Chapter 5

Federating Autonomic Network Management Systems for Flexible Control of End-to-End Communications Services

Brendan Jennings, Kevin Chekov Feeney, Rob Brennan, Sasitharan Balasubramaniam, Dmitri Botvich, and Sven van der Meer

[1] *The work described in this chapter is supported by Science Foundation Ireland, via the "Federated, Autonomic Management of End-to-end Communications Dervices (FAME)" strategic research cluster, grant no. 08/SRC/I1403. For a description of the FAME cluster see.* AQ:1

INTRODUCTION

Over the past decade the research community has been actively investigating how network management systems can be developed to achieve the flexibility and adaptability required to efficiently operate increasingly complex, heterogeneous, and interconnected networks. In the forefront of these efforts has been the vision of *autonomic network management* [1] in which a network has the ability to self-govern its behavior within the constraints of human-specified business goals. Autonomic network management promises a much more flexible approach to management that seeks to both allow systems to automatically adapt offered services or resources in response to user or environmental changes and to reduce operational expenditure for network operators.

While significant progress has been made, work to date has focused on autonomic management in the context of single, well-defined network domains. Relatively little work has been done on how autonomic management systems can be *federated* across management domains to provide end-to-end management of communications services. Current network management systems do perform some coordination across domains, but this is limited in scope to a small number of predefined information exchanges. In this chapter we argue that autonomic network management systems should be designed to incorporate capabilities supporting the negotiation and life-cycle management of federations involving two or more networks, service providers, and service consumers.

Autonomic Network Management Principles. DOI: 10.1016/B978-0-12-382190-4-4.00005-X

We first briefly review previous work on autonomic network management, focusing on aspects relating to federation. We then discuss our view of federation, introducing a layered federation model. Next we discuss the challenges that must be addressed to achieve federation of networks, their management systems, and the organizations/individuals that operate them. Finally, we discuss an example scenario based on end-to-end management of the delivery of Internet Protocol-based Television (IPTV) content that illustrates the benefits to be gained from adopting a federated autonomic management approach.

AUTONOMIC NETWORK MANAGEMENT: AVOIDING NEW MANAGEMENT SILOS

The term *autonomic computing* was coined by IBM as an analogy of the autonomic nervous system, which maintains homeostasis (which means essentially maintaining the equilibrium of various biological processes) in our bodies without the need for conscious direction [2]. Autonomic computing attempts to manage the operation of individual pieces of IT infrastructure (such as servers in a data center) through introduction of an autonomic manager that implements an autonomic control loop in which the managed element and the environment in which it operates are monitored, collected data is analyzed, and actions are taken if the managed element is deemed to be in an undesirable (suboptimal) state. The IBM *MAPE-K* control loop is made up of *Monitor*, *Analyze*, *Plan*, and *Execute* components, all of which rely on a common *Knowledge* repository.

The autonomic computing vision can be summarized as that of a "self-managing" IT infrastructure in which equipment will have software deployed on it that enables it to self-configure, self-optimize, self-heal, and self-protect. That is, it will exhibit what has come to be known as "self-*" behavior. Clearly this is a powerful vision, and it was therefore natural that the networking community would extend this vision from autonomic management of individual elements to autonomic networking—the collective (self-) management of networks of communicating computing elements. As surveyed by Dobson et al. [3], autonomic networking is a burgeoning research area, integrating results from disciplines ranging from telecommunications network management to artificial intelligence and from biology to sociology.

From the network management perspective, one of the successors to the IBM MAPE-K architecture is the *FOCALE* architecture [1], which seeks to address the particular challenges of managing communications network devices. FOCALE implements two control loops: A *maintenance loop* is used when no anomalies are found (i.e., when either the current system state is equal to the desired state, or when the state of the managed element is moving toward its intended goal); and an *adjustment loop* is used when one or more policy reconfiguration actions must be performed and/or new policies must be codified and deployed. The adjustment loop is supported by a dynamically updateable knowledge base—one that can be modifed to reflect new knowledge at runtime as new knowledge is discovered by reasoning and machine learning components. Furthermore, FOCALE is designed as a distributed architecture in which autonomic control loops can be embedded within autonomic managers, control single or groups of managed elements, and in which semantic models and model-based translation are used to support a common management *lingua franca* that can be used to normalize monitoring information and configuration commands.

Besides FOCALE, other proposed autonomic networking architectures have addressed some aspects of coordination of management activities between networked devices. For example the *Meta-Management System* (MMS) [4] provides robust management-plane communications incorporating self-configuration, self-healing, self-optimization, and self-protection capabilities. Another example is the *Autonomic Network Architecture* (ANA) [5], which aims to provide genericity and flexibility at all levels of an autonomic network AQ:2 by supporting the coexistance of different network styles and communications protocols. While efforts toward development of effective distributed autonomic management systems promise significant benefits, we believe there is a need for a more holistic approach in which communication and coordination of management activities between devices is governed by organizational federations that mandate management systems and autonomic managers to cooperate to achieve business goals.

OUR VIEW OF FEDERATION

We employ *federation* as a general term for describing cross-organizational capability sharing. More specifically, we define a federation as a persistent organizational agreement that enables multiple autonomous entities to share capabilities in a controlled way.

We make a number of observations regarding the implications of this definition. First, a federation brings together *autonomous entities*—organizations or individuals endowed with sovereign decision-making power over the resources they own or control. Hence, there is no single governing authority for the federation. Second, given the lack of a single governing authority, the federation must exist by virtue of the *agreement* of its members. If the federation is to be viable, it is important that its nature, structure, and evolution be agreed upon and that the value delivered to federation members by virtue of their participation be transparent and clear to each of them.

