Federated Management of the Future Internet: Status and Challenges

J. Famaey*and F. De Turck

Department of Information Technology, Ghent University – IBBT, Ghent, Belgium

SUMMARY

The Internet's original static services have been superseded by rich multimedia services with stringent end-to-end QoS requirements. Additionally, there has been a trend from simple applications offered by a single provider, towards service compositions, managed across the bounds of multiple domains. It is widely accepted that the end-to-end requirements of multimedia and composed services cannot be satisfied by the current Internet, which does not support inter-domain collaboration. The network federation paradigm was advanced to address these limitations. It envisions the automatic negotiation and management of dynamic agreements between network domains, allowing them to collaborate to achieve goals they cannot achieve alone. This article presents an overview of state of the art research in the area of federated network management. Specifically, existing definitions are compared and aligned. Moreover, the most important efforts towards an architecture for a federated Future Internet are discussed. Finally, we have identified several important research challenges that need to be tackled before the federated Future Internet vision can be fully achieved. For each of these challenges, existing research efforts are surveyed and remaining open issues identified. Copyright © 2011 John Wiley & Sons, Ltd.

1. INTRODUCTION

Since its inception, the Internet has evolved from a traditional packet-switched communication network towards a service-oriented delivery platform. Its original static and best-effort services, such as email and the World Wide Web (WWW), have been superseded by rich and complex services with stringent end-to-end requirements. Specifically, multimedia services, such as Internet television (IPTV) and Voice-over-IP (VoIP), have recently grown to become the most prevalent source of traffic on the Internet [1]. This type of services requires strong guarantees on the end-to-end Quality of Service (QoS), in order to satisfy end-user quality requirements. In parallel, the Internet's services are evolving from simple applications, offered by a single provider, towards complex end-to-end service compositions managed by multiple providers across administrative domains.

The trend towards composite services, distributed across multiple networks and with stringent end-to-end requirements, strongly contributes to the need for coordination and collaboration across independent network domains. In the current Internet, coordination across management domains exists only on a very limited scale, consisting of long-term and static collaborations with manually negotiated contracts [2]. Additionally, their scope is limited primarily to the participation in end-to-end routing protocols, network peering arrangements, leasing of computing or storage resources

---

*Correspondence to: Jeroen Famaey, Ghent University – IBBT, Gaston Crommenlaan 8/201, B-9050 Gent, Belgium, jeroen.famaey@intec.ugent.be

Prepared using **nemauth.cls** *[Version: 2010/05/13 v2.00]*

and the exchange of limited management information (e.g., for charging or billing purposes). For example, the relationship between content providers and Content Distribution Networks (CDNs), or service providers and cloud providers can be considered rudimentary federations, with (semi-)statically negotiated agreements. In our vision, such static agreements are inadequate in light of ever-changing end-user needs, network dynamics and evolving service requirements. In order to address this issue, Future Internet research has spawned the *federated network management* paradigm. Federated network management supports the coordination, interaction and collaboration of independently managed network domains through automatically negotiated and managed agreements. It aims to facilitate the delivery of value-added end-to-end services across the Internet. The first applications of federation concepts have recently started appearing. For example, the IETF Cdni working group[†] aims to standardise a set of interfaces that allow CDNs to set up federations among each other. Additionally, one of the goals of the GENI (Global Environment for Networking Innovations) and FIRE (Future Internet Research and Experimentation) research initiatives is to set up federated testbeds, which consist of multiple smaller testbeds with varying capabilities, allowing researchers to perform large-scale Future Internet experiments [3].

Let us consider the delivery of a multimedia service from an over-the-top content provider across the Internet, to show the strength of federations. In the current Internet, such services are delivered on a best-effort basis, which results in a lack of support for QoS guarantees. In the envisioned federated Future Internet, the service provider could negotiate agreements with transit Internet Service Providers (ISPs) on the path towards its customers, in order to set up a QoS guaranteed end-to-end path. Additionally, CDNs or cloud providers could be included in the federation, allowing the service provider to lease storage resources in the Internet's edge nearby its customers, in order to deploy caches and thus reduce delivery costs. This type of multimedia service federation clearly benefits all involved parties. The service provider can improve the quality of its delivered service and will thus enjoy increased customer satisfaction and higher revenues, the transit ISPs, CDNs and cloud providers will get a share of those revenues and the end-user quality of experience will be significantly increased. As the end-user expectations, service requirements and network characteristics may change over time, the negotiation and management of federations obviously needs to be automated and dynamic.

Although the ideas behind federated network management have existed for some time, several open issues and challenges remain to be solved. This article discusses the most important challenges associated with federated management of the Future Internet, surveys the current state of the art in research and identifies the remaining open issues. More specifically, the remainder of this article is structured as follows. Section 2 analyses the plethora of existing definitions of network federations and attempts to merge them into a unified definition encompassing different views on the topic. Several architectures have been proposed to incorporate federations into the design of the Future Internet. An overview of the most important federated Future Internet architectures is given in Section 3. Subsequently, Section 4 presents the important challenges and evaluates the status of current research efforts concerning them. Finally, the article is concluded in Section 5.

2. DEFINITION

Even in the narrow context of communications networks, the term *federation* has been defined in many ways over the years. This section lists some well known definitions from literature and identifies the common characteristics in order to align them. Originally, the term federation stems from political jargon. The Oxford English dictionary defines it as follows:

> *"The formation of a political unity out of a number of separate states, provinces, or colonies, so that each retains the management of its internal affairs"*

[†]http://tools.ietf.org/wg/cdni/

The adoption of the term in the context of communications networks has given rise to a plethora of definitions, derived from the original political definition. Panlab [4], a federated European test-bed facility, defines a federation as follows:

> *"A model for the establishment of a large scale and diverse infrastructure for the communication technologies, services, and applications and can generally be seen as an interconnection of two or more independent administrative domains for the creation of a richer environment and for the increased multilateral benefits of the users of the individual domains".*

Additionally, several definitions have been advanced in the context of network management specifically. Serrano *et al.* [5, 6] define a federation as:

> *"A set of domains that are governed by either a single central authority or a set of distributed collaborating governing authorities in which each domain has a set of limited powers regarding their own local interests"*

Finally, Jennings, Feeney *et al.* [7, 2, 8] have come up with an alternative definition:

> *"A persistent organizational agreement that enables multiple autonomous entities to share capabilities in a controlled way"*

Based on these definitions, we define a federation, in line with our vision, as follows. It is an agreement between a set of independent entities or network domains, that retain the responsibility over their internal management. As explicitly stated by Serrano [6], the federation is either governed by a central authority or by the independent entities themselves in a distributed manner. The agreement pertains to the (possibly restricted) sharing of a set of capabilities between the federation partners. Jennings and Feeney [2, 8] clarified that the term *capability* should be interpreted broadly, and might range from the usage of network infrastructure to the configuration of a specific device or software component. Finally, the federation should be persistent, which means that it should outlive individual interactions and transactions. Note that this does not imply that it should be in any way permanent.

Federation agreements can be *vertical* or *horizontal* [9]. In vertical federations, capabilities of one partner are leveraged by another. There is thus no real collaboration, but rather the provisioning of services or resources by one party in return for some kind of reward. This type of federation is already in use today (in a limited and static way), such as for example the use of CDNs by content providers, or cloud resources by service providers. Celesti *et al.* [9], among others, believe that there will be an evolution towards horizontal federations, in which two or more parties collaborate in order to achieve a set of common goals by the mutual sharing of capabilities. Therefore, the vertical federation can be considered a special case of the horizontal federation. The challenges and open issues discussed throughout the remainder of this article, thus apply to both types equally.

3. ARCHITECTURES AND MODELS

This section describes state of the art Future Internet management architectures and models that incorporate support for federations among independent network domains. Specifically, two important architectural models towards federated management of the Future Internet are explored: the *Layered Federation Model* from the FAME research cluster, as well as the *Autonomic Internet architecture* from the European AutoI project.

3.1. Layered Federation Model

The Layered Federation Model (LFM) [8, 2] was proposed within the context of the Federated, Autonomic End-to-End Communications Services Strategic Research Cluster (FAME)‡, a project
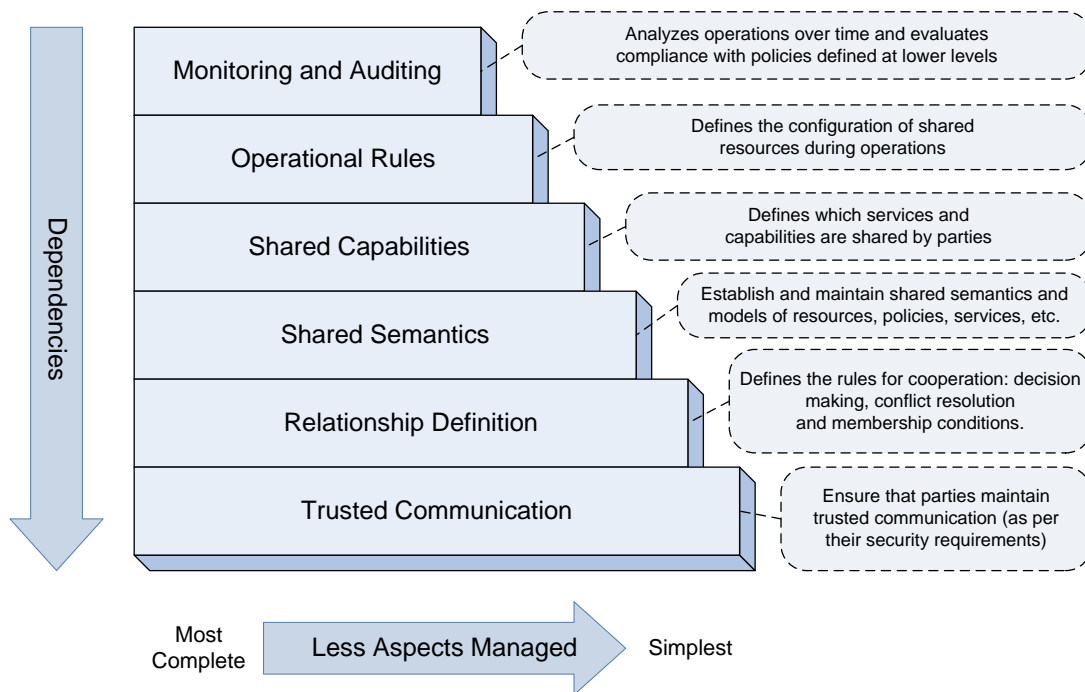
---

‡http://www.fame.ie

Figure 1. Layered Federation Model (from [8])

funded by Science Foundation Ireland. The LFM is a general purpose high-level conceptual model of the components of a federal agreement. It captures and reflects the factors that may vary across federal arrangements and models their evolving, dynamic nature. Figure 1 depicts the model's six layers.

