

Policy-Based Integration of Multiprovider Digital Home Services

**Rob Brennan, David Lewis, John Keeney, Zohar Etzioni, Kevin Feeney, and Declan O'Sullivan,
Trinity College Dublin**
Jose A. Lozano, Telefónica I&D
Brendan Jennings, Waterford Institute of Technology

Abstract

The digital home is both the nexus of a new wave of user-centric service integration and the front line of competition between device vendors, connectivity providers, and added-value service providers. Vertical integration of provider offerings via bundled services and devices is no longer sufficient to satisfy customer demands for best of breed elemental services and services composed from traditional service provider offerings, networked consumer devices, Web applications, and social networks. New digital home services spanning these domains must mediate between highly dynamic and overlapping sources of configuration and execution authority, while respecting the desires of the customer and simplifying their interactions with the system. In this article we present a policy-based federated service management architecture that addresses these concerns for digital home devices participating in end-to-end communications services. We also describe an OSGi-based gateway prototype of the architecture and investigate the scalability of our approach via simulation.

In terms of service delivery and network connectivity, the key differentiators of the digital home or home area network (HAN) compared to previous commercial communications infrastructures are the diffusion of ownership, heterogeneity of devices, and a high level of dynamism in the constituents and organization of the network. Service providers can no longer assume exclusive access to networked devices or even full inventories of the network's capabilities. For example, it is not clear how an e-health service for at risk patients could make use of an energy company's environmental monitoring sensors in the customer's home to augment medical monitoring without complex agreements between the providers.

Even though digital home services are an area of active development for many companies and standardization bodies, most of this work is focused on solving the details of equipment interoperability in the HAN, under the assumption that services will be delivered in the traditional way. This is unsurprising as historically there has been very little progress toward inter-service-provider management systems. This is despite the International Telecommunication Union's (ITU's) Telecommunications Management Network (TMN) series of Recommendations, which recognized the need for interdomain management interfaces as early as 1992 [1], and the TeleManagement Forum's ongoing work in this area [2]. Traditional communications service models are largely based on vertical integration of monolithic service provider offerings and equipment with perhaps some highly standardized (and therefore interoperable) customer-owned terminals. This constrained environment naturally simplifies the management of the service life cycle.

However, instead of a relatively static hierarchical system with a limited number of service providers as the only effective sources of authority, HAN management requires a move toward a more fluid ad hoc system consisting of sets of overlapping, complementary, and role- or context-specific negotiated networks of authority (for service execution, deployment, and configuration). It is envisaged that such a system will empower users to make use of composed services based on best of breed solutions, and will confer competitive advantages on agile and customer-focused service providers.

Motivation

The aim of this work is to provide a management architecture to simplify access to end-to-end services or service compositions that leverage the capabilities of the increasingly flexible devices deployed in the digital home. However, this approach should neither burden the end user with the need to understand technical details, or lock them in to one provider's offerings. In our view this necessitates an exchange of context-dependent decision making authority between service providers and HAN owners. In this scenario HAN owners release some management authority over their local network and devices in return for ease of use and access to new services that can maximize the use of HAN and device capabilities, perhaps across multiple service providers. Service and connectivity providers gain access to advanced local capabilities to maximize their ability to deliver end-to-end added value to customers. In return, providers accept the ability of HAN owners to place requirements on their networks in terms of capacity, services, or resources offered. In addition,