The third observation is that federations exist in order to enable the *controlled sharing of capabilities* between autonomous entities. We use the term *capability* in the widest possible sense It could range from the availability of a communication channel to the ability to perform device configuration changes. Controlled capability sharing means that federation members are granted (possibly constrained) access to capabilities they would not otherwise possess. Finally, we observe that federations should be *persistent*—which is not meant to imply

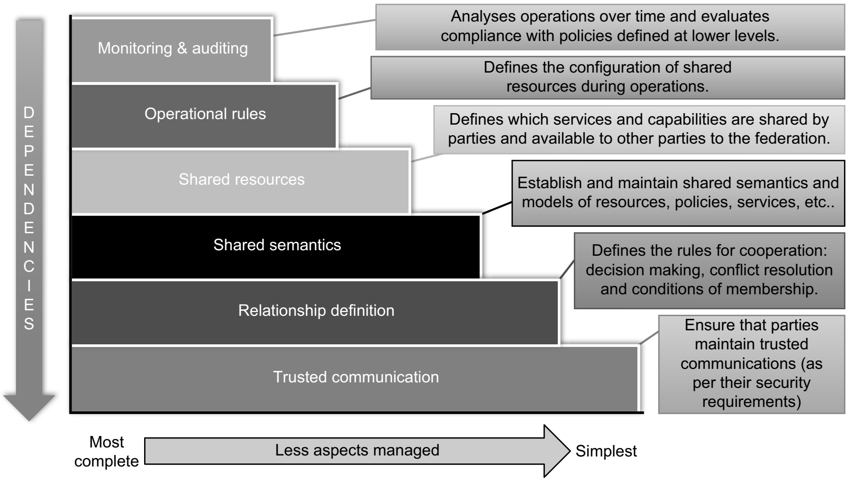
Monitoring & auditing
Analyses operations over time and evaluates compliance with policies defined at lower levels.
Operational rules
Defines the configuration of shared resources during operations.
Shared resources
Defines which services and capabilities are shared by parties and available to other parties to the federation.
Shared semantics
Establish and maintain shared semantics and models of resources, policies, services, etc..
Relationship definition
Defines the rules for cooperation: decision making, conflict resolution and conditions of membership.
Trusted communication
Ensure that parties maintain trusted communications (as per their security requirements)
DEPENDENCIES
Most complete
Less aspects managed
Simplest

FIGURE 5.1 Layered Federation Model (LFM).

that they are permanent or must last any minimum time period. They should, however, outlive individual transactions or interactions between federation members and will thus have an associated *federation life cycle*.

Following from the above discussion of our understanding of federations, we now introduce the *layered federal model* (LFM), depicted in Figure 5.1. The LFM is intended as a general-purpose high-level conceptual model of the components of a federal agreement. The model is decomposed into layers, with each layer representing particular aspects of a federation agreement on the organizational arrangement. Its main purpose is to serve as a useful model for the decomposition of federal relationships in order to render their definition and maintenance more tractable and transparent. The layers represent the most important aspects of successful persistent, cross-organizational relationships, with the relative positioning of these aspects representing the dependencies between the various elements that constitute a federal agreement. We now briefly review the six layers of the layered federation model.

- Trusted Communication Layer

  A basic requirement for any sort of communication between autonomous entities is a channel with sufficient security measures to satisfy both parties' requirements for the current dialogue. This requires them to agree on communications protocols, security mechanisms, and even applications that can deliver appropriate facilities for identification, authentication, privacy, and accountability. These concerns form the most fundamental layer of our federation model because all higher level agreements and interactions depend on them;
- Federal Relationship Definition Layer

  The relationship definition layer supports the definition and transmission of the basic rules that govern each organization's relationships with other organizations within the federation. This provides a generic layer in which rules concerning membership of the federation and sharing of capabilities (and their revocation) can be negotiated and agreed. For example, to support a peering relationship organizations may create a federation requiring their networks to carry traffic that originated in the networks of other federation members;

- Shared Semantic Layer

  Federations, as we define them, exist to allow autonomous organizations to share capabilities. However, any particular organization will generally have its own addressing mechanisms and semantics for describing the resources that it controls and the capabilities that they support. The shared semantic layer provides mappings between the semantics used internally by each federation member to describe their resources and capabilities to those used by the other parties. This could be achieved by means of a standardized federal semantic language, or it could be achieved by each party mapping directly between their internal semantics and those of the other parties to the federation.
- Shared Capabilities Layer

  Having established sufficiently secure communications, a general resource-sharing regime and shared semantics with respect to resources, the prerequisites are in place to allow capabilities to be shared. The capability-sharing layer is concerned with enabling members of federations to manage the dynamic set of capabilities they share. This includes providing a means whereby members of the federation can add and remove capabilities from the pool available to other federation members, as well as allowing other parties to discover which capabilities are available for use at any particular time.
- Operational Rule Layer

  The operational rule layer augments the capability-sharing layer by providing federation members with the ability to view and/or change the configuration of resources provided by other federation members. For example, a communications service provider may be granted the ability to configure a home user's set-top box in order to provision an IPTV service. Clearly, the AQ:3 degree to which resource configurations can be viewed and modified needs to be fully specified in the federation agreement;
- Auditing and Monitoring Layer

  In many cases, the lower layers of the federal model can adequately manage their own auditing, reporting, and compliance assurance. However, federal agreements may be formulated in such a way that compliance and verification is only possible through observing aggregate activity over significant periods of time. So, for example, a federal agreement might include a clause that specifies that each member should more or less provide as many useful resources to the federation as they use. Due to random variations, traffic spikes, and so on, such agreements can only be meaningfully checked over significant periods of time. The monitoring and auditing layer is thus the top layer of this federal model. It is responsible for providing members of the federation with detailed monitoring of their compliance and that of counterparties to federal agreements.