Each layer of the model represents an aspect of a federal agreement. Each layer builds upon the underlying layers and cross-layer interactions may occur. Additionally, in some agreements, there may be empty layers. Specifically, the model is composed of the following layers:

- **Trusted communication layer**: In order to facilitate the communication between independent management domains, a communication channel must be configured that satisfies the security and trust requirements of both parties. They must agree on communication protocols and security mechanisms. This is the most fundamental layer of the model, as all higher level agreements and interactions make use of it.
- **Federal relationship definition layer**: This layer supports the definition and transmission of the basic rules that govern the relationships between federal partners. This provides a generic methodology to negotiate on the rules concerning membership of a federation and sharing of capabilities.
- **Shared semantic layer**: The goal of a federation is to share capabilities among network domains, in order to obtain some added advantage for all participating parties. However, the involved parties will usually have their own mechanisms for addressing and describing their internal capabilities. The goal of this layer is to align these diverging mechanisms and descriptions by providing a mapping between the internally used semantics.
- **Shared capabilities layer**: On top of a sufficiently secure communications channel, a relationship agreement and a semantic mapping to facilitate mutual understanding, the actual sharing of capabilities can be supported. This consists of operations to add or remove capabilities to a shared pool, as well as discovery mechanisms to allow other parties to find the capabilities available for use at any particular time.
- **Operational rule layer**: An extension of the capability sharing layer that allows federation partners to view and configure the capabilities shared by others.

- **Monitoring and Auditing layer**: Although the layers of the LFM are expected to manage their own auditing, reporting and compliance assurance, this layer adds additional facilities. It supports long term monitoring and auditing through aggregation, as this might be required by some federal agreements.

In order to add support for the abstract LFM in actual network management systems, Feeney *et al.* proposed the Federal Relationship Manager (FRM) [8]. It is designed to interconnect existing network management systems, through the implementation of several aspects of the LFM. In order to reduce the complexity and cost to deploy the FRM, it minimizes the necessity for common technologies, protocols, models and processes within the participating network domains. The goal of the FRM is to adopt the set of common technical aspects that an organization must adopt in order to manage and maintain federal relationships. It incorporates two components, a semantic ontology mapping framework and a Community Based Policy Management System (CBPMS). The ontology mapping framework enables the efficient and effective creation and management of mappings between domains in order to increase understanding of shared capabilities across federations. The CBPMS provides secure authority management capabilities that are policy language and information-model neutral.

A further application of the LFM was introduced by Brennan *et al.* [10], in the form of the multi-domain relationship management architecture. The LFM is renamed the Layered Relationship Model (LRM) as it is extended to support hierarchical (i.e., domain compositions) in addition to peer-to-peer (i.e., domain federations) relationships. The multi-domain relationship management architecture is an IT architecture to manage the federation of next-generation communications service providers (CSP). It consists of three components: the domain relationship map (DRM), trusted community-based policy management system (TCBPMS) and relationship traceability map (TM) tool chain. The DRM models the federal relationship from the perspective of an individual participating domain. It provides an instantiation of the operational rules, shared capabilities, shared semantics and relationship definition layers of the LRM. The TCBPMS is an extension of the previously discussed FRM's CBPMS, which also incorporates an instantiation of the LRM's trusted communication layer. The relationship TM tool chain automatically generates TMs, which document the relationship between interacting software components within and across network domains. This architecture thus incorporates the features of the FRM, as well as several novel aspects.

3.2. Autonomic Internet Architecture

The autonomic Internet (AutoI) architecture was proposed within the context of the European FP7 AutoI project [11]. The architecture's aim is to support self-managing virtual resources that can span across heterogeneous networks. Although the project's focus is on autonomic and self-managing next-generation service-aware networks, it incorporates the federation aspect by supporting the management and sharing of resources across the bounds of administrative domains. The AutoI architecture is composed of five layers, the OSKMV planes [12]:

- **Virtualization Plane**: Virtualizes physical resources in order to support on-the-fly migration and reconfiguration of network resources.
- **Management Plane**: Deals with the creation and management of individual autonomic control loops. These loops are realized by Autonomic Management Systems (AMSs), which represent organisational or administrative boundaries.
- **Service Enablers Plane**: Responsible for the discovery, deployment, and composition of services.
- **Knowledge Plane**: A fully distributed information service, responsible for the timely dissemination of information to the other planes. Its inferencing capabilities allow it to derive new knowledge from the gathered information.
- **Orchestration Plane**: Orchestrates the interactions between network domains, AMSs and services.
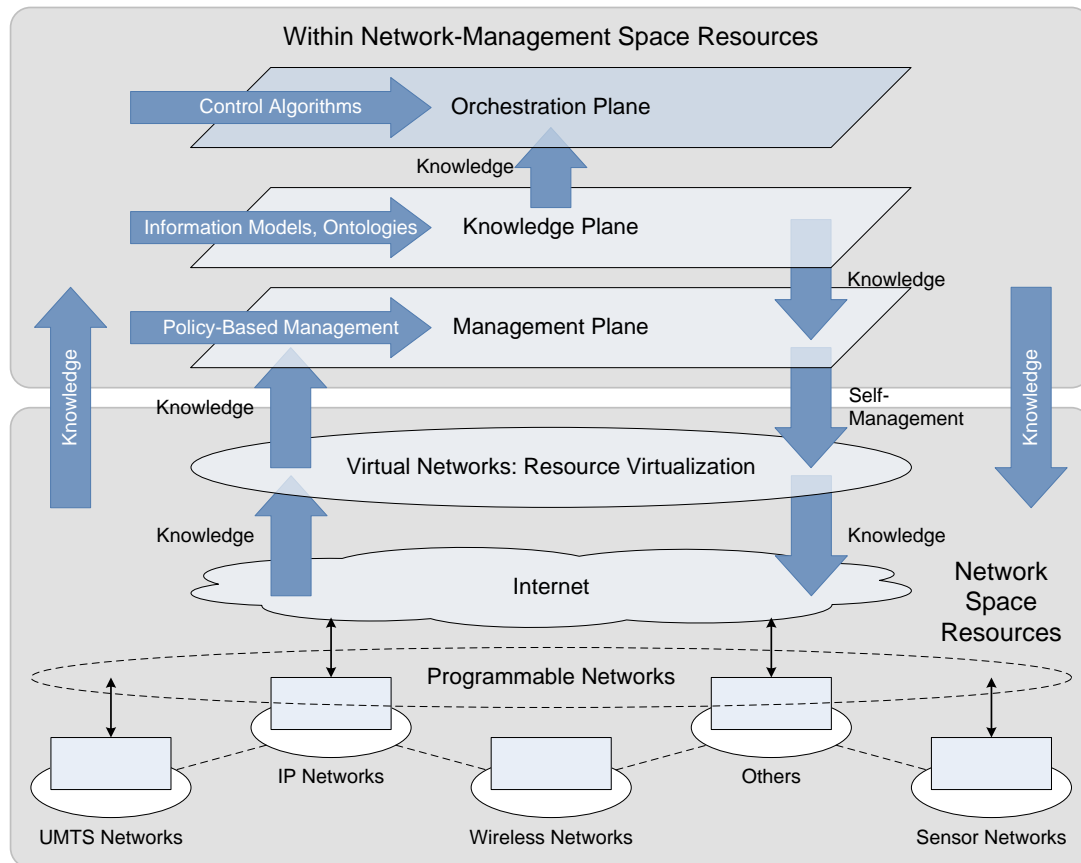
Figure 2. Autonomic Internet (AutoI) Architecture (from [11])

Figure 2 depicts an overview of the architecture and the OSKMV planes. As the AutoI architecture's ability to federate networks stems from its Orchestration Plane (OP) [12], the remainder of this section focusses on the OP and its capability to federate AMSs.

The OP governs the behaviour of the system in response to changing context and in accordance with applicable business goals and policies. It supervises all other planes' behaviour, in order to preserve integrity across the architecture. Its self-governing behaviour is achieved through a set of control algorithms, which can be plugged into it. The OP hosts one or more AMSs. Every AMS represents an independent administrative domain and is responsible for its own internal management. Specifically, the OP enables the federation of AMSs, negotiation of policies, distribution of management tasks and monitoring of AMS behaviour.

The OP is made up of a set of Distributed Orchestration Components (DOCs). A DOC is responsible for a single orchestration domain, which is in itself made up of multiple AMSs. It enables the AMSs of the orchestration domain to communicate and collaborate with each other. Additionally, DOCs can collaborate with each other, in order to provide end-to-end QoS. The DOCs thus serve as the facilitators of inter-domain federations. A DOC performs several tasks that together implement the ability of configure and manage federations. Specifically, these tasks are:

- **Distribution**: This component enables management tasks to be split across AMSs and executed concurrently, both within and across different network domains.
- **Negotiation**: The DOC enables the AMSs to negotiate their business objectives, in order to align them and achieve to a common set of goals for the federation. Two negotiation protocols are currently supported; coalition formation and bargaining.
- **Federation**: Allows a set of independent domains to be combined into a larger virtual domain, with a set of converged high-level goals (obtained through negotiation). The negotiation

process additionally aligns the internal domain Service Level Agreements (SLAs) and policies with the high-level federation-wide goals.