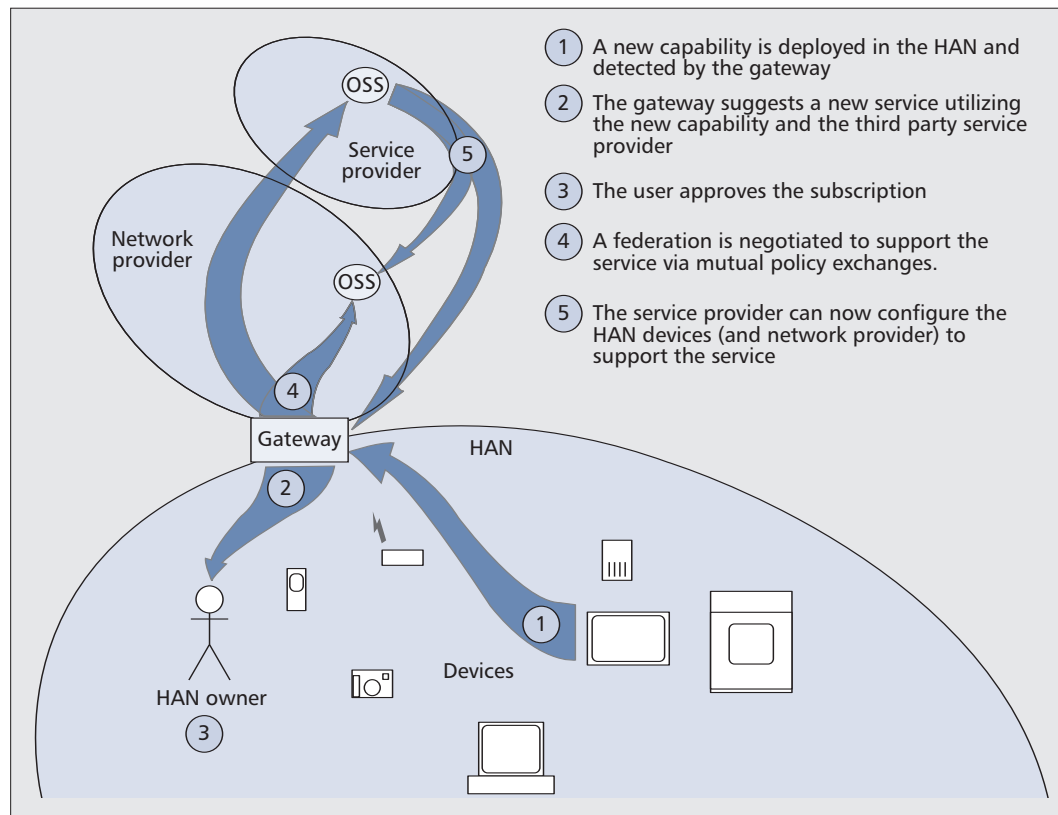


Figure 1. *New HAN capabilities triggering federation.*

service providers can no longer expect exclusive end-to-end access or control of services, networks, or terminals; instead, this access will be potentially shared with multiple service providers, and mediated by the HAN owner's own preferences or policies.

An example use case is illustrated in Fig. 1, where a new media renderer device, a television or monitor, is deployed in the HAN. The gateway, in its role as a device controller, detects this new device. In its role as a service broker application, it identifies new service orchestrations, deployments, or subscriptions that could now be enabled by combining the new HAN capability with previously discovered HAN capabilities, or services from a service catalog. The HAN owner is presented with the option to subscribe to a third party digital media streaming service, and this offer is accepted. To maximize the ability of the streaming service provider to ensure the customer quality of experience, it is desirable to ensure suitable bandwidth allocations during streaming, and to have access to the HAN device's configuration and capability descriptions, thereby requiring that the service provider can negotiate with other network and service providers. In addition, the HAN owner would like to be able to set parental access controls for the streaming service and apply restrictions on the service delivery quality, perhaps based on variable tariffs for the service. This is enabled by the service provider, network provider, and HAN forming a *federation*, whereby management authority rights can be pooled and a common view of the distribution of authority maintained. In this article we define a federation as autonomous entities that have a persistent agreement, which enables them to share capabilities in a controlled way, where the autonomous members also have some independent goals and freedom to pursue their goals. Such federations are persistent agreements (to distinguish from transactions with no evidence of agreement), but the life cycle of such federations can be relatively short.

In the example given, by nature of the federation agree-

ment, whenever a content streaming session is requested, the service provider will be able to ensure end-to-end quality of experience by exercising its authority to reserve bandwidth not only in its own network, but also that of the network provider and the HAN itself. The dynamic federation and subsequent pooling of management authority in this use case can be extended to more complex scenarios. For example, visitors to HAN environments might trigger federation formation between local and remote resources to enable novel multiparty service delivery, or HAN owners might serve as micro-providers as part of a wider public access network. The basic mechanisms of decentralizing management authority and automating federation to support collaboration presented here dramatically increase network flexibility.

Evolution of Home Area Networks

The HAN Landscape

A customer survey in 2008 by a HAN equipment retailer specializing in the technical hobbyist market recorded that an average of seven devices were on their customers' networks. These early adopters provide some indicators for the average home of the future. However, most of the devices were relatively high-capability general-purpose computing equipment, whereas the real growth in HAN population is being driven by new developments in consumer electronics and sensor networks. This second wave of HAN residents will outnumber general-purpose devices by at least an order of magnitude. Given that there will be a large number of relatively resource-constrained devices, it is inevitable that some form of consolidated control architecture will be deployed on more capable devices.