FEDERATION OF NETWORKS

As noted earlier, the main focus of the literature on autonomic networking has been on the development of algorithms and processes that allow networks to self-organize, self-optimize, and self-protect. As networks grow in complexity, autonomic capabilities such as this are increasingly seen as a necessity. For example, the anticipated deployment of very large numbers of Femto-cells attached to the mobile access networks will require a significant rethinking of traditional centralized operations and management approaches that are not designed to effectively handle management of very large numbers of network elements. In developing these algorithms and processes, researchers have drawn inspiration from myriad sources, notably the self-management behavior exhibited by biological systems. This has resulted in proposals of varying complexity and with varying responsiveness profiles, many of which have been shown to be suitable for deployment in large-scale network domains. Nevertheless, the question of how self-management algorithms and processes deployed in interconnecting network domains can be coordinated to deliver the best possible end-to-end behavior has received little attention.

Current network domains are predominantly managed on an individual basis, with the aim being to optimize the transfer of information within administrative boundaries. Although cordination across network boundaries does exist, its scope is limited primarily to participation in end-to-end routing protocols, network peering arrangements, and exchange of certain management information (in particular for charging and billing purposes). Furthermore, such coordination is highly static in nature, with changes requiring time-consuming negotiations and coordinated deployment by network operators.

We argue that there are significant benefits to be harnessed through adoption of approaches to coordinating the operation of self-management algorithms and processes across network administrative boundaries. The means of achieving this coordination can be either direct or indirect. In the direct approach, self-management processes would use agreed-upon protocols to directly transfer control information that allows them to synchronize their behavior. In the indirect approach, coordination would be effected by coordination of the parameterization of self-management processes by the management systems associated with the individual networks, with cross-domain interaction being handled by the management systems. Both approaches necessitate the availability of some of the facilities provided by the layered federation model, in particular: the presence of trusted communications across network boundaries, the definition of a federal relationship and some degree of shared semantics between the network operators.

FEDERATION OF MANAGEMENT SYSTEMS

As alluded to above, network management systems typically seek to deliver the best performance from a collection of network elements directly under their control. Even within a single network domain, network elements typically conform to different programming models and are controlled via a myriad of management protocols and interfaces. Management systems themselves rely on system models that, though nominally adhering to industry standards, are in reality very proprietary in nature and are based on divergent semantic models. In this context the task of federating at the management system level is very challenging, but, we believe, necessary to deliver end-to-end autonomic management solutions.

A key requirement for autonomic management systems is the ability to monitor their own environment so that they can react to changes and gather information to populate or specialize their own internal environmental models for planning and prediction. To provide end-to-end autonomic management, it is not sufficient for individual management systems to simply monitor their own domains. Instead, to close the federated, end-to-end autonomic control loop it is necessary to deploy processes that provide service-level monitoring of the operational state of network elements across network boundaries. Such processes would need to provide on-demain and continuous monitoring facilities that can AQ:4 be used to both accurately identify the source of unexpected service degradation and support prevention of potential service degradation. Collected data would need to be shared across different administrative boundaries and processed in a distributed, scalable manner that is conducive to analysis and contextualization using the aid of information or semantic models.

Monitoring data harmonization, as envisaged here, goes far beyond syntactical interworking: It requires mechanisms for mapping and federation of the large volume of low-level data produced by monitoring systems into semantically meaningful information and knowledge that may be discretely consumed by value network members. Any complex system has domain experts who are familiar with how the constituent parts of the system can be managed, and they are particularly aware of the end-to-end operating constraints of those constituent parts. Encoding this expertise in a manner that can be utilized by other stakeholders is a difficult task [6], but one we believe can be solved through the application of knowledge representation and engineering techniques. The benefit of enabling such expertise to be encoded, aggregated, and interpreted across diverse domains is that it will allow federation members to monitor other parts of the system, using knowledge gained to inform management decisions made locally. This clearly implies the presence of well-defined federal relationships, sharing of capabilities, and well-formed operational rules and federation monitoring and auditing facilities, as outlined in the layered federation model.

The flipside of end-to-end service level monitoring is how to affect configuration changes in response to sensed changes in a manner that is coordinated across an end-to-end service delivery path. Autonomic management systems seek to analyze monitoring information and use knowledge inferred from it to trigger policy management analysis and decision processes that result in the generation and deployment of sets of device configurations that best align a network's behavior with business goals. However, current network management systems are significantly impaired by their inability to automate the translation of business information to network device configuration commands. Automating the translation of business-level policies (specified in terms of entities such as products, services, and customers) through a number of levels of abstraction into corresponding device instance-level policies (configuration commands specified in terms of entities such as packet marking rules or firewall configuration rules) is hugely challenging, since information needs to be added (and removed) as the policies become more specific in nature. Ensuring the consistency of configurations across a federation is further complicated by the requirement for shared semantics, the need for synchronization, and the presence of appropriate federation operational rules.