- **Governance**: AMSs are self-governing entities, that may decide to change their internal policies or SLAs. This might trigger incompatibilities between internal and federal policies and goals. The DOC's governance component monitors this, and takes appropriate action if such incompatibilities arise (e.g., it might trigger a re-negotiation of the federal agreement).

3.3. Summary and comparison

The presented architectures aim to achieve a similar vision; to enable the negotiation, configuration and management of dynamic network federations in the Future Internet. However, their approach and emphasis differs significantly. The LFM and its instantiations, such as the FRM, focus on the semantic interoperability and compatibility of models and information between independently managed network domains. To achieve this, they employ a semantic ontology mapping framework, which aligns the semantics of information models and context information in order to facilitate unambiguous communication. On the other hand, the AutoI architecture's OP concerns itself with the alignment and compatibility of intra-domain SLAs and policies with the federation-wide goals and high-level policies. Through a set of monitoring processes and SLA negotiation protocols it ensures this compatibility throughout the life-cycle of federations. We believe that, in order to achieve the vision of a Future Internet that supports dynamically adapting network federations, aspects of both architectures will need to be incorporated, guaranteeing both semantic interoperability as well as policy alignment. Additionally, several other challenges will need to be tackled, which will be identified and discussed throughout the next section.

4. STATUS AND CHALLENGES

The architectures discussed in the previous section conceptually describe how federated network management could be incorporated into the Future Internet. However, concrete algorithms, protocols and solutions are needed to implement the described architectural components. In this section, we identify the, in our opinion, most important challenges that need to be tackled before the federated network management vision can be fully achieved. Additionally, state of the art research that addresses these challenges is evaluated and the remaining open issues are discussed. The following technical challenges are considered:

1. **Security and trust**: Partners in a federation exchange (possibly sensitive) information about their internal management and operations, and allow external parties to access and modify their internal resources. Additionally, the federated Future Internet is expected to support federation agreements without any form of centralized authority. As such, there is a need for decentralized security and trust mechanisms capable of operating in a fully distributed setting without any centralized governing authority.
2. **Semantic interoperability**: The internal semantic representations of models and context information might differ significantly across interconnected network domains. To allow independently managed network domains to cooperate in dynamic and automatically managed federations, these (possibly incompatible) internal representations need to be aligned. This indicates the necessity for semantic mapping techniques that translate between differing internal semantics in order to facilitate unambiguous understanding across network domains.
3. **Agreement negotiation**: Before a federation can be set up, the participating parties need to come to an agreement about the associated costs, benefits, shared capabilities, goals and federation-wide policies. As the involved network domains might have different expectations, requirements and internal (possibly conflicting) policies, their views and goals need to be aligned through the use of negotiation protocols. Once an acceptable compromise has been achieved, the federal agreement can be finalized.

4. **Resource discovery and matchmaking**: The translation of high-level federation goals into specific capabilities, resources and configurations postulates the need for matchmaking and discovery mechanisms. The abstract capabilities and configurations must be mapped unto actual physical or virtual resources that offer the required functionality. This is achieved through scalable and distributed mechanisms that allow shared capabilities and resources, which satisfy specific requirements, to be discovered.

5. **End-to-end resource configuration**: The goal of setting up network federations is to provide some sort of added value to service consumers, which the individual federation partners cannot offer by themselves. To achieve this, capabilities and resources are shared among them. The federation's high-level goals must thus be translated into concrete capability and resource configurations.

6. **Management coordination**: In a federation of network domains, resources and capabilities of the individual domains cooperate in order to satisfy one or more federation-wide goals. The management of the shared capabilities must be coordinated across the participating network domains in order to ensure their cooperation in achieving the expected service benefits. This necessitates the end-to-end monitoring of their state and performance, as well as scalable communications mechanisms that allow federated domains to exchange information about internal resources and policies.

Figure 3 depicts the identified challenges and shows where they are positioned within the envisioned federated network management architecture of the Future Internet. The figure denotes the vertical relationships (between the different architectural components within a domain) as well as the horizontal relationships (between same architectural components across domains). Throughout the remainder of this section, the challenges are discussed in more detail.

4.1. Security and trust

The network domains participating in a federation interact in different ways. They exchange sensitive information (e.g., detailed topology information, router IP addresses or credentials to configure resources) about their internal management and operations, and allow external parties to access and modify their internal resources. To ensure the integrity of the federation's participants and prevent malicious tampering with internal resources and information, these interactions should be sufficiently secure and trust should be guaranteed. This challenge thus consists of two parts. First, secure communications channels should be provided as to guarantee secrecy, integrity and authentication of exchanged messages [10]. This can be achieved through public key authentication and encryption (e.g., IPSec, SSL, TLS). As this is a well researched topic, this aspect will not be considered further. Second, trust should be guaranteed between communicating parties in order to prevent the malicious or accidental propagation of fraudulent information and policies. In a centralized system, trust is relatively easily enforceable, since all security policies are centrally managed by a secure and trusted authority. However, in a fully distributed environment, such as a distributed federation of networks, several trust issues arise.

In a federation, the authorization to use internal resources and capabilities is given to external parties. In turn, these authorizations might be delegated to other parties. This gives rise to a specific trust problem called *authorization subterfuge*, which is highly relevant in the context of network federations. Foley and Zhou [13] define it as the problem where "*delegation chains that are used to prove authorization may not actually reflect the original intention of all the participants in the chain.*" Specifically, authorization subterfuge is caused by incompetence, confusion or dishonesty. Through incompetence, a delegated authorization credential might be ambiguously defined, given the recipient more access rights than intended. Confusion refers to unintended side effects, caused by authorization policies of which the delegator has no knowledge. Finally, dishonesty means providing third parties unauthorized access by intentionally exploiting incompetence and confusion or denying accountability by claiming to be incompetent or confused. Zhou and Foley [14] propose the Distributed Authorization Language (DAL) to overcome these issues. They argue that existing frameworks for guaranteeing security and trust in distributed systems rely on a centralized security
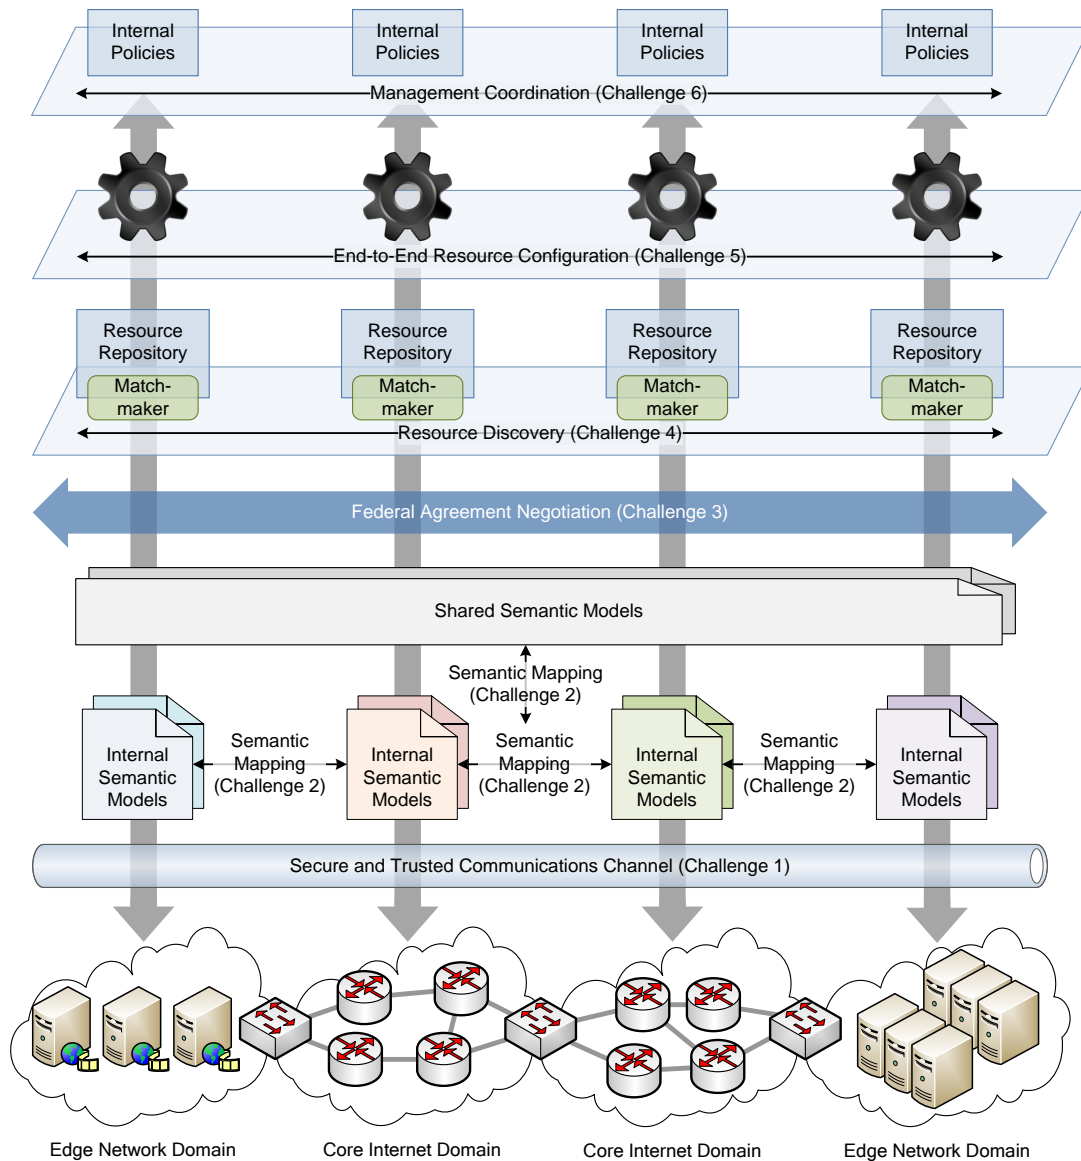
Figure 3. The identified challenges and their relationship to each other and the underlying network

administrator and do not consider authorization subterfuge. The DAL language, however, is the first step towards a fully distributed framework, without centralized control, that is subterfuge-safe.