One obvious candidate for service management software hosting within the HAN are home gateway devices such as set-top boxes, cable/digital subscriber line (DSL) modems, energy management gateways, or networked game consoles.

In reality there may be a variety of such gateway devices in a HAN, but current approaches tend to ignore this uncomfortable fact. 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 for longer timeframes than many other 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. All of these properties make gateways suitable locations for the exercise and control of management authority, as both a policy enforcement point and/or a management proxy for more transient, specialized, or limited devices.

There are a number of key technologies for gateway devices in the HAN, and each has a strong impact on the capabilities available for digital home services. Three of the most prevalent (Universal Plug and Play [UPnP], Open Services Gateway Initiative [OSGi], and Web services) of these technologies are discussed below.

UPnP — UPnP is a service-oriented set of interface definitions and protocols for consumer electronic devices, including mobile devices and consumer networking equipment, standardized by the UPnP Forum [3]. UPnP allows devices to dynamically join networks, advertise their capabilities and discover the presence and capabilities of other devices on the network. Each UPnP network is administered by one or more controller nodes. Standard Internet technologies like Domain Name Service (DNS) and Dynamic Host Control Protocol (DHCP) are the foundations of UPnP protocols; interface definitions use Extensible Markup Language (XML); and services are invoked via Simple Object Access Protocol (SOAP). UPnP defines collections of service interfaces grouped into devices, which themselves can be composites, such as a combined media streamer and display unit. There is an especially strong emphasis on support for media devices such as displays and media streamers. A subset of the UPnP specifications focusing on audio and video support have been developed by the Digital Living Network Alliance (<http://www.dlna.org>) together with profiles for WiFi and digital rights management (DRM) technologies to enhance interoperability. The scope of UPnP is firmly within a single administrative domain (the home), to the extent that authentication support is both optional and subject to a variety of implementations. The specifications have enabled Network Address Translation (NAT) traversal for UPnP services, but these have been criticized as extremely vulnerable.

In summary, UPnP provides a widely deployed mechanism for HAN devices to form dynamic local service and capability compositions. However, support for direct orchestration with non-UPnP devices and remote service providers is lacking, and device manageability is very limited.

OSGi — OSGi [4] provides a component model for Java, initially targeted at limited resource devices such as home gateways, but now more generally used in both desktop applications and enterprise application servers. It is both a runtime framework and a service oriented architectural pattern for Java applications. Key features of the framework are that it manages the life cycle of Java-based software components and supports loose coupling of these components through a common service model. In practice, OSGi acts as a dynamic module system that loads modules at runtime, limiting resource consumption to only those essential pieces required for proper operation. This makes it an extremely

efficient way to install, start, stop, update, and uninstall modules on an as-needed basis. The implementation approach is to enhance the Java Virtual Machine (JVM) in these areas by providing infrastructure on top of the JVM rather than changing the JVM itself.

OSGi is a mature set of specifications first published in 1999. The specifications are developed by the OSGi Alliance, an industry consortium. It has enjoyed wide penetration in the embedded systems marketplace with large deployments worldwide. It was also ratified as a final specification through the Java Community Process as JSR 291 in March 2007. OSGi-based component management is reused as part of the JSR 232 Mobile Operational Management specification that defines how mobile devices based on the J2ME Connected Device Configuration can evolve and adapt their capabilities by installing new components on demand.

Thus, OSGi is a mature technology with many attractive features for HAN equipment vendors, especially software life cycle management. However it is still a Java-centric platform with limited direct support for distributed systems or telecommunications-style management capabilities.

Web Services — Web services present the most promising realization of the service oriented architecture (SOA) paradigm [5]. As such they offer autonomous units of functionality available on the network that are loosely coupled; can be described in XML, published, discovered, and composed; and are designed for ease of integration via a set of standard commonly used protocols across platforms, programming languages, and enterprises. Web services can be developed in a variety of programming languages and platforms; yet, as they are exposed by common standard protocols such as XML, SOAP, and HTTP, they can be dynamically discovered over the Web and invoked. The Devices Profile for Web Services (DPWS) is a competing specification to UPnP that is under consideration by the OASIS standardization process. There are already a variety of device vendors offering DPWS interfaces on their devices. It is likely that higher-resource computing nodes will increasingly offer Web services interfaces, and these devices will continue to play an important role in the HAN. However, whether we will see widespread penetration of the DPWS specifications is still an open question.