FEDERATION OF ORGANIZATIONS AND THEIR CUSTOMERS

To respond to an increasingly deregulated environment, there is a strong impetus for network operators to support more flexible business models. One trend in particular is the proliferation of virtual operators that own little or no traditional infrastructure of their own, so that service delivery is always based on crossing multiple administrative domains and there is never completely in-house service instance handling. Flexibility for providers of carrier networks to deal with multiple virtual providers, perhaps even on a per-service basis, can drive down costs and significantly increase business agility. The benefits for smaller operators are obvious, but even for the big players, support of such flexible infrastructure is becoming a source of competitive advantage as the market for wholesale services continues to grow, driven by the increasing numbers of small operators that use their services. Ultimately, these complex value chains of business relationships also need to be managed themselves. Other trends, including the emergence of prosumers (users who are both producers and consumers of content and services), reinforce the need for operators to explicitly empower their networks to support the creation and management of federations with other network providers, service providers, and their own customers.

To efficiently support shorter service life cycles within dynamic federations of service providers and their customers, autonomic network management systems will need the ability to dynamically negotiate and manage federations with other management systems, while achieving an appropriate balance between satisfying local and federation-wide goals. This will require major extensions to current approaches for expression of business goals, negotiation mechanisms, and distributed security. A particular challenge relates to the difficulty of modeling the differing nature of diverse management approaches. Different priorities will be placed on the value of federation membership by different resources and services, as well as by the managers and/or management systems of those resources and services. This will influence the degree to which actors may wish to exert their management authority at any one time or to delegate it to other actors in the federation.

Increased management flexibility and extensibility can be provided by policy-based management (PBM) and the semantic modeling of management information [1]. PBM is an increasingly popular method for combining flexibility with efficiency in systems and network administration [7]. In PBM systems, administrators encode operational management decisions as rules, which are then mapped into concrete device configurations by the policy system. However, structural abstractions, such as roles and domains, used in many policy rule languages, reflect hierarchical organizational thinking that has been shown to be insufficient for modeling the complex interrelationships between individuals within human organizations [8]. This results in the need for complex tools to manage changes to these policy rule languages, which themselves become obstacles to interoperability and to change. Therefore, handling the authoring and maintenance of a coherent set of policies across a federation of providers presents significant problems for existing PBM systems with regards to policy conflict resolution and collaborative policy consensus forming.

There are significant open issues surrounding how federations are formed, what level of trust is required to support this formation, and what patterns of decentralized authority would best support different forms of federations or value networks. Securing a federation is also challenging, especially as service refinement and innovation may spawn subfederations of providers and customers. In order to ensure that disparate policies are consistently enforced across federations, most existing approaches rely on centralized administrative authority. We believe that noncentralized approaches, such as the distributed authorization language (DAL) [9]—which is designed specifically for decentralized coalition policies and supports dynamic coalition formation, merging, and delegation—will be required to deliver the powerful federation management facilities envisaged for the Operational Rule and Auditing and Monitoring layers of the layered federation model.

Many problems in dynamically federating independent management domains will not be resolved simply by resolving rule (or policy) conflicts because there is no underlying logic for resolving the conflict and, as importantly, verifying that the conflict has been resolved. Rather, such problems will require novel advances in semantic analysis, whereby heterogeneous representations of knowledge describing the state and behavior of managed resources can be harmonized and decisions enforced using unambiguous policy rules. The use of formal semantics, including mappings, in structured knowledge has already been demonstrated for autonomic networking approaches [10, 11]. Using these techniques it is possible to model the resources being managed, the context used in communications services, and the policies used to express governance. For a federated environment, semantic mapping is essential to allow the sharing and common understanding of contracts and policies, which together govern the interactions between different management elements.

We note that traditional semantic mapping approaches, as surveyed in [12], generally assume that the mapping task is performed by a knowledge engineer, whose task is to generate a full static mapping between models that will be used by several applications. This is a constrained and static view of the semantic AQ:5 mapping process that would not deliver the on the requirement for development of a full methodology for integration of disparate knowledge as envisaged in the Shared Semantics layer of our layered federation model. Furthermore, in the context of the LFM Shared Semantics layer, mappings in support of federations will need to be generated to be task-specific and context-sensitive, be able to represent partial knowledge, and need to be tracked, managed, and maintained over time [13].

EXAMPLE SCENARIO: END-TO-END MANAGEMENT OF IPTV SERVICES

We now explore a potential deployment scenario for flexibly supporting federations of networks, their management systems, the organizations that control them, and the users of their services. The scenario focuses on end-to-end delivery of Internet Protocol-based Television (IPTV) content from IPTV service providers' data centers to devices located in users' Home Area Networks (HANs). Deployment of IPTV is a major growth area in the telecommunications market at present, with many network operators seeing IPTV services as the next major driver of revenue. However, large-scale deployment of IPTV presents significant technical challenges. For television services, users have expectations of high levels of Quality of Service (QoS) and Quality of Experience (QoE), however, it is very difficult to ensure that sufficient bandwidth is available across an end-to-end path spanning multiple management domains and heterogeneous device types.

