Federated Autonomic Management of HAN Services

Rob Brennan, Zohar Etzioni, John Keeney, Kevin Feeney, Declan O'Sullivan
FAME & KDEG, School of Computer Science and Statistics
Trinity College Dublin
Ireland
{rob.brennan, etzioniz, john.keeny, kevin.feeney, declan.osullivan}@cs.tcd.ie

William Fitzgerald, Simon Foley
FAME & Department of Computer Science,
University College Cork,
Cork,
Ireland
{wfitzgerald, s.foley}@cs.ucc.ie

Abstract— **Management of a heterogeneous "outer edge" network ---where a Home Area Network (HAN) inter-operates with external network and service providers--- is complex and error prone. Configuration of a HAN is typically performed by non-technical end-users. As a consequence, an effective HAN configuration may be hampered by a poor understanding and/or management of HAN service requirements. Mis-configuration, may result in the failure to adequately provide HAN services. Thus, the challenge becomes one of autonomically deploying and maintaining meaningful and error-free heterogeneous HAN configurations. This paper explores an integrated solution to address the following requirements: managed capability sharing, usability and security. A prototype HAN gateway architecture that builds upon an explicit user-centric semantics, enables autonomic management of shared UPnP services with appropriate access controls is outlined.**

Keywords- HAN, Semantics, Federation, Access Control

I. INTRODUCTION

An emerging characteristic of communications networks is the growing complexity and heterogeneity of the "outer edge" domain – the point of attachment of Home Area Networks (HANs) and other restricted private networks to commercial access networks [Chow08]. Management of the components of these networks, outer edge devices, such as femto base stations, home gateways, and so on, is today provided on a piecemeal basis, with different devices having a wide range of management functionality, from none to proprietary or, at best, it conforms to one of a range of competing standards [Bottaro07]. Furthermore, management operations must be performed by end-users, so there is huge potential for mis-configuration that can impact significantly upon the delivery of services on an end-to-end basis with consequent operational and support costs for service providers [ieeeNetwork].

This paper proposes new approaches for services deployment and management in a HAN context. Our primary research focus is on exploring methods through which HAN-based network management systems can assume autonomically, the responsibility for appropriate and secure configuration of HAN devices as new services are deployed and delivered to end-users. A significant challenge in this regard is that as the diversity and capabilities of HAN devices increases, it becomes increasingly difficult to capture and exchange the knowledge required to facilitate delegation of management technology between management domains. A second focus is on end-user enablement through the application of autonomic techniques and semantic policy-based control in the HAN. This includes compliance-driven network access control configuration. These goals give rise to the following research questions:

1. What methodologies and techniques are appropriate to capture semantic models and mappings that will enable the exchange of management capabilities?

2. How can the flow of authority and knowledge in network management systems shape appropriate and secure configuration in HAN environments?

To address these questions an end-to-end service for sharing UPnP capabilities between federated HANs [LasVegasPaper] is developed. This service builds upon previous research by utilizing the Federal Relationship Manager (FRM) [computerNetworks Journal]. In this paper we discuss how that work has been extended with semantic descriptions of shared capabilities and autonomic access controls that build on both these semantic service descriptions and a security knowledge base. We present a revised gateway architecture based on [ieeeNetwork] that links these components together.

This paper is organised as follows. Section 2 discusses system requirements, in section 3 we provide an overview of existing solutions for HAN capability sharing, in section 4 we describe our approach. Finally section 5 describes our conclusions and plans for future work.

II. REQUIREMENTS

A. Capability/Management Capability Sharing

We define a capability as an abstraction of one or more useful aspects of one or more resources or services. Capabilities can be local or remote and must be actively shared to grant remote access. Federation is defined as *a "persistent organisational agreement which enables multiple autonomous entities to share capabilities in a controlled way"* [ieeeNetwork].

This research is supported by the Science Foundation Ireland (Grant 08/SRC/I1403) as part of the Federated, Autonomic End to End Communications Services (FAME) Strategic Research Cluster (www.fame.ie).