Multiprovider Service Delivery in HAN

The trend toward device heterogeneity and highly dynamic network structures in the HAN combine to make traditional service delivery models impractical and uncompetitive in terms of operational expense, ability to satisfy consumer demands, and market agility. Now, different parties need to access the HAN infrastructure, and this implies a need for multidomain service management and control mechanisms. Ideally, such a system should allow for coherent end-to-end management approaches across service compositions to enable consumer control at a suitable level of abstraction from network details, limit unforeseen service interactions, maximize agility, and reduce costs for service providers. To maximize scalability and support the new complexity of many-to-many provider/consumer relationships, decentralization of decision making authority and service management control is essential.

Due to the diversity of actors, the heterogeneity of the communications infrastructure, the historical lack of interdomain management standards, and the continuing fragmentation of standardization bodies, it is unlikely that static syntactic approaches to policy interoperability, negotiation, and exchange will suffice. Instead, open semantic models are

needed to facilitate such ad hoc control and management of devices and services by cooperating multiprovider federations of autonomous management systems. It is worth noting that in contrast to the ongoing diversification of communications standardization approaches and bodies, the domain of semantics or knowledge representation is converging on a number of key technologies, such as the World Wide Web Consortium's (W3C's) resource description framework (RDF), that are both well supported by tools and gaining commercial traction in the form of initiatives such as the Linked Data movement.

Policy-Based Service Integration Architecture

Our approach is to provide a federated, business or consumer goal-oriented, semantically enhanced policy-based management architecture. The use of semantically enhanced policies enables devices and service providers to offer machine interpretable descriptions of capabilities and constraints. This will therefore enable automated negotiation and semantic interoperability to handle the potential diversity of capability representations, which may be considerable even for a limited domain such as digital home services. The sections below describe the approach in more detail.

Federated Policy-Based Management

The service integration approach described here overlays a federated policy-based management architecture over a conventional SOA. The key function of this federated management service is to perform management decisions based on the authority granted to the potential user of a service. These management decisions are made on the basis of service invocation and resource usage permissions granted by the federated parties that have authority over the different resources consumed by the service.

To implement the federated policy-based management service, we use an existing policy-based management mechanism, the Community Based Policy Management (CBPM) System [6]. The CBPM model is designed in order to incorporate a naturalistic model of organizations. Here, managed resources, are modeled as hierarchical trees of capabilities. CBPM maintains a distributed map of the hierarchical and federation relationships between individuals, groups, and the resources that each manages. Instead of using centrally defined authority roles (Role-Based Policy Management [RBPM]) [7, 8], organizations are modeled as a hierarchy of authority, where high-level communities (groups) have more high-level, wide ranging, but less specific authorities and responsibilities, and where low-level communities (subgroups or individuals) have more specific and concrete authorities and responsibilities. All members of the organization are members of the root community (but authority to make decisions may be delegated to a smaller subcommunity).

Communities can be progressively subdivided into subcommunities with more fine-grained authorities and responsibilities. Subcommunities can then be delegated some authorities to perform some actions on some subset of the super-community's resources. Individuals can be members of any number of communities, and implicitly remain members of their super-communities. Rather than simply representing atomized roles, communities also contain a collective identity. This collective identity is expressed in the authorities and resources delegated to the community and the community's own collective decision making policies.

In this model the hierarchical community structure is itself a managed resource with capabilities to spawn/merge subcommunities or further delegate authorities. Therefore, authority for community management can be distributed throughout the

federation like any other resource. This means that subcommunities can be allowed to self-manage, with autonomous control over their resources and community structure, including the ability to spawn subcommunities. However, it is important to note that any policies (about resources or communities) specified at a higher level in the community hierarchy have specific precedence over conflicting policies specified by subcommunities.