The FRM framework (cf. Section 3.1) is based on a trust model for secure delegation of capabilities in federated systems [15], referred to as the *capability authority model*. This model describes the capabilities provided by a domain and how authority to invoke and manage them is distributed to domains though ownership relations and delegation. It supports two modes; non-trusted and trusted. In non-trusted mode, every domain keeps track of the domains that have been given access rights to specific capabilities. The usage of capabilities always follows the chain of delegation. However, following the chain of delegation might become a performance bottleneck in large-scale scenarios with many embedded federal agreements. The trusted mode can be used to overcome this. It allows delegated capabilities to be directly invoked by any domain that holds its access rights. Through guaranteed uniqueness of permission identifiers, delegation subterfuge is prevented.

Another problem that arises in the context of distributed security and trust is the interoperability of heterogeneous access control mechanisms. Although the generic semantic interoperability problem is discussed further in Section 4.2, we consider here the specific interoperability problems arising in the context of access control. Traditionally, the interoperation problem of access policies is solved by enforcing a common vocabulary for policy attributes [16]. Martínez-García *et al.* [17] argue that this cannot be assumed in independently designed and managed systems. They chose another approach to tackle the problem, which converts policy attributes from one representation to another. This allows access control mechanisms to interoperate without the need for a shared vocabulary. Concretely, a fuzzy set theory approach is proposed, which determines the relationship between attributes using membership functions. The approach is extended to the interoperability problem between more than two domains, where the number of possible mappings increases exponentially. This is alleviated by chaining attribute conversion mappings together. Every domain maintains mappings for a small subset of trusted other domains. Creating a mapping to another, unknown, domain can be done by first combining the mapping of a known domain, with one or more other mappings maintained by other domains.

Traditional frameworks and mechanisms for security and trust in heterogeneous and distributed environments were based on a trusted central management component that maintains and checks access rights and security policies. Additionally, they do not consider federation-specific issues, such as delegation subterfuge and interoperability. Recently, some advances have been made towards fully distributed trust management systems, without the need for a central managing entity and built-in methods to detect and overcome delegation subterfuge. These research efforts should be combined to provide a framework for configuring and maintaining dynamic federal trust relationships, capable of interoperating with heterogeneous intra-domain access control systems.

4.2. Semantic interoperability

Automatically configuring and managing dynamic federations among network domains requires autonomic agents located within these independently managed networks to interact and communicate. They must be capable of discovering and configuring shared resources and negotiate mutually beneficial agreements. This is only possible if these autonomic agents are capable of understanding each other and the information they exchange. Unambiguous understanding between autonomic entities is generally achieved through the use of shared semantic models [18]. This facilitates the autonomic understanding and interpretation of policies [19], context information [20], services [21] and shared resources [8]. However, it is infeasible to assume the same set of semantic models can be shared among all network domains in a federated Future Internet. This necessitates the need for mapping the semantics used internally by the different network domains [8, 2]. The importance of this challenge was first indicated by Jennings *et al.* [7] in 2009. It is also reflected in the LFM (cf. Section 3.1), where the *shared semantic layer* is responsible for mapping the diverse internal semantic models unto a standardized semantic language or unto one another (if such a standardized language is missing).

The creation of mappings between semantic models has been the topic of much research in the recent past [22, 23], and is often referred to as *ontology mapping*. According to Choi, Song and Han [23], ontology mapping techniques can be classified into three categories; (1) mapping a global into a set of local ontologies, (2) mapping between local ontologies and (3) mapping on ontology merging and alignment. Semantic mapping in the context of network federations is mainly related to the second category. In the context of this category, Choi defines ontology mapping as "*the process that transforms the source ontology entities into the target ontology entities based on semantic relation.*" It is most useful for highly dynamic, open and distributed environments, such as network federations in the Future Internet. Traditional ontology mapping tools rely on significant intervention of domain experts or knowledge engineers to aid in the mapping process [24, 25, 26]. This results in constrained and static mappings, which makes these approaches ill-suited for dynamic and automatically configured and managed network federations [7]. Besana *et al.* [27] propose an algorithm specifically focussed on mapping diverse and dynamic ontologies with only partially overlapping knowledge domains. They claim that it is infeasible to create all

possible mappings in advance, due to the huge amount of possible combinations. Additionally, mapping complete ontologies at runtime is a computationally expensive task. Although they note these problems in the context of semantic peer-to-peer networks, they also arise in the context of network federations, where many network domains can potentially interact with one another. To solve these problems, Besana proposes a novel algorithm that maps only those ontological concepts that are relevant for the interaction at hand. Specifically, the mapping framework dynamically maps concepts when they are first encountered during an interaction. This mapping process is iterative and consists of three steps; hypothesis generation, filtering and selection. Although the framework aims to automate the entire mapping process eventually, the described algorithm does not offer a solution for all steps, necessitating the intervention of human experts.

The semantic interoperability problem has also been studied in the more specific context of Future Internet network management. Strassner, Serrano *et al.* introduced the inference plane [28, 29] an evolution of Clark's knowledge plane [30]. The original knowledge plane does not consider the heterogeneity of technologies, devices and information models in independently managed network domains. The inference plane extends the knowledge plane to be able to cope with this through semantic interoperability. Specifically, they propose the use of a *lingua franca*, based on a set of common information models and ontologies. This set of common models can then serve as a lexicon to translate the different semantic models into a mutually understandable format. Although this approach does not require administrative domains to change their internally used semantics and models, it does obligate them to be able to translate their internal models into the globally agreed upon *lingua franca*. Wong *et al.* [31] proposed a semi-automatic ontology mapping algorithm for the communication between network domains. Although their approach is capable of automatically mapping different ontologies unto one another, the algorithm has several parameters that need to be manually configured by domain experts. As the configuration of these parameters depends on the nature of the interaction, human intervention is required whenever the context within which the network domain interacts changes. Finally, the FRM (cf. Section 3.1) incorporates a semantic mapping framework [8, 32], based on the ontology mapping approach presented by O'Sullivan *et al.* [33]. When a new federal relationship is created, the FRM attempts to re-use and adapt known mappings to facilitate the interoperability between the domains taking part in the relationship. The original mappings that are used as a basis for this process could be created through manual or (semi-)automated processes. They argue that re-use avoids unnecessary redundancy and prevents an explosion in the number of created and deployed mappings. Additionally, this allows the FRM to efficiently deal with dynamism of data and schemata, which they argue has been ignored in research to date on semantic interoperability.

In conclusion, we found that existing methods for the facilitation of semantic interoperability require significant intervention by human experts. This obviously hinders the automatic and dynamic configuration and management of network federations. However, it has been argued that the fully automatic generation of mappings between semantic models is difficult, if not impossible, due to the uncertainty related to matching two ontologies or other semantic models [34]. To accomplish the vision of fully automated federation life-cycle management, human interventions in the creation of mappings should be performed in an offline manner. This can be achieved by employing a common set of information models, for which mappings can be created in advance [28]. This would allow network domains to employ internal semantics for intra-domain management, while the shared semantic models would serve as a *lingua franca* to facilitate semantic interoperability in inter-domain affairs. This has the added advantage of preventing an explosion in the number of mappings, which would have to be created if no common shared models exist. Finally, most existing semantic mapping approaches are concerned with mapping entire models unto one another, which is a computationally expensive operation. In order to guarantee scalability, semantic mapping approaches should be able to determine the relevant subset of models, based on the federation's context. Some promising early work on this topic [27] was performed in the context of peer-to-peer networks.

4.3. Agreement negotiation

As stated in the definitions presented in Section 2, a network federation is characterised by an agreement between the participating parties. This agreement formally stipulates the rights and obligations of the involved network domains, which usually pertain to a set of shared resources or capabilities. Additionally, the agreements should specify the revenue sharing strategy, which determines how monetary gains, if any, are split across the participants. Today, the negotiation of federal agreements between network domains is a manual and time-consuming process, where business managers, lawyers, and network operators define the business, legal and technical aspects that the participating parties must adhere to. In the Future Internet, where federations will be dynamically and automatically initiated based on changing needs and requirements, this manual process should be (at least partly) automated. We envision an agreement negotiation mechanism that is automatically executed by autonomic agents. Nevertheless, they perform these negotiations within the bounds specified by high-level business, legal and technical policies, defined by human managers and operators. As such, humans are no longer directly involved in these negotiations, but can influence and govern them through the definition of high-level policies.

Automated negotiation protocols have been most commonly proposed in the context of SLA negotiation. An SLA is a formal agreement between a service provider and its customer. It specifies the terms under which a service is delivered, such as the quality, availability, or QoS guarantees. Several standardization efforts exist for the negotiation of Web Service SLAs. The Web Service Level Agreement (WSLA) specification was first proposed by IBM in 2001 [35]. It addresses the specification, creation and monitoring of SLAs. Concretely, the SLAs specify the obligations of the service provider, in terms of IT-level service parameter guarantees (e.g., availability, response time and throughput). Additionally, WSLA specifies the measures that should be taken in case the provider fails to meet its obligations. Although WSLA is equipped to incorporate an automated negotiation protocol, the actual protocol is outside its scope. More recently, the Web Services Agreement (WS-Agreement) [36] specification was introduced by the Open Grid Forum. It is a Web Service protocol for establishing agreements between parties, and has goals similar to WSLA. The agreements themselves are specified in an XML-based language. WS-Agreement incorporates three main components; a schema for specifying agreements, a schema for specifying agreement templates and a set of operations for managing their life-cycles (i.e., creation, expiration and monitoring). In contrast to WSLA, WS-Agreement does propose its own negotiation protocol [37]. It allows two parties to negotiate on the terms of an agreement. If a compromise is reached, the negotiation results in the creation of an actual agreement using the WS-Agreement specification. Hudert et al. [38] extended the WS-Agreement specification, adding support for multilateral, in additional to bilateral, negotiations. However, they do not propose an actual multilateral negotiation protocol.