Today, IPTV service provision typically only takes place within a single network operator's domain or between an operator's domain and the domain of a closely associated third-party service provider. A typical current deployment is described by Hu et al. [14], who outline how IPTV services are delivered in Chunghwa Telecom's Taiwanese network. The advantage for operators of controlling all aspects of IPTV service provision is that ownership of the access network to the customer premises gives them the fine-grained control of network capabilities required to provide the necessary QoS/QoE guarantees for real-time media streaming. Furthermore, direct access to configure set-top boxes in customer premises allows them support for enhanced features such as live program guides and interactive, value-added services. Although this provides a workable solution for today's deployments, it is a largely static and tightly coupled solution, necessitated in part by the lack of flexibliity of present management systems, particularly with respect to their marginal, if any, support for interdomain capability sharing.

We argue that a more dynamic service provider federation-based approach would enable the customer to pick best of breed third-party service providers and also allow the network operator to partner with a varied array of third parties. To illustrate how this could be achieved, we will now outline how a set of autonomic management techniques could be applied together to provide some of the key functionalities of a federated management solution for IPTV deployments.

Coordinated Self-Management for IPTV Content Distribution

As discussed at the start of this chapter, much of the focus of autonomic network management research has been on the development of self-management algorithms and processes. One of the key challenges is how to coordinate the behavior of self-management processes operating either at different layers within a given network or across network boundaries. In this section we outline two self-management processes: The first is a gradient-based routing process operating in the core network, while the second is a content management process operating at the IPTV service level. The two processes are designed to operate *symbiotically* to provide effective end-to-end delivery of IPTV content to those parts of a network where demand is high. This symbiotic relationship is implicit rather than explicit in that management systems are not required to communicate to ascertain the appropriate parameterizations of processes to allow them to cooperate efficiently. The particular IPTV deployment we consider in this chapter is depicted in Figure 5.2. It includes a single network provider, two independent IPTV service providers, and a number of end-user home area networks (for simplicity, only one is shown in the figure).

The Gradient-Based Routing (GBR) process is a bio-inspired process that, in contrast to existing routing processes, is highly distributed and dynamic, allowing routes to be discovered on a hop-by-hop basis as traffic flows through a network. Central to the process is a gradient solution that mimics the chemotaxis motility process of microorganisms, whereby a microorganism senses the chemical gradient toward a food source and migrates by following the increasing gradient path. In an analogous manner, the GBR process creates a gradient field across the network topology, which is then used to ascertain the best routes between source and destination network nodes. A major strength of the process is that as the load patterns in the network change, so too will the gradient field. This results in efficient use of network bandwidth and increased survivability in the face of failure events. More information on the approach can be found in Balasubramaniam et al. [15].

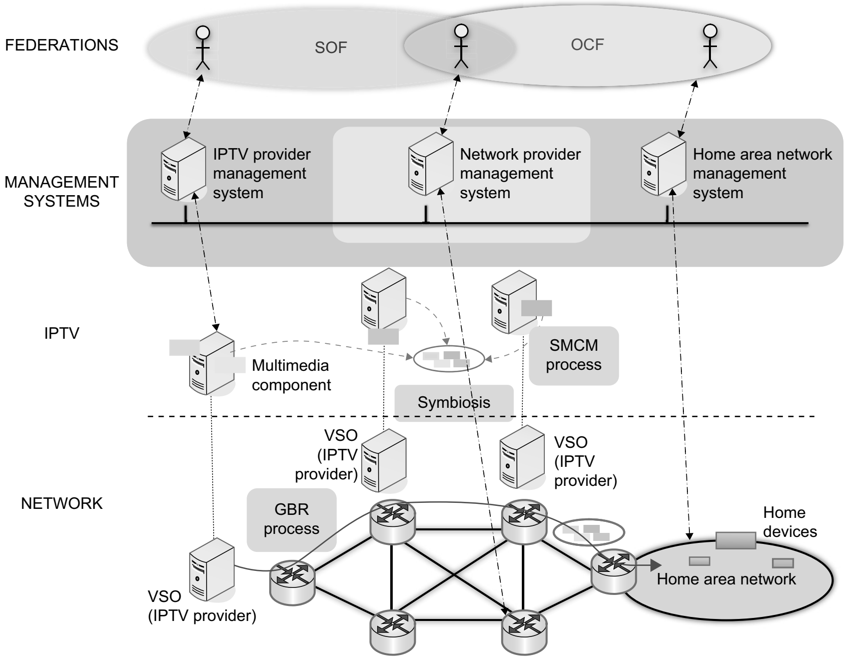

FIGURE 5.2 Federated IPTV Use Case.

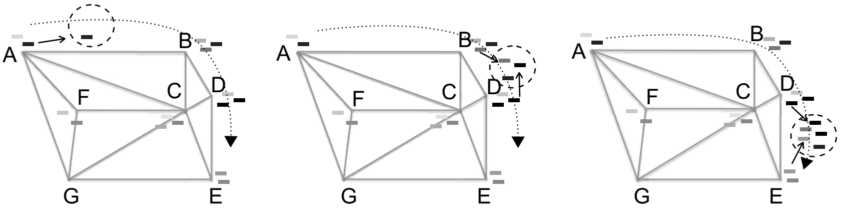

AQ:6

FIGURE 5.3 Example Slime-Mold Inspired Component Composition.

The second process we outline operates at the IPTV service layer; it controls how items of IPTV content are distributed around a set of content servers attached to different network end points and all controlled by a single IPTV service provider. Once again the process is bio-inspired, this time by the behavior of slime molds—organisms that have funguslike properties and that usually live on dead plant material (see Figure 5.3). A typical property of slime molds is their ability to coordinate among themselves and move as a unit in the event of unfavorable environmental changes, in particular the scarcity of food. An example is Dictyostelium cells, which, when starved, are able to develop into agar that forms fruiting bodies. These fruiting bodies mobilize as a unit and move to a new environment, where, upon arrival, they revert into individual cells.