The CBPM model also supports the concept of a federation, where two or more communities that do not share a common parent can agree to share resources and authorities. Here, each community delegates some resources and authorities to a newly formed orphan federal community. The federal community then has control over the delegated resources, but has no control over the other resources in the participating communities. If allowed, the federal community can then manage and perhaps delegate to subcommunities the resources and authority granted to it. A federated policy decision point function (FPDF) then can support policy decision requests based on its local policy, community, and resource knowledge or appropriately chain requests to other members of the federation for resolution before the final policy decision is returned. Since policy decision responses can themselves contain constraints rather than just simple allow/deny messages it is possible to decompose and recompose these chained decentralized decisions.

A major advantage of the CBPM system is that when a CBPM policy decision point (PDP) receives a policy decision request to perform a specific action on a target resource, the CBPM policy search algorithm utilizes delegations of authorities to eliminate branches of the community hierarchy from the policy search. Hence, it only needs to search for applicable policies within communities that have been delegated the required authority over the relevant resource. In addition, since high-level policies have precedence, the search algorithm can stop once an applicable policy has been found, rather than continuing to search all subcommunities. This can vastly reduce the size of the policy search space. Policy conflicts between subcommunities can also be automatically detected and resolved by elevating any policy request to their nearest common parent.

Overcoming Service Domain Heterogeneity

Given the diversity and dynamism of multiprovider digital home services, devices, or resources, it is unlikely that traditional syntax-based approaches to interoperability will be either sufficiently deployed or flexible enough to ensure interoperability across heterogeneous domains and models. Our approach has been to augment XACML policy specifications and traditional management information models with graph-based metadata using the W3C's RDF specifications. This provides a common interchange format for service, configuration, and monitoring information, but more significant, it provides an open, distributed, extensible, and easily processed framework for querying and making assertions about services and their execution environment. Simple querying of (distributed) RDF graphs through SPARQL provides ample power for current service integration tasks such as end-to-end device selection, composition, and configuration.

A graph-based meta-data approach naturally lends itself to capturing context-specific facades of underlying data stores, for example, by presenting incomplete yet consistent information in compliance with authorization restrictions. Such models can also easily be transformed or merged to suit the data producer's or consumer's requirements. In addition RDF, or *linked data*, currently represents the middle ground between prevalent relational database (RDB) models and newer

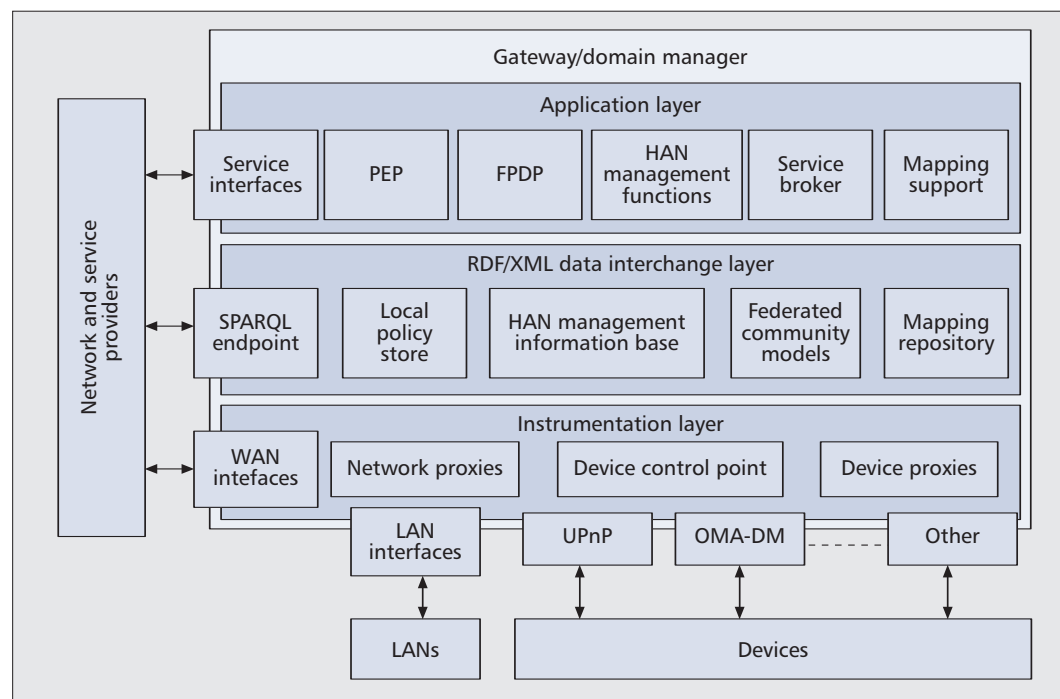


Figure 2. OSGi-based HAN gateway prototype architecture.

knowledge-based representations targeting machine reasoning, such as W3C's Web Ontology Language (OWL), which is itself built on top of RDF.