The composition of resources or services positioned within independent managed domains is especially relevant within the context of Grid and Cloud Computing. As such, the automatic negotiation of SLAs has been a topic of interest within these fields for several years. Hasselmeyer *et al.* [39] presented a framework for SLA management in Grids, which incorporates a simple one-phase *discrete-offer-protocol*. It lets the customer send a request for an offer to the service provider. The provider then decides whether to reply with an offer or not (based on its available WS-Agreement SLA templates). If it replies with an offer, the customer can either accept or reject. The protocol thus supports bilateral negotiations, without the possibility for compromise (i.e., the customer can only accept or reject the initial offer, not make a counter-offer). Recently, Parkin *et al.* [40] extended the framework with a more elaborate SLA negotiation protocol. The protocol is a multi-round re-negotiation protocol. The multi-round aspect means that if the initial offer is not accepted, one or more additional offers can be made in an attempt to reach a suitable compromise. Additionally, the protocol supports re-negotiation of existing SLAs, to accommodate changing requirements and business goals. Like its predecessor, the protocol focusses on bilateral negotiations, between a service provider and its customer.

Yan *et al.* [41] propose a generic SLA negotiation protocol, that supports multiple service providers. Specifically, it allows a single customer to negotiate the provisioning of a complex

service, of which the individual components are offered by different service providers. The consumer is represented by a set of agents who negotiate with the individual service providers. A coordinating agent makes sure that the individually negotiated SLAs together satisfy the customer's end-to-end QoS requirements. The agents use the *FIPA iterated contract net interaction protocol* [42]. It is a bilateral one-to-many agent negotiation protocol. A single agent (the initiator) requests an offer from a set of other agents (participants). The participants may reply with an offer, or refuse. The initiator then iteratively repeats this process with the remaining participants, until a suitable offer is made or all offers have been refused. As such, Yan's framework does not actually support multilateral negotiation, but rather transforms it into a set of coordinated bilateral negotiations.

The importance of agreement negotiation, within the context of a federated Future Internet, was first recognised by the AutoI project consortium. AutoI's OP (cf. Section 3.2) incorporates support for the negotiation of agreements between AMSs [12]. Rubio-Loyola *et al.* [43] propose an algorithm to negotiate service provider coalitions, for AutoI's OP. The algorithm is based on an electronic marketplace, where every service provider publishes its service offerings and associated guarantees. The algorithm acts as a centralized manager, that coordinates the negotiations on behalf of the customer. The algorithm thus supports the negotiation of multilateral agreements, but requires a trusted central management entity. More recently, Chai *et al.* [44] proposed an alternative negotiation protocol for the AutoI OP. The protocol supports bilateral negotiation between two AMSs and additionally assumes the OP plays a coordinating role. Specifically, it is based on the concept of alternating offers. The two negotiating parties take turns in making an offer. The other party accepts, rejects, or opts out. If the offer is rejected, the other party must make a counter-offer, otherwise the negotiation ends. The probability that a participant accepts an offer is based on a utility function, its patience (which decreases over time) and estimated risk. The utility is a function of the amount of requested resources, the agreements duration, the expected benefits, and the current offer.

The envisioned Future Internet should support federations consisting of many (i.e., more than two) independent network domains, without any need for centralized control or management. Existing agreement negotiation protocols are usually designed for bilateral negotiations [37, 39, 40, 42, 44], which is inconsistent with our vision of large-scale federations. Recently, some protocols for the automated negotiation of multilateral agreements have been presented [38, 43]. They, however, expect a centralized trusted management component that coordinates the negotiation process. Before multi-party federations without any form of centralized control can become a reality, the gap towards an automated, multilateral, fully distributed agreement negotiation protocol must thus be filled. Additionally, existing protocols that support iterative negotiations are usually based on a set of mathematical utility functions that model their satisfaction as a function of the current offer and time. These functions need somehow be mapped unto the human specified business, legal and technical high-level policies that should constrain and guide the negotiation process. This aspect has not been discussed in state of the art research and remains an open issue.

4.4. Resource discovery and matchmaking

A federation of network domains is set up in order to achieve a common goal that the federation partners cannot achieve alone. However, before the federation agreement can be negotiated, the initiator must determine the set of network domains and shared resources/capabilities that can achieve this goal. As the Internet consists of many thousands of network domains and many more shareable resources and capabilities, scalable and distributed resource discovery mechanisms are needed that are capable of mapping generic federation goals into concrete physical and virtual resources and capabilities. This challenge thus consists of two closely related topics: (1) the actual discovery of resource and capability descriptions, and (2) matching or mapping goals unto those descriptions. The first topic is referred to as resource discovery, while the second is called matchmaking.

Scalable resource discovery has been a well researched topic in the areas of Grid and, more recently, Cloud Computing. It has become especially relevant within these areas since the introduction of federated grids and clouds. Ranjan *et al.* stated that a resource discovery mechanisms for global, or federated, grids should should be scalable, fault tolerant and impart a limited overhead

on the underlying network [45]. Traditional resource discovery approaches for federated Grids used centralized or hierarchical resource indexing services. Especially the centralized, but also the proposed hierarchical methods are prone to central points of failure and scale poorly to a large number of resources. As a solution, peer-to-peer-based resource indexing and discovery protocols were proposed. Early attempts used unstructured peer-to-peer networks in combination with flooding to broadcast the set of available resources [46, 47]. However, flooding-based peer-to-peer protocols are known to scale poorly in terms of generated network overhead [45]. In an attempt to reduce network overhead, solutions based on structured peer-to-peer networks were proposed [48, 47]. They often use a Distributed Hash Table (DHT) [49, 50], as an underlying routing substrate. In a DHT, data is stored as (key, value) pairs, which can be looked up in a logarithmic number of hops. Resource discovery mechanisms based on DHTs are usually based on mapping a d-dimensional logical key to the 1-dimensional DHT key-space. Every dimension then corresponds to a specific attribute of a grid or cloud resource. In a computational grid, these attributes would be for example CPU, memory, bandwidth and cost. Although this is a viable solution for grids and clouds, where the resource types and attributes are limited and known at design time, this information is unknown in the context of generic network domain federations, where a huge number of different types of resources and capabilities, with widely varying attributes, are available for sharing. Heine *et al.* [51] solve this problem by using an ontology to represent resources and their attributes. This approach allows new attributes and resource types to be added at runtime. However, the d-dimensional queries are split up into d 1-dimensional queries and merged after the results have been returned. This causes potentially huge amounts of useless information to be propagated through the network. Additionally, their approach only supports the equality operator, and not more complex comparison operators. Pipan [52] identified some further drawbacks of DHT-based resource discovery, such as it lack of adaptability in highly dynamic scenarios (i.e., where resources and their attributes change often) and its difficulty to handle rich resource descriptions with many attributes. He proposes a novel overlay network, called TRIPOD, which combines the advantages and reduces the disadvantages of existing structured overlays. Specifically, it is capable of finding resources based on proximity, efficiently processes complex queries and handles dynamics in resource characteristics well.

Matchmaking is most often defined in the context of services, and is concerned with determining the set of services that match a given set of requirements. Traditional methods employ keyword-based matching. However, this has been shown to lead to low matching precision, due to lack of semantics [53]. More recent algorithms employ semantic service descriptions to improve matching precision and make them machine-understandable. These novel semantic matchmaking algorithms are also more suitable for use in federations, as semantics are an important mechanism to guarantee interoperability between network domains. Several semantic matchmaking algorithms have been proposed for web services. Most early work focussed on matching inputs and outputs [54]. However, more recent algorithms have started taking into account preconditions and effects. Preconditions model the state the environment must be in before the service is executed, while effects define the changes that the service will have on the environment. These algorithms are based on various semantic techniques, such as description logics [53], SWRL rules [55, 20] and SPARQL queries [56]. The described works on semantic matchmaking focusses heavily on matching software services. However, resources and capabilities in a generic network federation should be interpreted more broadly, and could for example be physical server resources, network paths, device configurations or software components. Nevertheless, some matchmaking algorithms have been proposed for the matching and discovery of generic resources. Islam *et al.* [57] devised a matchmaking framework for generic resource discovery. The proposed matchmaking algorithm uses a tree-based search method, which results in a linear computational complexity. They argue that this is a significant reduction compared to other state of the art matchmaking algorithms, which have a cubic computational complexity. However, its lack of support for semantics impedes its use in a federated management scenario.

Jennings, Feeney *et al.* [8, 2] recognised the importance of service discovery for network federations in their LFM (cf. Section 3.1). The shared capabilities layer is, among other things,

responsible for discovering the capabilities that are available at any particular time. Additionally, the FRM [8] contains an instantiation of this functionality, in the form of the "Capability Publication and Discovery" component. It uses an authenticated SPARQL endpoint to find capabilities based on RDF descriptions. The RDF documents describe the web service entry points of the actual underlying capabilities. Although this approach is useful for finding suitable capabilities offered by a known federation partner or candidate. It does not support capability discovery from an unknown source. The authors explicitly stated that this problem is outside the scope of their work.