Our Slime Mold Content Management (SMCM) process applies an analogous process to enhance the management of IPTV content distribution. As the market for IPTV grows, we can expect high levels of volatility in demand for particular pieces of content. Peer-to-peer networks partially address this scenario by segmenting large multimedia content items (such as movies) into short segments that are then distributed around the network. Peer-to-peer techniques are subsequently used to concatenate and stream the content to the end user. This results in significantly improved bandwidth utilization profiles as compared to the traditional approach of having individual flows for each content streaming session. However, such peer-to-peer approaches generally are not aware of the parts of the network experiencing high demand for particular content items or of the prevailing load of different parts of the network. Hence, there is no guarantee that content is distributed optimally from the perspective of the operation of the network.

The SMCM process takes into account the location of user demands and migrates content toward content servers closest to those parts of the networks where those users are attached. As illustrated in Figure 5.3, based on user demands that are diffused and propagated between the nodes, content segments will "absorb" these demands and coordinate and form agarlike units and migrate toward the destination closed to those users' point of attachment. As shown in the figure, as the unit is formed and moves from node, it will pick up segments that are relevant to other segments within the unit. This process will lead to a distributed migration of content items toward parts of the networks having high demand for popular content. This has the effect of lessening the load on the network in comparison to the case where content does not migrate, which makes the task of the GBR process less onerous. Similarly, the operation of the GBR maximizes available bandwidth between content servers and users' points of attachment, which increases the capacity of the IPTV service provider to deliver service to its customers.

As described above, the GBR and SMCM processes implicitly aid each other's management task, providing a form of *mutualistic symbiosis*. This symbiotic relationship could be further enhanced by enabling explicit coordination of the processes. For example, a network provider could provide information relating to the status of the network topology (for example, regions currently experiencing congestion), which the SMCM process could use to divert content migration around congested regions. Similarly, a HAN could provide information relating to user preferences and requirements to the IPTV service provider that could be used to configure how IPTV content items are adapted for consumption by particular users or groups of users. In the next section we explore how network management systems could effect this form of explicit coordination.

Federating Network and IPTV Provider Management Systems

Closer cooperation between management systems is key to realizing flexible federated management. In this context cooperation encompasses the exchange of monitoring and configuration data and possibly the delegation of authority to manage resources. We believe that this form of cooperation is most readily achievable where management systems are based on the PBM paradigm. In previous work we [16] addressed the concept of a *Policy Continuum*, in which policies at different levels of abstraction are linked in order to allow high-level goals expressed using business concepts (e.g., customer, service) to be translated into appropriate low-level device configurations (e.g., CLI-based configuration of router traffic shaping and queuing). Implementation of the policy continuum enables these constituencies, who understand different concepts and use different terminologies, to manipulate sets of policy representations at a view appropriate to them and to have those view-specific representations mapped to equivalent representations at views appropriate for manipulation by other constituencies. If, as illustrated in Figure 5.4, federating management systems each realize a policy continuum, then the governance of the federation can be effected by a set of "federation-level" policies.

Federation-level policies would be formulated to allow federation participants access to information provided by other participants and/or allow participants to appropriately configure resources in other participants' networks. In previous work [15] we outlined how the creation of a DEN-ng model of the GBR process facilitates its reparameterization by policies. As an example we showed how the GBR could be reparameterized to prioritize optimal load

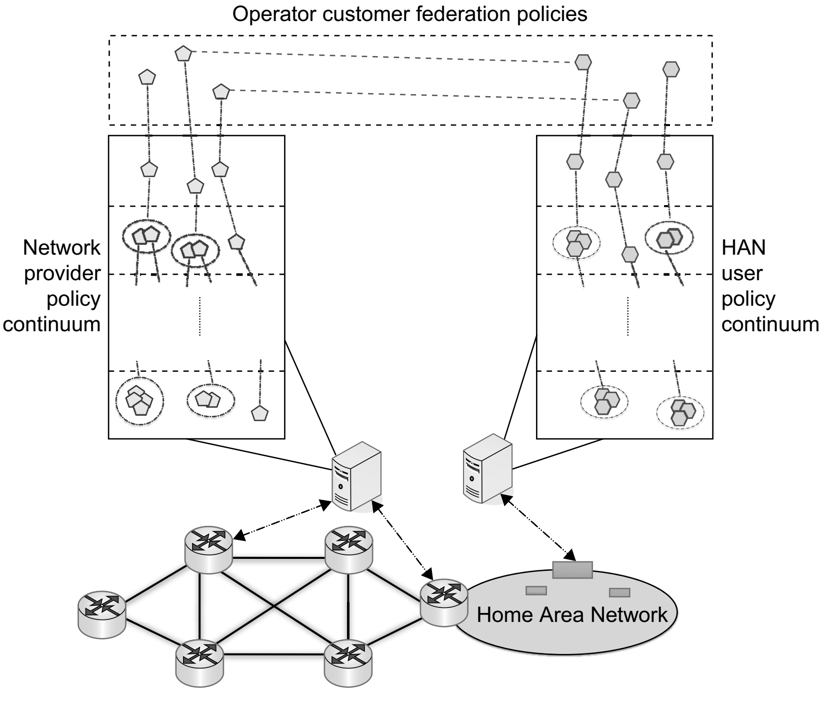
Operator customer federation policies
Network provider policy continuum
HAN user policy continuum
Home Area Network

FIGURE 5.4 Operator Customer Federation Policies.

balancing over achieving shortest paths following a link failure event. Similarly, in our IPTV federated management system scenario, a network management system could be allowed to reconfigure GBR parameters in a neighboring network to prioritize end-to-end load balancing when a significant failure occurs in its own network, the objective being to deliver the best possible end-to-end QoS to IPTV service users. Similarly, as discussed above, a network provider management system could have a policy that triggers sharing of network topology and link load information with the management system of an IPTV service provider once loads in parts of the network go beyond a specified threshold. This information would allow the IPTV service provider management system to reparameterize its SMCM process to migrate content away from highly loaded regions of the underlying network.