A key advantage of this approach lies with the generation and usage of partially automated mapping functions to help mediate between different domain models, network resource models, or resource capability descriptions. Even though graph transformation and merging facilitates semantic interoperability across domains, specific mappings will be required to link individual graphs together consistently. Generation of these mappings will require enormous effort, and the current state of the art in this field suggests that complete automation of this process is unlikely to be achieved. However, once generated, it is envisioned that mappings can easily be shared. This, in turn, implies the need for rich metadata describing mappings, their intended application context, standard mapping formats, and so on. Mapping evolution mechanisms will also feature as a way to support the longer-term adaption of services and devices to a changing environment. This requires significant new infrastructure to monitor and manage mapping life cycles and their dependencies on source and target graphs.

OSGi HAN Gateway Prototype Architecture

A key component for realizing federated local autonomy between the HAN and external network and service providers is the gateway device. Our current prototyping work focuses on a gateway implementation based on the OSGi platform and UPnP device connectivity. However, the gateway architecture presented here is itself more general.

Conceptually the gateway architecture, illustrated in Fig. 2, is based on three layers: the application layer, the data interchange layer, and the instrumentation layer. The lowest layer is the instrumentation layer, which mediates between different device or network technologies and local gateway functions (or authorized remote users of local devices and services). It is responsible for natively communicating with devices, collecting service, network, and device information, and acting as a device controller or proxy for individual devices. The instrumentation layer locally communicates with diverse devices

using protocols such as UPnP, OMA-DM, DPWS, and even proprietary or specialized protocols. In addition, it encapsulates gateway WAN interfaces for transport or network layer interaction with remote service or network providers.

The data interchange layer lies above the instrumentation layer. It acts as a generalized repository for federation, gateway, HAN, device, and application layer management data. It abstracts data that comes from the local network and uses RDF to maintain graph-based representations. In general, remote access to this layer is through a SPARQL endpoint controlled via a policy enforcement point (PEP) and mapping support functions in the application layer, which together provide faceted access to the underlying data. Thus, only relevant, permissible, and appropriately transformed data is available to remote users, such as other federation members. This preserves RDF's ability to represent distributed graphs of data without compromising data security or semantic interoperability between federated domains.

The data interchange layer hosts a variety of specialized repositories that expose knowledge and context for local applications running above it or remote applications with appropriate access rights. For example, the HAN management information base maintains a semantically enriched version of a traditional network management function management information base (MIB) for the entire HAN. Thus, there is a composite and unique view of all local network and device capabilities together with federated capabilities. The knowledge model includes the concept of a *semantic connector*, which is a semantic model for devices to describe how they should be handled. This directly supports application layer management functions such as inventory and configuration management, but also plays a role in service monitoring, self-diagnosis, reconfiguration planning, policy conflict identification and resolution, service brokering, and policy authoring. In addition, this layer hosts a set of stores and repositories for policies, federation community information, and mappings. There is a service catalog from which the set of services are proposed to the home user based on the local devices, preferences, and policies.

The application layer hosts a set of management applica-




Figure 3. Scalability analysis of the CBPM model, compared to role-based model.

tions and remote service interfaces (e.g., for devices or federated management functions). The HAN management functions component includes both traditional Fault, Configuration, Accounting, Performance, and Security (FCAPS) management systems for the HAN, but also new knowledge-based capabilities for local autonomous management. One such capability would be security auditing and suggested reconfiguration actions based on machine reasoning over local policies, and threat and configuration models. The PEP is where actions and policies are enforced. The federated PDP (FPDP) is the local instance of the federated CBPM service. It is responsible for both local decision making and invocation of federated policy decision requests. The basic service provided is federated, community/resource/action-based access control; but the ability to embed RDF or OWL descriptions or constraints in policies, and to return such constraints as part of a policy decision, means that local plug-ins can provide enhanced policy conflict detection, resolution, refinement, and decision capabilities. The service broker is an application for matching HAN capabilities to potential service

offerings or orchestrations. Finally, mapping support applications provide both end-user-focused mapping interfaces and federation-oriented mapping discovery, negotiation, and deployment mechanisms.