There are multiple, overlapping reasons for network and service providers to engage in capability sharing (federation) in the context of the HAN environment. For example, they gain access to individual HAN capabilities in order to maximize their ability to deliver end-to-end services to HAN owners (customers). From a HAN user perspective, having the flexibility to deal with multiple providers, may drive down costs and increase business agility. Finally, the flexibility engendered by pervasive social networking and other similar advances in media production democratization on the Web means that people wish to connect directly their digital infrastructure with that of their friends, on a peer to peer (HAN to HAN) basis.

B. Usability (Managability)

The shift in value towards products' ability to be used in concert with the rest of the digital ecosystem means that consumers must be able to manage (or delegate the management of) federated multi-device, multi-user, multi-network deployments once only the remit of traditional operator's network control centers. Thus the key to the success of these networking features will be the ease with which ordinary users can access them.

Empowering non-technical service consumers and managers depends on making complex systems comprehensible and in a manner that makes it much easier for the user to elicit knowledgeable conclusions from the information presented. This is also essential to empower non-technical service consumers or suppliers to make sense of complex tasks, where users need to be able to understand the information that informs the task and be able to abstract and contextualize this possibly unfamiliar information from a viewpoint that makes sense to them [Novak07].

There are several sources of system complexity that must be tamed to create usable federated HANs. Self-configuring autonomic systems are one approach to increasing usability [Salehie07], policy-based management has been shown to be an intuitive governance model for systems [Barrnett04].

C. Security

While HAN services may provide their own security, for example access control, it is considered best practice to rely on multiple layers of security, for example the deployment of firewalls [10]. The Home Area Network Access Control (HANAC), for example a gateway firewall, provides an important point of demarcation between networks of different levels of trust. The challenge is to generate a HANAC configuration that is aligned with the HAN service security requirements, that is, it permits valid service traffic, and, preferably no more and no less.

Management of the HANAC configuration is complex and requires the HAN administrator (home user) to have a deep knowledge of the high-level security requirements of each HAN service and how those requirements may be upheld as low-level HANAC configurations. Effective configuration may be hampered by a poor understanding and/or management of each service's security requirements, which in turn, may unnecessarily expose the HAN to known threats. Typical HAN administrators do not possess the expert knowledge of a security expert who draws upon best practice and standards, in order to synthesize an effective HANAC configuration that is aligned with the high-level security requirements of HAN services.

III. Existing Solutions To HAN Capability Sharing

For the purposes of this discussion we split the prior work into the specific case of UPnP capability sharing and more general federation solutions.

A. UPnP Capability Sharing

Several researchers have proposed mechanisms for extending UPnP across multiple networks. Lee et al. [Lee07] suggest an architecture for content sharing among UPnP devices, based on *HomeConnectors* communicating with remote *HomeConnectors* in other home area networks via a connection manager. A local SSDP manager listens to the local network and relays local SSDP announcements to remote HANs where they are repeated. However, this architecture does not traverse NAT or firewalls and assumes that all UPnP devices have public IP addresses. Chowdhury et al. [Chow08] present a solution for connecting multiple UPnP networks based on a protocol for establishing trust groups of home networks. Once a group of home networks has been established, users can define which devices they wish to share with the group. Remote devices are represented as embedded devices in the home gateway device. This approach requires dynamic modifications to router and firewall configurations to enable sharing, which makes it less portable and resilient. Kang et al [Kang05] present an architecture based on UPnP and OSGi that allows users to consume multimedia services from multimedia servers outside their home network. The home gateway acts as a proxy media server from multimedia providers reachable outside the HAN. However, the approach is specific to multimedia services and is not general to UPnP services. Kim et al [Kim07] suggest using a SIP-UPnP bridge in the home gateway for allowing remote access to UPnP devices. In this solution secure VPN connections are established to support sharing between HANs.