Resource discovery research to date has mostly focussed on Grid and Cloud Computing scenarios. Although there are some parallels with the envisioned federated Internet scenario, there are also several differences that make a direct application of existing methods difficult. In Grid and Cloud Computing, there are a limited number of resource types and attributes. State of the art resource discovery methods thus assume a limited set of resource types and attributes, that are additionally known at design time. In contrast, a huge variation in resource types and attributes is expected to be encountered in the federated Future Internet. Work to date on semantic matchmaking has focussed on matching objectives to software services. However, the shared resources in a federation of networks could also refer to, for example, hardware components, network capabilities, or device interfaces. As such, existing description methodologies, as well as matchmaking algorithms, need to be adapted to encompass this. In the context of generic network federations, the discovery of shareable capabilities within a known network domain, and linking them to federation goals has been studied to some extent. However, it is assumed that the candidate domains to include within a federation are known. Determining this set of candidates is an important open question that remains to be solved.

4.5. End-to-end resource configuration

The negotiated high-level federation goals need to be mapped unto concrete resource and capability configurations. In the envisioned service-driven Future Internet, such goals usually relate to the added value of delivered end-to-end services. An important driver for federations is the provisioning of end-to-end QoS across a set of independently managed network domains, which cannot be guaranteed in the current best-effort Internet. Network domains will need to cooperate in order to provision end-to-end paths that satisfy bandwidth and other network parameter requirements. Several evolutionary and revolutionary methods have been proposed to extend the Internet with QoS reservation capabilities, through the federation of Internet routing domains.

Evolutionary approaches usually propose extensions to the de-facto standard inter-domain routing protocol BGP (Border Gateway Protocol). Kumar and Saraph [58] propose such an evolutionary solution, based on the Routing Control Platform (RCP). They propose the *Alliance Network* model, which allows Autonomous Systems (AS) to join into federations in order to provide end-to-end QoS for end-users. The QoS-guaranteed paths through the participating ASs are identified and configured using the Virtual Space (VS) routing algorithm [59]. However, their model assumes prior agreements on revenue sharing and information exchange have been negotiated. Due to its semi-static nature, the model therefore does not scale up to the global Internet.

Pouyllau and Douville [60] identified several shortcomings associated with extending BGP to support inter-domain QoS-guaranteed routing, such as scalability and confidentiality problems. To alleviate this, they propose a revolutionary approach, based on the negotiation of SLAs. Every intermediary network domain, referred to as a carrier, offers a set of Service Level Specifications (SLSs) (i.e., the shared capabilities). Every SLS is related to the reservation of a QoS-guaranteed path through the associated carrier domain. The proposed algorithm maps the customer's QoS requirements to a chain of SLSs that together form a QoS-guaranteed end-to-end path. The carriers associated with the selected SLSs subsequently negotiate to form a federation. The composition problem is modelled using the game theory approach and solved using an algorithm based on Q-learning [61]. Their approach, however, assumes that the federation is negotiated and configured with the help of a third party, which has complete knowledge about the capabilities of each candidate carrier.

Recently, we presented the FedRR algorithm [62]. Its goal is to set up network federations to support end-to-end QoS-guaranteed paths across multiple network domains. It identifies the network domains that need to be included within the federation, as well as the capabilities (i.e., QoS classes and network paths) that need to be shared and reserved within each identified domain. Additionally, it allows cloud providers to be included within the federations supporting the dynamic deployment of content caches inside the network. In line with to Pouyllau's approach, it assumes the set of candidate federation partners, as well as their capabilities, are known by a central governing entity.

Work to date on the transformation of federation goals into specific capability configurations has heavily focussed on specific scenarios (e.g., end-to-end QoS), where the goals, and thus the required capabilities, are known at design time. This circumvents the need to dynamically map goals to specific shared capabilities. However, in order to support dynamic, automatically configured and managed federations with varying goals and requirements, there is need for more intelligent translation mechanisms. They should be able to determine a set of candidate federation partners based on their offered capabilities and the specific requirements of services and end-users. Additionally, existing federated resource configuration methods assume the complete set of network domains and capabilities is known. However, in a network as large as the Internet, consisting of tens of thousands of independently managed network domains as well as millions of soft- and hardware resources, this is an infeasible assumption. To guarantee scalability of federation management architectures, resource configuration will thus need to be combined with scalable mechanisms to discover and match shared capabilities, as discussed in Section 4.4. Finally, it is often assumed that a central governing entity oversees the configuration and coordination of resources and capabilities. However, this assumption is inconsistent with the vision that the Future Internet should support fully distributed federations, without the need for a central governing body [2]. To support this vision, resource configuration algorithms and processes need to be adapted to operate in a fully distributed environment, without centralized control and management.

4.6. Management coordination

The network domains participating in a federation are expected to collaborate in a coordinated fashion in order to achieve the federation-wide goals. To achieve this, there is a need for distributed management coordination mechanisms that govern the behaviour of individual participating network domains, and monitor the end-to-end state of federated resources and services. Specifically, it should be possible to detect and solve conflicts between federation-wide goals and policies on one hand, and intra-domain policies on the other. Additionally, scalable information dissemination substrates are needed in order to facilitate the end-to-end monitoring of internal domain states.

The dissemination of aggregated monitoring information and context is necessary in order to guarantee the continuous satisfaction of federation-wide goals. Additionally, to ensure inter-domain understanding and interoperability, the exchanged information should be semantically annotated. The publish-subscribe paradigm is well suited to offer these functionalities. It allows interested parties to subscribe to specific types of events. When an event that matches the subscription is published, it is routed accordingly. The paradigm has been successfully applied to the dissemination of semantic information in large-scale networked environments under the banner of *knowledge based networking* (KBN) [63]. It is an extension of content based networking (CBN) [64], which involves the forwarding of events across a network based on subscription filters based on the semantics of the (meta-)data of the event's contents. KBN extends this and states that the semantics of messages play an important part in the matching of publications to subscriptions. To this end, the Sienna publish-subscribe system, originally devised for CBN, was extended with more expressive semantics for the specification of subscriptions [65] to satisfy the KBN vision. Messages in the original Siena take the form of a set of typed attributes (i.e., name-value pairs). Filters specify constraints on the values of those attributes. The filtering process is thus purely based on the syntactical form of messages. Siena additionally supports patterns, which allows matching on combinations of messages. Carzaniga *et al.* [64] additionally propose a set of efficient and scalable routing strategies to forward messages from publishers to interested subscribers. The KBN extension [65], proposed by Keeney *et al.*, adds limited support for semantic messages and filters. Specifically,

three new attribute types are added: ontological properties, concepts and individuals. Through basic ontological reasoning, filtering can be done based on semantic equivalence, super- and subconcept relationships, and property relationships. Throughout the years, several other semantic publish-subscribe frameworks have been proposed. They have varying degrees of semantics and inferencing power, ranging from simple RDF graph matching [66, 67], to full-fledged OWL [68, 20] and SWRL-based reasoning [20]. However, it is difficult to find a good balance between expressive and inferencing capabilities on one hand, and routing performance and scalability on the other. Increasing the expressiveness of the messages and filters, will generally lead to significantly reduced scalability.

An aspect of context dissemination that has been mostly ignored, is the actual specification of subscriptions or filter rules. In the envisioned Future Internet, where networks are autonomously managed and federations are automatically created, the generation and adaptation of subscriptions and filter rules needs to be automated as well. This process should take into account the requirements and goals of a specific federation, as well as changes in the state of the environment and requirements. Latré *et al.* [69] first identified this problem, and proposed an algorithm to generate semantic filter rules, using an OWL-based reasoning approach. The algorithm takes into account the dynamic requirements of autonomic management components, as well as the changing state of the managed environment. Based on these inputs it generates and adapts filter rules for use in semantic context dissemination frameworks.

Another important aspect of the end-to-end coordination of federations, is the alignment and compatibility of intra-domain with federation-wide policies and goals. This is an important focal point of the AutoI OP (cf. Section 3.2) [12]. The AutoI DOC component hosts a set of behaviours, which describe the specific orchestration tasks it performs. The *governance behaviour* is concerned with the alignment and compatibility of goals and policies. Network domains might independently change their internal policies or requirements, which might lead to federation-wide inconsistencies. Specifically, the DOC performs three tasks to detect and correct such inconsistencies. First, it monitors the management actions performed within domains and verifies the alignment between internal configurations and federation goals. Second, if a conflict is detected, the DOC informs the management components of the offending domains. Third, if necessary the DOC will trigger a renegotiation of the federal agreement in order to ensure continued smooth operation of the network. The analysis and resolution of conflicts between policies is a complex topic in itself. Research to date has mostly focussed on the policy conflict analysis within single network domains. In an attempt to extend this research topic to federations, Barron *et al.* extended the DEN-ng policy model to support federation-wide policies [70]. More recently, they outlined a novel policy conflict analysis algorithm for federal management policies [71] based on this extended model.

An important aspect of the federation of network domains is the end-to-end coordination of management behaviour, as well as the alignment of internal and federation-wide goals and policies. To guarantee inter-domain coordination, there is a need for scalable mechanisms to exchange and correlate semantic monitoring information. Existing work on semantic context dissemination either offers good scalability with limited expressiveness or the other way around. It is widely believed good routing performance and scalability cannot be combined with extensive expressive power. The combination of both aspects still requires further study. Additionally, relatively little research has been done on the topic of automatic generation of subscriptions and filter rules. In order to achieve the fully automated configuration and management of federations, this topic should be further explored. Another important aspect of management coordination is the detection and resolution of local and federal policy conflicts. Research in the area of policy conflict analysis between local and federal policies is very limited. Additionally, the complex topic of conflict resolution has not been addressed at all in the context of federations.