Realization of management system federation at the policy level will require sophisticated processes for authoring and analyzing policies. Of particular importance is ensuring that deployed policies do not conflict and are consistent with both an organization's own goals and the commitments it makes as part of a federation with other organizations. Therefore, any proposed introduction, withdrawal, or modification of one or more local or federation-level policies must be assessed on this basis. Moreover, as we discuss in the next section, facilities are required to manage the life cycle of the federations themselves.

Interprovider and Provider-User Federations for IPTV Services

As described earlier, the LFM requires the presence of trusted communication and the definition of the federal relationship between participants. There is a substantial body of work on providing functionality of this type (e.g., see [17], which we will not discuss here. Assuming the presence of such facilities, we will now explore the provision of the other aspects of the LFM in the context of our IPTV service delivery scenario.

As illustrated in Figure 5.2, we envisage the deployment of *Federal Relationship Managers (FRMs)* that augment the already deployed management system. For a core network, the FRM may be implemented as a component of the traditional Operations Support Sytem (OSS), while from a HAN the FRM functionality may be integrated as a domain controller in a set-top box. The FRMs act as interconnectors between the management systems of the various organizations participating in a federation, in essence providing the functionality described in the Shared Semantics, Shared Capabilites, and Operational Rules layers of the LFM. They thus enable the establishment of semantic interoperability for the type(s) of integration/interworking and supporting capabilities currently desired by the federation members and enable the identification, publication, discovery, and description of specific capabilities as shared within the federation by individual federation members.

For our IPTV service delivery scenario we address two types of federation: that between the network operator and its customers (OCF) and that between service providers and network operators (SOF). The OCF provides a simple means by which the customer can delegate management authority over the capabilities in the HAN, on a case-by-case basis, to their network provider in the interest of receiving improved customer support, an increased range of services, and ease of use for their HAN. The SOF allows the network operator to make the capabilities shared between the OF available to service providers, including those representing the capabilities that exist in the HANs of individual subscribers. These federations provide the ability to optimize the delivery of IPTV content items by providing a controlled mechanism by which users can allow service providers access to information required to adapt content to that user's device constraints and preferences. Similarly, the network operator and service provider can allow each other access to the information required to maximize the symbiotic relationship between the GBR and SMCM self-management processes.

We view the FRM as an interconnector between existing management systems and existing semantic spaces, rather than a universal model that must be applied across the entire network of relationships. It encapsulates the sum-total of common technical infrastructure that an organization must adopt in order to manage and maintain an arbitrarily complex set of federal relationships. It does not mandate any particular policy language, information model, or management structures or process across a set of federal relationships. In our ongoing work to create FRM prototypes we are utilizing outputs of previous work on an ontology mapping framework [13] and on the *Community Based Policy Management System (CBMPS)* [18]. We now briefly describe these technologies and their application in the context of an FRM.

The CBPMS is a general-purpose policy management framework designed to be policy-language and information-model neutral. Rather than focusing on the specific semantics of policies or resources, it provides a flexible and secure authority management capability. The CBPMS supports decentralized management through delegation of capability authorities. Capability authorities are references to nodes on a resource-authority tree, and this tree is implemented as a service that can be deployed by the owner of any resource that is to be shared. What makes the CBPMS particularly suitable for application in an FRM is that these capability authorities are higher-level constructs than are permissions—the standard unit of most access control and management policy systems. This allows, for example, a telecommunications service provider to grant access to an IPTV service provider to all of their customers (or whatever subset they require) via a single delegation of a capability authority rather than having to specify individual permissions for each user, which is impractical on such a scale. In the CBPMS schema, policies are also associated with capability authorities, which serve as filters on policy searches. When compared to the standard role-based approach to policy management, this yields considerable gains in terms of the size of the policy search space that must be traversed for each policy decision. In large, complex, end-to-end service delivery scenarios, such gains are extremely important, since services such as IPTV have relatively low tolerance for delays in establishing the connections and even lower tolerance for latencies and jitter.

The ontology mapping framework we apply is an extended version of that described in [13] which spans ontology mapping creation, through use and reuse, to evolution and management. The goal of applying this famework within the FRM is to enable the effective and efficient creation and management of mappings to increase understanding of shared capabilities across the federation. For example, in our IPTV scenario this would include the capabilities shared within federations (typically network or service resources) and the federation-level policies used to express governance over the capabilities. Current ontology mapping approaches can be characterized as follows: "knowledge engineers" engage in "one-shot" processes that result in static "one size fits all" mappings, which are then published for indiscriminate use. In contrast, our ontology mapping process is designed to: (a) cope with the diversity of actors involved in managing a federation (i.e., not always specialist knowledge engineers with specialist tools); b) allow for the diversity of ontology mapping execution deployments; (c) enable rich annotation of ontology mappings through meta-data (see below); and (d) enable sustainable and scalable deployment of mappings through dependency modeling.