B. Management Federation

Historical approaches to federation or at least interoperability of telecommunications systems have always emphasized interoperability at the bearer and control planes,. Unfortunately, this only gives very limited flexibility when offering services and provides virtually no support for managing the service lifecycle (although billing is always addressed) that is key to efficient leveraging of the network infrastructure. Although they are widely studied, service level agreements (SLAs) are most often a part of the legal framework for interworking as much as a technological issue. A critical feature of most prior attempts to federation has been the assumption or imposition of a single, unified management model of the network and services with the consequent constraint on the supported business models. It is the current authors' contention that rather than imposing shared models, that management model heterogeneity is an axiomatic

property of any realistic ecosystem supporting dynamic federation formation.

Proposing the use of semantic web technology for OAM is not new, see [Lopez09] for a recent survey. However new aspects of our approach are the emphasis on RDF rather than OWL (although see [Feridun10] for a recent "linked data" approach to OAM), the lack of unified, complete knowledge models of the network(s), the central role of a dynamic approach to semantic interoperability and the combination of FRM-style organization-centric policy (rules) with semantics.

IV. Our Approach

In the following sub-sections we describe the three main technical contributions of this paper: our HAN Gateway/Domain controller architecture, our novel XMPP-based approach to UPnP inter-domain capability sharing, semantic capability graphs for management of the shared capabilities and an autonomic access control configuration manager that provides robustness and usability for typical HAN users.

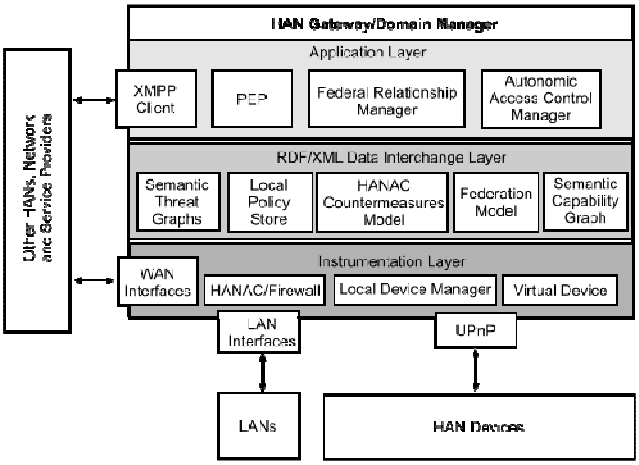

Figure 1. Gateway Architecture

A. Gateway/Domain Controller Architecture

Home gateway devices such as set-top boxes, cable/DSL modems, energy management gateways or networked games consoles are an obvious candidate for situating domain management software within the HAN. Gateways provide natural locations of mediation between the HAN and external actors such as service providers. They often support multiple service plane interactions and thus are more likely to be powered and available on longer time-frames than many HAN devices. They often assume super-peer or controller roles in their associations with other local devices. Finally they tend to be built on more general purpose computing platforms with more extensive computational resources.

Our current prototyping work focuses on a Java-based gateway implementation with UPnP device connectivity. However the gateway architecture presented here (fig. 2) is itself more general. Conceptually the gateway architecture is based on three layers; the application, data interchange and instrumentation layers. The lowest layer is the instrumentation layer that mediates between different device or network technologies and local gateway functions or authorized remote users of local devices and services. The data interchange layer acts as a generalized repository for federation, gateway, HAN, device and application layer management data. The use of RDF/XML for our knowledge and information models simplifies selective re-use and merging of repositories traditionally kept separate in management systems. It has also enabled parallel development of our work at the application layer since most of the knowledge is kept in self-describing ontologies. The application layer hosts a set of management applications and remote service interfaces, e.g. for multimedia capability sharing or federated management functions.