Scalability Study of CBPM Policy-Based Integration

Policy-based management systems that depend on real-time evaluation of policy rules can suffer from significant performance degradation as policy sets grow. This problem is particularly significant in multidomain management problems as large numbers of policies defined by multiple organizations may have to be evaluated in order to arrive at a decision. The CBPM partitions policies by both capabilities and organizational scope, which combine to allow the system to scale well, even across autonomous domains. In order to demonstrate the relative advantages this policy system exhibits, a series of simulation-based experiments were performed.

These experiments measured the number of policies that must be evaluated and the number of conflicts that can be automatically resolved in a given context. This evaluates the

efficiency of the CBPM system at partitioning a policy search space. The experiments were carried out on a simulated five-layer community model with 14 communities, a resource hierarchy of eight resources arranged in three layers, and a three-layer hierarchy of nine types of actions that could be performed on those resources. For the purpose of the experiment, the policy specification language consisted of simple positive and negative authorizations, including no constraints or conditions. A series of 1000 policy requests were then passed into the CBPM PDP. In order to compare the CBPM model to an alternative policy model, the experiments were repeated using a simulated RBPM PDP, modeling the same experimental setup. For each simulation, using both the CBPM and RBPM PDPs the average number of policy evaluations per request was measured. In addition, the number of modal conflicts was measured, where two applicable policy rules have the same target and action, and opposite results.

The first experiment was conducted by randomly distributing policies throughout the community hierarchy. A second experiment was also performed where the policies were distributed to each community according to the depth of the community in the hierarchy, so most policies were distributed to the leaf communities. This is intended to reflect a situation where there are relatively few policies defined in the higher-level communities, since there are few rules that apply uniformly across the entire organization, and most rules are closely linked with the specific function, resource, or action. Here subcommunities were twice as likely as their parent to contain a policy rule.

Based on the results in Fig. 3, several observations can be drawn. For a given community model and resource model, as policies are added, the CBPM policy search and evaluation algorithms can exploit the hierarchical nature of the community and resource models, and so scale better. This is especially obvious where more of the policies are found toward the leaf communities. Indeed, as the organization grows and the number of managed resources grow, the applicability of a given rule for a given request will be lower, so the efficiency differences between RBPM and CBPM models will become even more pronounced.

Based on the relatively small experimental setup described, with only eight resources and nine actions, the probability of conflicts occurring in 5000 random policy rules grows toward 100 percent. However, the probability of conflicts still increases much more quickly, and with a much smaller number of policies in the system, for the RBPM model than for the CBPM model. This initial work demonstrates a significant potential advantage in deploying the CBPM model.

Summary and Outlook

The paths and modes of digital home service provision will continue to fragment; hence, it is necessary to look toward flexible service architectures based on the pooling of management authority, such as the federated model described here. However, the work presented in this article only deals directly with enabling local autonomy (i.e., independent decision making) as part of a (federated) multiprovider environment. Future work will have to explore autonomic management of communications and services [9, 10] in the home to both empower end users to easily interact with home capabilities and to set boundaries on their behavior in predictable ways.

Acknowledgments

This work was partly funded by Science Foundation Ireland via grant 08/SRC/I1403 — Federated, Autonomic Management of End-to-End Communications Services (FAME).

References

- [1] ITU Rec. M.3010, "Maintenance: Telecommunications Management Network, Principles for a Telecommunications Management Network," 1992.
- [2] The TeleManagement Forum; <http://www.tmfforum.org>
- [3] The Universal Plug and Play Forum; <http://www.upnp.org>
- [4] The OSGi Alliance, "OSGi Service Platform, Core Specification r4," Aug. 2005; <http://www.osgi.org>
- [5] M. P. Papazoglou *et al.*, "Service-Oriented Computing: State of the Art and Research Challenges," *IEEE Computer*, vol. 40, no. 11, Nov. 2007, pp. 38–45.
- [6] K. Feeney, D. Lewis, and V. Wade, "A Service Oriented Policy Architecture for Managing Services Provided by Web-Application Frameworks," to appear, *IEEE Internet Comp.*
- [7] K. Feeney, D. Lewis, and V. Wade, "Roles Considered Harmful in Policy-based Management for Dynamic Organizations," *Proc. 10th IFIP/IEEE Int'l. Symp. Integrated Net. Mgmt.*, Munich, Germany, May 21–25, 2007, pp. 741–44.
- [8] M. Solman and E. Lupu, "Security and Management Policy Specification," *IEEE Network*, vol. 16, no. 2, Mar. 2002, pp.10–19.
- [9] S. Dobson *et al.*, "A Survey of Autonomic Communications" *ACM Trans. Autonomous Adaptive Sys.*, vol. 1, no. 2, Dec. 2006, pp 223–59.
- [10] B. Jennings *et al.*, "Towards Autonomic Management of Communications Networks," *IEEE Commun. Mag.*, vol. 45, no. 10, 2007, pp. 112–21.