5. CONCLUSION

This article presents an in-depth survey of state of the art research on federated management of the Future Internet. The article offers several distinct contributions. First, the plethora of

definitions of the term *network federation* introduced throughout the years were compared and aligned. Specifically, we attempted to come to a unified vision of federated network management, by combining the important aspects of these existing views. Second, an overview was given of the two most influential Future Internet architectures that include support for federations of network domains; the Layered Federation Model (LFM) and the Autonomic Internet (AutoI) architecture. Third, several important challenges related to the envisioned federated Future Internet were identified and discussed. The state of the art research that addresses these challenges was thoroughly evaluated and remaining gaps were identified. This led to several pertinent conclusions, of which we consider the following to be the most important:

- High-level federation goals should be mapped to specific capability and resource configurations within the participating domains. However, work to date on this topic has focussed on very specific scenarios where the goals, and consequently the mapping to resources, are known at design time. This circumvents the need for dynamic algorithms that map generic federation goals unto specific resource configurations. However, we believe such generic algorithms are needed in order to support generic federations of network domains, with widely varying goals and tasks.
- Additionally, existing end-to-end resource configuration frameworks often assume the existence of a centralized management entity that coordinates the configuration effort. This is incompatible with the vision of fully distributed federations without the need for centralized governance or control.
- The automatic configuration and management of federations requires significant communication and collaboration between automated management components. To allow these components to interact with mutual understanding, the internal semantics of the associated network domains need to be aligned. Existing methods to facilitate semantic interoperability (e.g., based on ontology mapping) rely on significant at-runtime interventions by domain experts. By adopting a standardized set of shared information models and semantics, the semantic mapping process can instead be done in an offline manner. This would negate the need for at-runtime human intervention, but would require all involved network domains to adopt standardized models and translate their internal models and semantics in advance.
- Existing methods to discovery resources across the bounds of administrative domains were designed specifically for Grid and Cloud Computing scenarios. In such scenarios, the types of resources and their attributes are limited and known at design time. Existing resources discovery methods exploit this assumption and can therefore not be directly applied to a generic federation scenario. Novel resource discovery algorithms and protocols thus need to be designed, capable of handling a huge number of different resources types and attributes, not necessarily known at design time.
- Moreover, the network domains in which discoverable resources reside are assumed to be known. In a large-scale scenario with many thousands of potentially collaborating network domains, this assumption is infeasible. There is thus need not only for efficient resource discovery mechanisms in known domains, but also scalable techniques to identify and select the set of potential federation partners among a huge number of available network domains.
- In order to successfully configure a federation of networks, they should be able to negotiate a federal agreement. Existing protocols for agreement negotiation often support only bilateral negotiations. The few protocols that do support simultaneous multilateral negotiations, rely on a centralized trusted component that governs the negotiation process. As federations are expected to possibly contain a large number of network domains without any centralized governing entity, such existing agreement negotiation protocols cannot be directly applied. Novel protocols that support fully distributed multilateral negotiations, need to be devised.
- An important aspect of federation life-cycle management is the alignment of internal and federation-wide policies and goals. Most research on the topic of policy conflict analysis and resolution considers only single-domain scenarios. Although some research has been done

in the area of conflict analysis of federation policies, the subject of conflict resolution in federated scenarios is yet to be addressed.

ACKNOWLEDGEMENT

REFERENCES

1. Cisco Systems. Cisco visual networking index: Forecast and methodology, 2010–2015. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf 2011. Last accessed: 8 March 2012.
2. Jennings B, Feeney K, Brennan R, Balasubramaniam S, Botvich D, van der Meer S. Federated autonomic network management systems for flexible control of end-to-end communications services. *Autonomic Network Management Principles: From Concepts to Applications*. Elsevier, 2011; 101–118, doi:10.1016/B978-0-12-382190-4.00005-X.
3. Szegedi P, Figuerola S, Campanella M, Maglaris V, Cervello-Pastor C. With evolution for revolution: Managing FEDERICA for future internet research. *IEEE Communications Magazine* 2009; **47**(7):34–39, doi:10.1109/MCOM.2009.5183470.
4. Wahle S, Harjoc B, Campowsky K, Magedanz T, Gavras A. Pan-european testbed and experimental facility federation – architecture refinement and implementation. *International Journal of Communication Networks and Distributed Systems* 2010; **5**(1):67–87, doi:10.1504/IJCNDS.2010.033968.
5. Serrano M, van Der Meer S, Holum V, Murphy J, Strassner J. Federation, a matter of autonomic management in the future internet. *Network Operations and Management Symposium (NOMS)*, 2010; 845–849, doi:10.1109/NOMS.2010.5488357.
6. Serrano M, Davy S, Johnsson M, Donnelly W, Galis A. Review and designs of federated management in future internet architectures. *The Future Internet*. Springer Berlin Heidelberg, 2011; 51–66, doi:10.1007/978-3-642-20898-0.
7. Jennings B, Brennan R, Donnelly W, Foley SN, Lewis D, O'Sullivan D, Strassner J, van der Meer S. Challenges for federated, autonomic network management in the future internet. *International Symposium on Integrated Network Management (IM)*, 2009; 87–92, doi:10.1109/INMW.2009.5195942.
8. Feeney K, Brennan R, Keeney J, Thomas H, Lewis D, Boran A, O'Sullivan D. Enabling decentralised management through federation. *Computer Networks* 2010; **54**(16):2825–2839, doi:10.1016/j.comnet.2010.07.006.
9. Celesti A, Tusa F, Villari M, Puliafito A. How to enhance cloud architectures to enable cross-federation. *IEEE 3rd International Conference on Cloud Computing (CLOUD)*, 2010; 337–345, doi:10.1109/CLOUD.2010.46.
10. Brennan R, Feeney K, Keeney J, O'Sullivan D, Fleck J, Foley S, van der Meer S. Multidomain IT architectures for next-generation communications service providers. *IEEE Communications Magazine* 2010; **48**(8):110–117, doi:10.1109/MCOM.2010.5534595.
11. Galis A, Denazis S, Bassi A, Giacomin P, Berl A, Fischer A, de Meer H, Strassner J, Davy S, Macedo D, *et al.*. Management architecture and systems for future internet networks. *Towards the Future Internet – A European Research Perspective*. IOS Press, 2009; 112–122, doi:10.3233/978-1-60750-007-0-112.
12. Macedo D, Movahedi Z, Rubio-Loyola J, Astorga A, Koumoutsos G, Pujolle G. The AutoI approach for the orchestration of autonomic networks. *Annals of Telecommunications* 2011; **66**(3):243–255, doi:10.1007/s12243-010-0187-x.
13. Foley S, Hongbin Z. Authorization subterfuge by delegation in decentralized networks. *13th International Conference on Security Protocols*, 2005; 97–102, doi:10.1007/978-3-540-77156-2_12.
14. Zhou H, Foley S. A framework for establishing decentralized secure coalitions. *19th IEEE Computer Security Foundations Workshop*, 2006; 270–282, doi:10.1109/CSFW.2006.5.
15. Feeney K, Brennan R, Foley S. A trust model for capability delegation in federated policy systems. *6th International Conference on Risk and Security of Internet and Systems (CRiSIS)*, 2011; 1–8, doi:10.1109/CRiSIS.2011.6061828.
16. Gong L, Qian X. Computational issues in secure interoperation. *IEEE Transactions on Software Engineering* 1996; **22**(1):43–52, doi:10.1109/32.481533.
17. Martínez-García C, Navarro-Arribas G, Foley S, Torra V, Borrell J. Flexible secure inter-domain interoperability through attribute conversion. *Information Sciences* 2011; **181**(15):3491–3507, doi:10.1016/j.ins.2011.04.023.
18. Strassner J, Agoumine N, Lehtihet E. FOCALE – a novel autonomic networking architecture. *International Transactions on Systems Science and Applications* 2007; **3**(1):64–79.
19. Strassner J, Neuman de Souza J, van der Meer S, Davy S, Barrett K, Raymer D, Samudrala S. The design of a new policy model to support ontology-driven reasoning for autonomic networking. *Journal of Network and Systems Management* 2009; **17**(1):5–32, doi:10.1007/s10922-009-9119-3.
20. Famaey J, Latré S, Strassner J, De Turck F. Semantic context dissemination and service matchmaking in future network management. *International Journal of Network Management* 2011; doi:10.1002/nem.805.
21. Martin D, Burstein M, McDermott D, McIlraith S, Paolucci M, Sycara K, McGuinness DL, Sirin E, Srinivasan N. Bringing semantics to web services with OWL-S. *World Wide Web* 2007; **10**(3):243–277, doi:10.1007/s11280-007-0033-x.
22. Kalfoglou Y, Schorlemmer M. Ontology mapping: The state of the art. *The Knowledge Engineering Review* 2003; **18**(1):1–31, doi:10.1017/S0269888903000651.