B. XMPP-based UPnP Capability Sharing

In order to enable secure and simple sharing of UPnP devices (and their capabilities) we have extended the UPnP protocol to work over the eXtensible Messaging and Presence Protocol (XMPP) [XMPP] messaging infrastructure. XMPP is an open, XML-based protocol for near real-time messaging and presence. Using XMPP as an infrastructure for connecting multiple networks provides secure and standard communication, simple user roster-management and a powerful presence mechanism, which is useful for the dynamic nature of HANs. The home gateway runs an XMPP client that connects to an XMPP server in order to communicate with the user's defined friends (and shared capabilities/UPnP devices on their networks). Once a friend becomes available (online), capability sharing processes can be initiated.

This service architecture depends on two custom components: a local UPnP network manager and virtual remote devices (proxies for shared devices). The local UPnP network manager acts as a UPnP control point for local devices and acts as an endpoint for remote invocations of UPnP's simple service discovery protocol. This enables permitted remote networks to discover local UPnP devices. Each remote device has a local virtual device instance which is visible on the local UPnP network; this allows devices and managers to act on remote devices as if they were local. All UPnP SOAP requests or responses for remote devices are subjected to local access controls and filtering and then tunneled through XMPP to the remote network. For a full description of this architecture see [LasVegasPaper]. It acts as a flexible local capability definition and sharing infrastructure supporting services such as file sharing, the playing of media streams generated on one network on media renderers (e.g. HD TVs) on another network.

C. Semantic Capability Graphs

Our Federal Relationship Manager (FRM) provides a means for domains to manage capability sharing, e.g. through establishing a shared semantics, secure capability delegation and negotiating the operational rules for sharing, and so on. The overall FRM architecture is described elsewhere [computerNetworksJournal]. A key feature of the system is the distribution of self-describing capability authorities across

federated domains. Here we present for the first time the modeling approach employed to define a semantic capability graph to support shared capability models in federated systems.

1) Capability Authorities

Shared capabilities must map onto some local resources or services in a consistent way and there must be a mechanism to verify which local resources have been shared and to whom. Capability sharing is enabled by delegation of capability authorities between federated domains. A Capability Authority is both a well defined capability and an associated set of permissions and non-functional restrictions.

Any party that wishes to make capabilities available to third parties must construct a capability authority model to express how the capabilities that it is offering are bundled together into capability authorities – basic aggregations of sets of capabilities with the permissions to use them. The capability authority model is instantiated as a service that compares two capability authorities and answers the question as to whether the first capability authority encapsulates the second according to the model. This allows capability authorities that represent arbitrary aggregations of specific permissions to be distributed between federal participants. Whenever a third-party wishes to invoke a capability of a federal partner, the federal partner merely needs to establish whether the capability being invoked is encapsulated by a capability authority that has been issued to that third party. Capability authorities are abstractions that may map to specific resources, services or functions, but they may also map to sets of services with restrictions on parameters. So, for example, a capability authority named *AccessMediaStreamer* may map directly to a service of the same name, or it may map to a set of services (e.g. GetMediaInfo, SetPlayMode,…). By extension, the capability authority *AccessLoungeMediaStreamer* may map to these same services with their parameters restricted to only allow the services to be invoked on a specific device. Capability authority models thus serve to aggregate resources and services into bundles that are useful for distribution and are abstracted away from the underlying implementations.

2) RDF-based Capability Models

Delegation of capability authorities is a flexible and expressive means of applying access control to capability sharing in federal relationships without requiring all of the parties to a priori support common policy or information models. There are already a wide variety of RDF-based formats for describing service invocation, e.g. see [Roman05] . Thus, we adopt an agnostic attitude to semantic service description languages and adopt the simple assumption that the various services that constitute our capabilities may be described by arbitrary sets of RDF triples. This parallels the approach of the Linked Data community [LDtheStorysoFar] to encourage the publication of structured information that is interlinked into a wider web of data to give it context and the opportunity to leverage these other information sources. Note that this does not preclude using any particular formalism such as OWL-S within a particular domain, it just does not make it a perquisite for deploying the system. The objective of this structured capability description is to hold sufficient information to assist human intervention in interworking in the likely case that completely automated approaches to universal interoperability [Noy].