SUMMARY AND OUTLOOK

Network management systems today federate to provide end-to-end delivery of communications services. However, this federation is achieved through closed or static engineering approaches targeted toward well-defined application scenarios. While these solutions can be very effective, they typically involve significant recurring maintenance costs. In this chapter we argued that this approach will soon be no longer tenable. The continually evolving nature of large networks means that they are ever-changing in terms of the details of individual capability. Hence static federation approaches are fundamentally brittle and expensive in the medium term. In contrast the approach we advocate centers around dynamic creation and management of federations based on negotiating and enabling appropriate, minimal integration between deployed network and service management systems.

In the chapter we outlined a layered federation model that attempts to encapsulate and interrelate the various models, processes, and techniques that will be required to realize such a flexible approach to management system federation. Using an example use case based on end-to-end IPTV service delivery, we discussed the ongoing work of the Ireland-based FAME strategic research cluster [19] in pursuance of this vision. it is hoped that the outcome of this work will go some way toward making flexible federation of future autonomic network management systems a realistic possibility.

REFERENCES

[1] B. Jennings et al., "Towards autonomic management of communications networks," *IEEE Commun. Magazine*, Vol. 45, No. 10, pp. 112–121, October 2007.

[2] J.O. Kephart and D.M. Chess, "The vision of autonomic computing," *Computer*, Vol. 36, No. 1, pp. 41–50, January 2003.

[3] S. Dobson et al., A survey of autonomic communications, *ACM Trans. Auton. and Adapt. Syst.,* Vol. 1, No. 2, pp. 223–259, December 2006.

[4] H. Gogineni, A. Greenberg, D. Maltz, T. Ng, H. Yan, and H. Zhang, "MMS: An autonomic network-layer foundation for network management." *IEEE Journal on Selected Areas in Communications*, Vol. 28, No. 1, pp. 15–27, January 2010.

[5] G. Bouabene, C. Jelger, C. Tschudin, S. Schmid, A. Kelle, and M. May, "The autonomic network architecture," *IEEE Journal on Selected Areas in Communications*, Vol. 28, No. 1, pp. 1–14, January 2010.

[6] J. Strassner, *Policy-Based Network Management: Solution for the Next Generation*, Morgan Kaufmann, 2003.

[7] R. Boutaba, and I. Aib, "Policy-based management: A historical perspective," *J. Netw. Syst. Management*, Vol. 15, No. 4, pp. 447–480, December 2007.

[8] J. Moffett, and M. Sloman, "Policy hierarchies for distributed system management," *IEEE JSAC*, Vol. 11, No. 9, pp. 1404–1414, December 2003.

[9] H. Zhou and S.N. Foley, "A Framework for Establishing Decentralized Secure Coalitions," *Proc. IEEE Computer Security Foundations Workshop*, 2006.

[10] J.E. López de Vergara, V.A. Villagra, and J. Berrocal, "Applying the Web Ontology Language to management information definitions," *IEEE Commun. Mag.*, Vol. 42, No. 7, pp. 68–74, July 2004.

[11] J. Keeney, D. Lewis, D. O'Sullivan, A. Roelens, A. Boran, and R. Richardson, "Runtime semantic interoperability for gathering ontology-based network context," *Proc. IEEE/IFIP Network Operations and Management Symp. (NOMS 2006)*, pp. 55–66, 2006.

[12] N. Noy, "Semantic integration: A survey of ontology-based approaches," *SIGMOD Record*, Vol. 33, No. 4, pp. 65–70, December 2004.

[13] D. O'Sullivan, V. Wade, and D. Lewis, "Understanding as we roam," *IEEE Internet Computing*, Vol. 11, No. 2, pp. 26–33, March/April 2007.

[14] C. Hu, Y Hsu, C. Hong, S. Hsu, Y. Lin, C. Hsu, and T. Fang, "Home network management for IPTV Service Operations—A service provider perspective," Proc. 5th IFIP/IEEE International Workshop on Broadband Convergence Networks (BcN 2010), to appear.

[15] S. Balasubramaniam, D. Botvich, B. Jennings, S. Davy, W. Donnelly, and J. Strassner, "Policy-constrained bio-inspired processes for autonomic route management," to appear, *Comp. Netw.*, accepted August 2008. AQ:7

[16] S. Davy, B. Jennings, and J. Strassner, "The policy continuum—Policy authoring and conflict analysis," *Comp. Commun.*, Vol. 31, No. 13, pp. 2981–2995, 2008.

[17] S.N. Foley, and H. Zhou, "Authorisation subterfuge by delegation in decentralised Networks," *Proc. 13th Int'l Security Protocols Workshop*, 2005.

[18] K. Feeney, D. Lewis, and V. Wade, "Policy based management for Internet communities," *Proc. 5th IEEE Int'l Workshop on Policies for Distributed Systems and Networks (Policy 2004)*, pp. 23–34, 2004.

[19] FAME Strategic Research Cluster, [Online], Available (last accessed December 4, 2010): http://www.fame.ie

Author Queries

AQ:1 See what?
AQ:2 What is genericity?
AQ:3 Is this correct?
AQ:4 Did you mean on-demand?
AQ:5 Deliver the what? Text seems missing here?
AQ:6 Where is Fig. 5.3 title?
AQ:7 Please update?