23. Choi N, Song IY, Han H. A survey on ontology mapping. *ACM SIGMOD Record* 2006; **35**(3):34–41, doi: 10.1145/1168092.1168097.
24. Mitra P, Wiederhold G. Resolving terminological heterogeneity in ontologies. *15th European Conference on Artificial Intelligence (ECAI)*, 2002.
25. Doan A, Madhavan J, Dhamankar R, Domingos P, Halevy A. Learning to match ontologies on the Semantic Web. *International Journal on Very Large Data Bases* 2003; **12**(4):303–319, doi:10.1007/s00778-003-0104-2.
26. Ehrig M, Staab S. QOM – quick ontology mapping. *Third International Semantic Web Conference (ISWC)*, 2004; 683–697, doi:10.1007/978-3-540-30475-3_47.
27. Besana P, Robertson D, Rovatsos M. Exploiting interaction contexts in P2P ontology mapping. *2nd International Workshop on Peer to Peer Knowledge Management*, 2005.
28. Strassner J, Ó'Foghlú M, Donnelly W, Agoulmine N. Beyond the knowledge plane: An inference plane to support the next generation Internet. *First International Global Information Infrastructure Symposium (GIIS)*, 2007; 112–119, doi:10.1109/GIIS.2007.4404176.
29. Serrano M, Strassner J, Ó'Foghlú M. A formal approach for the inference plane supporting integrated management tasks in the Future Internet. *IFIP/IEEE International Symposium on Integrated Network Management (IM)*, 2009; 120–127, doi:10.1109/INMW.2009.5195947.
30. Clark DD, Partridge C, Ramming CJ, Wroclawski JT. A knowledge plae for the Internet. *ACM Conference on Applications, Technologies, Architectures and Protocols for Computer Communications (SIGCOMM)*, 2003, doi: 10.1145/863955.863957.
31. Wong AKY, Ray P, Parameswaran N, Strassner J. Ontology mapping for the interoperability problem in network management. *IEEE Journal on Selected Areas in Communications* 2005; **23**(10):2058–2068, doi:10.1109/JSAC.2005.854130.
32. Brennan R, Feeney K, Walsh B, Thomas H, O'Sullivan D. Explicit federal relationship management to support semantic integration. *1st IFIP/IEEE Workshop on Managing Federations and Cooperative Management*, 2011; 1148–1155, doi:10.1109/INM.2011.5990575.
33. O'Sullivan D, Wade V, Lewis D. Understanding as we roam. *IEEE Internet Computing* 2007; **11**(2):26–33, doi:10.1109/MIC.2007.50.
34. Keeney J, Lewis D, O'Sullivan D, Roelens A, Wade V, Boran A, Richardson R. Runtime semantic interoperability for gathering ontology-based network context. *10th IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2006; 56–65, doi:10.1109/NOMS.2006.1687538.
35. Ludwig H, Keller A, Dan A, King RP, Franck R. Web service level agreement (WSLA) language specification. http://www.research.ibm.com/wsla/WSLASpecV1-20030128.pdf 2003.
36. Andrieux A, Czajkowski K, Dan A, Keahey K, Ludwig H, Nakata T, Pruyne J, Rofrano J, Tuecke S, Xu M. Web services agreement specification (WS-Agreement). http://www.ogf.org/documents/GFD.193.pdf 2011.
37. Battré D, Brazier F, Clark K, Oey M, Papaspyrou A, Wieder P, Ziegler W. WS-Agreement negotiation version 1.0. http://www.ogf.org/documents/GFD.192.pdf 2011.
38. Hudert S, Ludwig H, Wirtz G. Negotiating SLAs – an approach for a generic negotiation framework for WS-Agreement. *Journal of Grid Computing* 2009; **7**(2):225–246, doi:10.1007/s10723-009-9118-3.
39. Hasselmeyer P, Mersch H, Koller B, Quyen HN, Schubert L, Wieder P. Implementing an SLA negotiation framework. *Expanding the Knowledge Economy: Issues, Applications, Case Studies (eChallenges)*, 2007; 154–161.
40. Parkin M, Hasselmeyer P, Koller B, Wieder P. An SLA re-negotiation protocol. *2nd Non Functional Properties and Service Level Agreements in Service Oriented Computing Workshop*, 2008.
41. Yan J, Kowalczyk R, Lin J, Chhetri MB, Goh SK, Zhang J. Autonomous service level agreement negotiation for service composition provision. *Future Generation Computer Systems* 2007; **23**:748–759, doi:10.1016/j.future.2007.02.004.
42. Foundation for Intelligent Physical Agents (FIPA). Fipa iterated contract net interaction protocol specification. http://www.fipa.org/specs/fipa00030/SC00030H.pdf 2002.
43. Rubio-Loyola J, Merida-Campos C, Willmott S, Astorga A, Serrat J, Galis A. Service coalitions for future internet services. *IEEE International Conference on Communications (ICC)*, 2009, doi:10.1109/ICC.2009.5199454.
44. Chai WK, Galis A, Charalambides M, Pavlou G. Federation of Future Internet networks. *IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2010; 209–216, doi:10.1109/NOMSW.2010.5486573.
45. Ranjan R, Harwood A, Buyya R. Peer-to-peer-based resource discovery in global grids: A tutorial. *IEEE Communications Surveys* 2008; **10**(2):6–33, doi:10.1109/COMST.2008.4564477.
46. Iamnitchi A, Foster I, Nurmi D. A peer-to-peer approach to resource location in grid environments. *11th IEEE International Symposium on High Performance Distributed Computing (HPDC)*, 2002; 419, doi:10.1109/HPDC.2002.1029949.
47. Trunfio P, Talia D, Papadakis H, Fragopoulou P, Mordacchini M, Pennanen M, Popov K, Vlassov V, Haridi S. Peer-to-peer resource discovery in grids: Models and systems. *Future Generation Computer Systems* 2007; **23**(7):864–878, doi:10.1016/j.future.2006.12.003.
48. Basu S, Banerjee S, Sharma P, Lee SJ. NodeWiz: peer-to-peer resource discovery for grids. *IEEE International Symposium on Cluster Computing and the Grid (CCGrid)*, 2005; 213–220, doi:10.1109/CCGRID.2005.1558557.
49. Rowstron A, Druschel P. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. *IFIP/ACM International Conference on Distributed Systems and Platforms*, 2001; 329–350, doi: 10.1007/3-540-45518-3_18.
50. Stoica I, Morris P, Liben-Nowell D, Karger D, Kaashoek M, Dabek F, Balakrishnan H. Chord: a scalable peer-to-peer lookup protocol for internet applications. *IEEE/ACM Transactions on Networking* 2003; **11**(1):17–32, doi:10.1109/TNET.2002.808407.
51. Heine F, Hovestadt M, Kao O. Towards ontology-driven P2P grid resource discovery. *Fifth IEEE/ACM International Workshop on Grid Computing*, 2004; 76–83, doi:10.1109/GRID.2004.61.

52. Pipan G. Use of the TRIPOD overlay network for resource discovery. *Future Generation Computer Systems* 2010; **26**(8):1257–1270, doi:10.1016/j.future.2010.02.002.
53. Shen G, Huang Z, Zhang Y, Zhu X, Yang J. A semantic model for matchmaking of web services based on description logics. *Fundamenta Informaticae* 2009; **96**(1):211–226, doi:10.3233/FI-2009-175.
54. Paolucci M, Kawamura T, Payne TR, Sycara KP. Semantic matching of web services capabilities. *First International Semantic Web Conference (ISWC)*, 2002; 333–347, doi:10.1007/3-540-48005-6_26.
55. Bener AB, Ozadali V, Ilhan ES. Semantic matchmaker with precondition and effect matching using SWRL. *Expert Systems and Applications* 2009; **36**(5):9371–9377, doi:10.1016/j.eswa.2009.01.010.
56. Sbodio ML, Martin D, Moulin C. Discovering semantic web services using SPARQL and intelligent agents. *Web Semantics: Science, Services and Agents on the World Wide Web* 2010; **8**(4):310–328, doi:10.1016/j.websem.2010.05.002.
57. Islam R, Islam Z, Leyla N. A tree-based approach to matchmaking algorithms for resource discovery. *International Journal of Network Management* 2008; **48**(5):427–436, doi:10.1002/nem.686.
58. Kumar N, Saraph G. End-to-end QoS in interdomain routing. *International Conference on Networking and Services (ICNS)*, 2006; 82–88, doi:10.1109/ICNS.2006.45.
59. Saraph G, Singh P. New scheme for IP routing and traffic engineering. *Workshop on High Performance Switching and Routing (HPSR)*, 2003; 227–232, doi:10.1109/HPSR.2003.1226709.
60. Pouyllau H, Douville R. End-to-end QoS negotiation in network federations. *12th IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2010; 173–176, doi:10.1109/NOMSW.2010.5486578.
61. Pouyllau H, Carofiglio G. Inter-carrier SLA negotiation using Q-learning. *Telecommunication Systems* 2011; doi:10.1007/s11235-011-9505-5.
62. Famaey J, Latré S, Wauters T, De Turck F. FedRR: A federated resource reservation algorithm for multimedia services. *13th IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2012.
63. Jones D, Keeney J, Lewis D, O'Sullivan D. Knowledge-based networking. *Second International Conference on Distributed Event-Based Systems*, 2008, doi:10.1145/1385989.1386034.
64. Carzaniga A, Rosenblum DS, Wolf AL. Design and evaluation of a wide-area event notification service. *ACM Transactions on Computer Systems* 2001; **19**(3), doi:10.1145/380749.380767.
65. Keeney J, Roblek D, Jones D, Lewis D, O'Sullivan D. Extending siena to support more expressive and flexible subscriptions. *Second International Conference on Distributed Event-Based Systems (DEBS)*, 2008; 35–46, doi:10.1145/1385989.1385995.
66. Wang J, Jin B, Li J, Shao D. A semantic-aware publish/subscribe system with RDF patterns. *28th Annual International Computer Software and Applications Conference (COMPSAC)*, 2004; 141–146, doi:10.1109/CMPSAC.2004.1342818.
67. Petrovic M, Liu H, Jacobsen HA. G-ToPSS: Fast filtering of graph-based metadata. *14th international conference on World Wide Web (WWW)*, 2005; 539–547, doi:10.1145/1060745.1060824.
68. Li H, Jiang G. Semantic message oriented middleware for publish/subscribe networks. *Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense III*, vol. 5403, 2004; 124–133, doi:10.1117/12.548172.
69. Latré S, van der Meer S, De Turck F, Strassner J, Won-Ki Hong J. Ontological generation of filter rules for context exchange in autonomic multimedia networks. *12th IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2010; 575–582, doi:10.1109/NOMS.2010.5488448.
70. Barron J, Davy S, Jennings B, Strassner J. A policy authoring process and den-ng model extension for federation governance. *5th International Workshop on Modelling Autonomic Communication Environments (MACE)*, 2010; 73–86, doi:10.1007/978-3-642-16836-9_7.
71. Barron J, Davy S, Jennings B. Conflict analysis during authoring of management policies for federations. *1st IFIP/IEEE Workshop on Managing Federations and Cooperative Management*, 2011; 1180–1187, doi:10.1109/INM.2011.5990579.