Based upon these assumptions, we construct our hierarchical capability authority models as RDF graphs themselves by adding a set of triples to whatever RDF triples that are available to describe the services that we wish to share as capabilities. We do this, firstly by defining the *HasAuthority* relation as a transitive relation in OWL. Then, we can define OWL or RDF classes to represent whatever collection of services that represent the most convenient aggregations of services for our sharing requirements. For example, we could define the *UserInfoService* Class as a class that enumerates a list of services that provide different types of information about users.

Thus, we can use an OWL reasoner to perform authentication on all requests to local capabilities – if the third party that attempts to invoke any given service has authority for that service then the request is permitted. Hence, we can inject our hierarchical capability authority models into any set of RDF triples simply be adding a small number of triples that represent the particular authority hierarchy that is most convenient for distributing the capabilities that we wish to make available to third parties.

An example RDF Capability Authority expressed in Turtle syntax (omitting standard RDF/RDFS prefixes):

```
@prefix frm: <http://fame.ie/federalrelationshipmanager> .
@prefix dc: <http://purl.org/dc/elements/1.1>/ .
@prefix dbpedia: <http://dbpedia.org/resource/> .
@prefix ex: <http://example.org#> .

ex:MyBigTV frm:hasAuthority ex:MySharedUpnpServices ;
        rdf:type frm:UpnpOverXmppService ;
        dc:creator ex:TheHanOwner ;
        dc:date 2010-01-14;
        rdfs:comment "Capability description for HD TV in front
room" ;
        frm:generatesEvent frm:ConfigurationError ;
        frm:generatesEvent
http://sw.opencyc.org/concept/Mx4rwJN-YpwpEbGdrcN5Y29ycA ;
        frm:serviceName "Display on my big TV" ;
        frm:serviceType dbpedia:Video_Mixing_Renderer ;
        frm:hasInput dbpedia:Streaming_media .
```

D. Autonomic Access Control Configuration

We argue that a framework is required in which one can uniformly represent and reason about the knowledge associated with HANAC security configuration to simplify user involvement in this complex area. We take an ontology engineering approach to modelling this HANAC security configuration knowledge. In [UCC5], [UCC7] the research focused on using ontologies to model network access control configuration for iptables [UCC12] and TCP-Wrapper [UCC23].

An ontology provides a conceptual model of a domain of interest [UCC20]. It does so by providing a formal vocabulary describing various aspects of the domain of interest and provides a rich set of constructs to build a more meaningful level of knowledge. In the case of HANAC

configuration management, an ontology provides the ability to make logical assertions and inferences with which to structure, share and infer new knowledge about the HAN service security requirement and the HANAC configuration domains.

A threat-based approach is proposed as a means of structuring the knowledge about the management of access control configuration. *Semantic Threat Graphs* [UCC8], a variation of the traditional threat tree, are encoded within the ontology-based framework in order to relate knowledge about HAN high level security requirements, best practice recommendations and HANAC access-control rules in terms of assets, threats, vulnerabilities and countermeasures. Threats are organized into a hierarchical structure such as a Microsoft STRIDE based [UCC9] hierarchy. Identifying threats in this way, for example Denial of Service attacks, facilitates the generation of appropriate access-control rules (countermeasures) such as, automatic whitelisting of permitted IPTV service providers (IP addresses) and connection-throttling. The semantic threat graph approach takes advantage of an ontology's ability to share and integrate knowledge within other ontologies. Thus, the iptables and TCP-Wrapper ontologies are reused to describe detailed HANAC countermeasure configurations.

A knowledge-base of best practice standards provide a basis for the generation and analysis of network access control configuration. Ontologies for best practice standards (e.g. [UCC24]) for firewalls, Email servers, Web servers and XMPP servers [UCC11], [UCC16], [UCC17] are developed [UCC4], [UCC6]. Future research will consider additional best practice standards applicable for the HAN environment.