Biographies

ROB BRENNAN (rob.brennan@cs.tcd.ie) is a research fellow in the Knowledge and Data Engineering Group (KDEG) within the School of Computer Science and Statistics, Trinity College Dublin (TCD), Ireland. His research interests include semantic mapping, distributed systems, and the application of linked data approaches to network and service management. He has contributed to 3GPP, TMF, IETF, OMG, and ETSI communications standards. In 2004 he obtained his Ph.D. in distributed intelligent networks from the Department of Electronic Engineering at Dublin City University. Prior to joining TCD he worked in the Ericsson Ireland Network Management Research Centre and a number of startups. He currently leads the Home Area Networks team in the SFI-FAME project.

JOHN KEENEY is a research fellow with the KDEG in the School of Computer Science and Statistics at TCD. His research focuses on the use of semantics in the management of autonomic adaptable systems, particularly networking and telecoms systems. He graduated from TCD in 1999 with an undergraduate degree in computer engineering. His Ph.D. in computer science, also from TCD, was completed in 2004. He has published in excess of 30 papers in significant journals, conferences, and workshops.


DAVID LEWIS has 17 years' R&D experience in academia and industry, with over 100 publications. He gained his Ph.D. from University College London, where he worked for 12 years before moving to TCD in 2002. He has an international research reputation in the engineering of open distributed systems for network and service management and the knowledge-driven engineering of autonomic pervasive computing and communication systems. He is on the editorial boards and program committees of numerous conferences and journals in the management systems and autonomic systems areas. He has worked as a senior consultant to a Danish SME developing network management platforms, and has contributed to international standards bodies such as ITU-T Study Group 12 and the TeleManagement Forum in the area of management system modeling.

DECLAN O'SULLIVAN is director of the KDEG at TCD, and has over 20 years' R&D experience in both industry and academia. He holds Ph.D., M.Sc., and B.A. (Mod) degrees in computer science from TCD. His particular research interest is in knowledge driven approaches to achieving semantic interoperability, especially applied to network and service management in distributed networks. During his time in industry, he was involved in industry and fora such as TeleManagement Forum and Object Management Group (OMG). He has over 70 publications, and has contributed to several organizing and program committees in this field.

ZOHAR ETZIONI is a Ph.D. student in the KDEG at TCD. His main research interests are network management, service management, distributed systems, swarm intelligence, software engineering, and agile methodologies. He has 11 years of extensive industry experience as a system and software architect in various domains including telecommunications and medical imaging. He has a Master's degree in computer science from the Open University Israel and a Master's in philosophy from Tel Aviv University, Israel.

KEVIN FEENEY is a research fellow with the KDEG at TCD. He holds Ph.D. and B.A. (Mod) degrees in computer science from TCD. He has experience as a researcher and software developer, designer, and architect in a number of national and international companies since 1997. His research on policy-based management has been widely published in several high-quality international journals and conferences.

BRENDAN JENNINGS [M] is a senior research fellow with the Telecommunications Software & Systems Group, Waterford Institute of Technology, Ireland. His



research interests include autonomic network management, charging and billing, and performance management. He regularly serves on the technical program committees of a number of network management related conferences and workshops, and is the current Secretary of the IEEE ComSoc Technical Committee on Network Operations and Management (CNOM).

JOSÉ A. LOZANO is head of the Autonomic Systems Division at Telefónica I+D. He holds a M.S.Eng. degree in telecommunication from the Universidad Politécnica de Madrid. He has more than 15 years' experience in the network and services management area, where he has been involved mainly in activities related to technological and strategic consultancy. His research interests include distributed management systems, and adaptive and self-managed infrastructure engineering. He is participating in some standardization bodies such as ETSI and the TeleManagement Forum.

