rubio-Loyola, J., Serrat, J., Berl, A., Meer, H., Davy, S., Movahedi, Z., Lefevre, L.: Self-organising management overlays for future internet services. In: van der Meer, S., Burgess, M., Denazis, S. (eds.) MACE 2008. LNCS, vol. 5276, pp. 74–89. Springer, Heidelberg (2008)
4. Bless, R., Hübsch, C., Mies, S., Waldhorst, O.: The Underlay Abstraction in the Spontaneous Virtual Networks (SpoVNet) Architecture. In: Proc. of 4th EuroNGI Conf. on Next Generation Internet Networks, NGI 2008 (2008)
5. Hübsch, C., Mayer, C.P., Mies, S., Bless, R., Waldhorst, O.P., Zitterbart, M.: Reconnecting the internet with ariba: Self-organizing provisioning of end-to-end connectivity in heterogeneous networks. In: SIGCOMM 2009, Demos (2009)
6. Hanka, O., Spleiss, C., Kunzmann, G., Eberspächer, J.: A DHT-inspired clean-slate approach for the Next Generation Internet. In: Fachgespräche Future Internet in Karlsruhe, ch. 2 (November 2008)
7. Menth, M., Hartmann, M., Klein, D., Tran-Gia, P.: Future internet routing: Motivation and design issues. Oldenbourg Wissenschaftsverlag it - Information Technology 50(6) (December 2008)

[5] Compare Email, Web, Peer-to-Peer, and now Youtube as primary sources of traffic over the lifetime of the Internet.

8. Jokela, P., Nikander, P., Melen, J., Ylitalo, J., Wall, J.: Host identity protocol (extended abstract). In: Wireless World Research Forum,
9. Trossen, D. (ed.), et al.: Conceptual Architecture of PSIRP Including Subcomponent Descriptions (D2.2) (June 2008), http://www.psirp.org/publications
10. Zahemsky, A., Esteve, C., Csaszar, A., Nikander, P.: Exploring the pubsub routing & forwarding space. In: ICC Workshop on the Network of The Future
11. Google Inc.: Google App Engine, http://code.google.com/intl/en/appengine/
12. Amazon Inc.: Amazon Web Services, http://aws.amazon.com/
13. Carle, G., Kinkelin, H., Müller, A., Niedermayer, H., Pahl, M.O., König, A., Luckenbach, T., Scholl, K., Schuster, M., Thiem, L., Petrak, L., Steinmetz, M., Niedermeier, C., Reichmann, J.: Autonomic Home Networks in the BMBF project AutHoNe. In: 8th Würzburg Workshop on IP EuroView 2008 (July 2008)
14. Luckenbach, T., Schuster, M., Pahl, M.O.: An autonomic home networking infrastructure. ERCIM News 77 - Special theme: Future Internet Technology, 41 (April 2009)
15. The Zfone Project (2008), http://zfoneproject.com
16. Rowstron, A., Druschel, P.: Pastry: Scalable, distributed object location and routing for large-scale Peer-to-Peer systems. In: Guerraoui, R. (ed.) Middleware 2001. LNCS, vol. 2218, pp. 329–350. Springer, Heidelberg (2001)
17. Holz, R., Niedermayer, H., Hauck, P., Carle, G.: Trust-rated authentication for domain-structured distributed systems. In: Mjølsnes, S.F., Mauw, S., Katsikas, S.K. (eds.) EuroPKI 2008. LNCS, vol. 5057, pp. 74–88. Springer, Heidelberg (2008)
18. Gutmann, P.: PKI: It's not dead, just resting. IEEE Computer 35(8), 41–49 (2002)