The advantage of taking a ontological approach to representing the semantic threat graph is that it provides a basis for extendability, interoperability and complex composition of other security domains of interest based on the principles of Open World Assumption (OWA) [UCC1]. For example, by including a model of an intrusion detection system within the semantic threat graph, one can more effectively reason about the HANAC recommendations been made based on both a top-down approach (best-practice standards) and a bottom-up approach (IDS rule signatures). In [UCC6], an ontology engineering approach to the management of heterogeneous security access control configuration is considered.

1) Automated HANAC Configuration Synthesis

Synthesis of an appropriate access control configuration relies on the existence of a knowledge-base of candidate HANAC access-control rules that are consistent with the high-level security requirements of each HAN service. These could, for example, represent considered best practice for the HANAC (for example, firewall best practice [UCC24]) that protect HAN services, for example IPTV-based services.

The following is a generic SWRL rule that examines the threats (*?threat*) and vulnerabilities (*?vul*) that each HAN service (*?srvc*) has and searches for suitable countermeasures (*?rule*) that may be implemented by the HANAC gateway (`hanacGW`).

HANService(?srvc) ^ *HANSecService(*`hanacGW`*)* ^
Threat(?threat) ^ *Vulnerability(?vul)* ^
Countermeasure(?rule) ^ *hasWeakness(?srvc, ?vul)*^
threatens(?threat, ?srvc) ^ *exploits(?threat, ?vul)* ^
mitigates(?rule, ?vul) ^ *protects(*`hanacGW`*, ?srvc)*
→ *implements(*`hanacGW`*, ?rule)*

a) Auto-Synthesis of Candidate Access-Control Rules:

As knowledge about assets, threats and vulnerabilities become known, it becomes possible to consider automatic synthesis of HANAC access-control rules as a basis for the previous SWRL rule. The following SWRL rule fragment will automatically populate the knowledge-base with a set of iptables firewall rules (using built-in *swrlx:makeOWLIndividual*), which considers a IPTV service provider HANAC whitelist. Knowledge about an IPTV service's IP address (variable *?iptvip*) and the source IP addresses in which the threat of not providing intended IPTV service provider access (*Threat(?noIPT V Access)*) is used to synthesise specific firewall rules (*?iptr*) from a template iptables rule *(*`iptrtemp`*)*.

HANService(?iptv) ^ *Threat(?noIPTVAccess)*^
Vulnerability(?noIPTVAllowRule)^
TemplateIPTRule(`iptrtemp` *)*^
hasWeakness(?iptv, ?noIPTVAllowRule)^
exploits(?noIPTVAccess, ?noIPTVAllowRule)^
mitigates(`iptrtemp`*, ?noIPTVAllowRule)*^
hasThreatSource(?noIPTVAccess, ?tip)^
hasIPAddress(?iptv, ?iptvip) ^ *hasPort(?iptv, ?iptvp)*
swrlx:makeOWLIndividual(?iptr, `iptrtemp`*,*
?noIPTVAccess, ?iptv)
→ *IPTRule(?iptr)*^
hasChain(?iptr, forward)^
hasSrcIPAddress(?iptr, ?tip)^
hasDstIPAddress(?iptr, ?iptvip)^
hasDstPort(?iptr, ?iptvp)^
hasAction(?iptr, accept)^
mitigates(?iptr, ?noIPTVAllowRule)

V. CONCLUSIONS & FUTURE WORK

The work described here has progressed our understanding of the research questions laid out at the start of this paper as follows:

1. We have investigated the approach of defining new looser, semantic models of HAN capabilities and threat-graph based models for automatically generating access control rules on both those capabilities and any HANAC-defined network resource, e.g. a specific protocol or port. These ideas have been tested by describing the specific capabilities flexibly shared by our UPnP over XMPP system and deploying them for an access control scenario.
2. We have outlined the approach to leverage capability models, threat modes and known best practice guidelines in HANAC configuration to automatically generate policy

rules to best protect a HAN with minimum security expertise required on behalf of the user.

Capability models themselves have previously been shown to provide a flexible and expressive means of applying access control to capability sharing in federal relationships without requiring all of the parties to support a pre-agreed common policy or information models. It is hoped that further experiments with our test-bed will evaluate the extent to which the semantic capability graphs enable semantic interoperability by allowing a range of service definition approaches to be combined and the relative costs of providing policy-based access controls at different levels of granularity of UPnP sharing.

In addition we hope to leverage our semantic HAN models to build dynamic visualizations of HAN activities at multiple levels of abstraction that will aid in HAN behavior comprehensibility for end-users. Future research will also consider self-optimisation and self-healing in conjunction with self-configuration.

REFERENCES

[1] [Chow08] R. Chowdhury, A. Arjona, J. Lindqvist, and A. Ylä-Jääski, "Interconnecting multiple home networks services," International Conference on Telecommunications (ICT 2008), pp. 1-7, June 2008.

[2] [Bottaro07] A. Bottaro, A. Gérodolle, P. Lalanda, "Pervasive Service Composition in the Home Network", 21st International IEEE Conference on Advanced Information Networking and Applications (AINA-07), Niagara Falls, Canada, May 2007.

[3] [ieeeNetwork] Brennan, R., Lewis, D., Keeney, J., Etzioni, Z., Feeney, K., O'Sullivan, D., Lozano, J. A. and Jennings, B.: Policy-based Integration of Multi-Provider Digital Home Services. IEEE Network, Nov/Dec, 2009

[4] [LasVegasPaper] Zohar Etzioni, Kevin Feeney, John Keeney, Declan O'Sullivan, Federated Homes: Secure Sharing of Home Services, to appear in Proc. IEEE Consumer Communications and Networking Conference (CCNC11), 9-11 January 2011, Las Vegas, Nevada, USA

[5] [computerNetworksJournal] Kevin Feeney, Rob Brennan, John Keeney, Hendrik Thomas, Dave Lewis, Aidan Boran, Declan O'Sullivan, Enabling Decentralised Management through Federation, to appear in Elsiver Computer Networks 2010

[6] [Barrnett04] Barrett, R. (2004) "People and Policies: Transforming the Human-Computer Partnership", in proc 5th IEEE int'l workshop on policies for distributed systems and networks (Poilicy'04), 7-9 June 2004, pp 111 – 114

[7] [Novak07] Novak, J.: "Helping Knowledge Cross Boundaries: Using Knowledge Visualization to Support Cross-Community Sensemaking", in Proc. of the Conference on System Sciences, HICSS-40, Hawaii, January 2007

[8] John Wack, Ken Cutler, and Jamie Pole. Guidelines on Firewalls and Firewall Policy: Recommendations of the National Institute of Standards and Technology. NIST-800-41, 2002.

[9] [Salehie07] Salehie, M. and Tahvildari, L. 2009. Self-adaptive software: Landscape and research challenges. ACM Trans. Auton. Adapt. Syst. 4, 2 (May. 2009), 1-42. DOI= http://doi.acm.org/10.1145/1516533.1516538

[10] [Lee07] H. Yong Lee; J. Won Kim; "An Approach for Content Sharing among UPnP Devices in Different Home Networks," IEEE Transactions on Consumer Electronics, vol.53, no.4, 2007.

[11] [Kang05] D. Kang, K. Kang, S. Choi, J. Lee, "UPnP AV architecture multimedia system with a home gateway powered by the OSGi platform", IEEE Transactions on Consumer Electronics, vol. 51, no. 1, 2005

[12] [Kim07] J. Kim, Y. Oh, H. Lee, E. Paik, K. Park, "Implementation of the DLNA Proxy System for Sharing Home Media Contents", IEEE Transactions on Consumer Electronics, Vol. 53, No. 1, 2007

[13] [Lopez09] López de Vergara, J.E., Guerrero, A., Villagrá, V. A., Berrocal, J.: Ontology-Based Network Management: Study Cases and Lessons Learned. J. Network and Systems Management Volume 17, Number 3 / September, 234--254 (2009)

[14] [Feridun10] Feridun, M., Tanner, A.: Using Linked Data for Systems Management. NOMS 2010. Available at: http://ftp.zurich.ibm.com/pdf/csc/FeridunTannerNOMS2010b.pdf

[15] [XMPP] P. Saint-Andre. (2004, October) Extensible messaging and presence protocol (xmpp): Core. IETF. {Online}. Available: http://www.ietf.org/rfc/rfc3920.txt

[16] [Roman05] Roman, D, Keller, U, Lausen, H, de Bruijn, J, Lara, R, Stollberg, M, Polleres, A, Feier, C, Bussler, C and Fensel, D: Web Service Modeling Ontology, Applied Ontology, 1(1): 77 - 106, 2005

[17] [LDtheStorysoFar] C. Bizer, T. Heath, and T. Berners-Lee, "Linked data - the story so far," International Journal on Semantic Web and Information Systems (IJSWIS), 2009.

[18] [Noy] Semantic integration: a survey of ontology-based approaches by Natalya F Noy SIGMOD Rec., Vol. 33, No. 4. (December 2004)

[19] [UCC5] William M. Fitzgerald and Simon N. Foley. Aligning Semantic Web Applications with Network Access Controls. *International Journal on* Computer Standards & Interfaces, Elsevier, Article in Press, October 2009.

[20] [UCC7] William M. Fitzgerald, Simon N. Foley, and Míchel Ó Foghlú. Network Access Control Configuration Management using Semantic Web Techniques. *Journal of Research and Practice in Information Technology, Volume 41 (2)*, May 2009.

[21] [UCC23] Wietse Venema. TCP Wrapper: Network monitoring, access control, and booby traps. *3rd UNIX Security Symposium*, September 1992.

[22] [UCC20] David Taniar and Johanna Wenny Rahayu. Web Semantics Ontology. Idea Publishing, 2006.

[23] [UCC8] Simon N. Foley and William M. Fitzgerald. An Approach to Security Policy Configuration using Semantic Threat Graphs. *23rd Annual IFIP WG 11.3 Working Conference on Data and Applications Security (DBSec), Springer LNCS, Canada*, July 2009

[24] [UCC9] Shawn Hernan, Scott Lambert, Tomasz Ostwald, and Adam Shostack. Uncover Security Design Flaws Using The STRIDE Approach. *http://microsoft.com/*.

[25] [UCC24] John Wack, Ken Cutler, and Jamie Pole. Guidelines on Firewalls and Firewall Policy: Recommendations of the National Institute of Standards and Technology. NIST-800-41, 2002

[26] [UCC11] Jeremie Miller, Peter Saint-Andre, and Philipp Hancke. XEP0220:Server Dialback. *http://xmpp.org*, March 2010.

[27] [UCC16] Peter Saint-Andre. XEP-0205: Best Practice to Discourage Denial of Service Attacks. *http://xmpp.org*, January 2009.

[28] [UCC17] Peter Saint-Andre and Peter Millard. XEP0178: Best Practices for Use of SASL EXTERNAL with Certificates. *http://xmpp.org*, February 2010.

[29] [UCC4] William M. Fitzgerald. An Ontology Engineering Approach to Network Access Control Configuration. *PhD Dissertation, Department of Computer Science, University College Cork, Ireland*, August 2010.

[30] [UCC6] William M. Fitzgerald and Simon N. Foley. Management of Heterogeneous Security Access Control Configuration using an OntologyEngineering Approach. *2nd ACM Workshop on Assurable and Usable Security Configuration, Chicago, USA*, October 2010

[31] [UCC1] Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter Patel-Schneider. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, March 2003.

[32] [UCC12] Netfilter. Netfilter: A framework that enables packet filtering, network address translation and packet mangling. http://www.netfilter.org.