

On the Management of Virtual Networks

Rafael Pereira Esteves and Lisandro Zambenedetti Granville, Federal University of Rio Grande do Sul
Raouf Boutaba, University of Waterloo

ABSTRACT

Network virtualization is a feasible solution to tackle so-called Internet ossification by enabling the deployment of novel network architectures in a flexible and controlled way. With network virtualization, it is possible to have multiple virtual networks running simultaneously on top of a shared physical infrastructure. Network management with virtualization support, however, poses challenges that need to be addressed in order to fully achieve an effective and reliable networking environment. These challenges have motivated the investigation of novel management solutions in recent years. This article surveys some of the prominent solutions for network virtualization management and identifies research opportunities in the area.

INTRODUCTION

Network virtualization has been considered a viable solution to enable novel network architectures and to overcome so-called Internet ossification [1]. In recent years, several network virtualization testbeds have been deployed, allowing researchers to propose and evaluate new solutions on a large scale with real traffic. Future Internet advocates claim that network virtualization environments (NVEs) enable diverse network architectures to coexist in a single infrastructure without affecting production services. Moreover, network virtualization is gaining attention from major industry players, and commodity network devices with virtualization support are becoming popular in the marketplace, such as routers supporting multiple virtual routing tables or programmable switches based on OpenFlow. From the commercial point of view, network virtualization redefines business relationships among participating entities. The traditional role of ISPs is split into infrastructure providers (InPs), mainly responsible for offering physical resources to service providers (SPs), who deploy virtual networks hosting a variety of services that can be accessed by end users on demand. This way, SPs have more flexibility to innovate and deploy new value-added services, which ultimately affects their revenue. NVEs also allow InPs to use dynamic pricing schemes by adjusting prices according to supply and demand.

Among the technical challenges to enable NVEs, management has special importance. Management of NVEs is crucial to guarantee the proper operation of the physical infrastructure, the hosted virtual networks, and the services supported by the virtual networks. Recently, the management of NVEs has started receiving special attention from the research community, motivating a variety of research projects over the past few years. These projects share the common goal of considering management at the network design phase, as opposed to addressing it after network deployment.

In this article, we present a survey of current advances in the management of NVEs, reviewing representative projects and highlighting their main features, benefits, and limitations. We classify research projects according to three different perspectives: *management targets*, *management functions*, and *management approaches*. A management target refers to the component being managed, such as physical and virtual nodes, intradomain links, and interdomain links. A management function denotes a specific capability supported by a management application, including resource provisioning and monitoring. Management approaches typically employed on an NVE vary from centralized and distributed management to autonomic and policy-based management.

This article differs from and complements previous surveys on network virtualization [1, 2] by exclusively focusing on the management of NVEs. The remainder of this article is organized as follows. First, we discuss management of NVEs, based on the three perspectives mentioned before (management targets, functions, and approaches). Next, research projects on management of NVEs are presented and compared. We then highlight some of the open challenges that are not yet addressed by current projects. Finally, we conclude this article.

MANAGEMENT OF NETWORK VIRTUALIZATION ENVIRONMENTS

To allow better understanding of the management of NVEs in terms of management targets, functions, and approaches, we first introduce a conceptual management model for network vir-

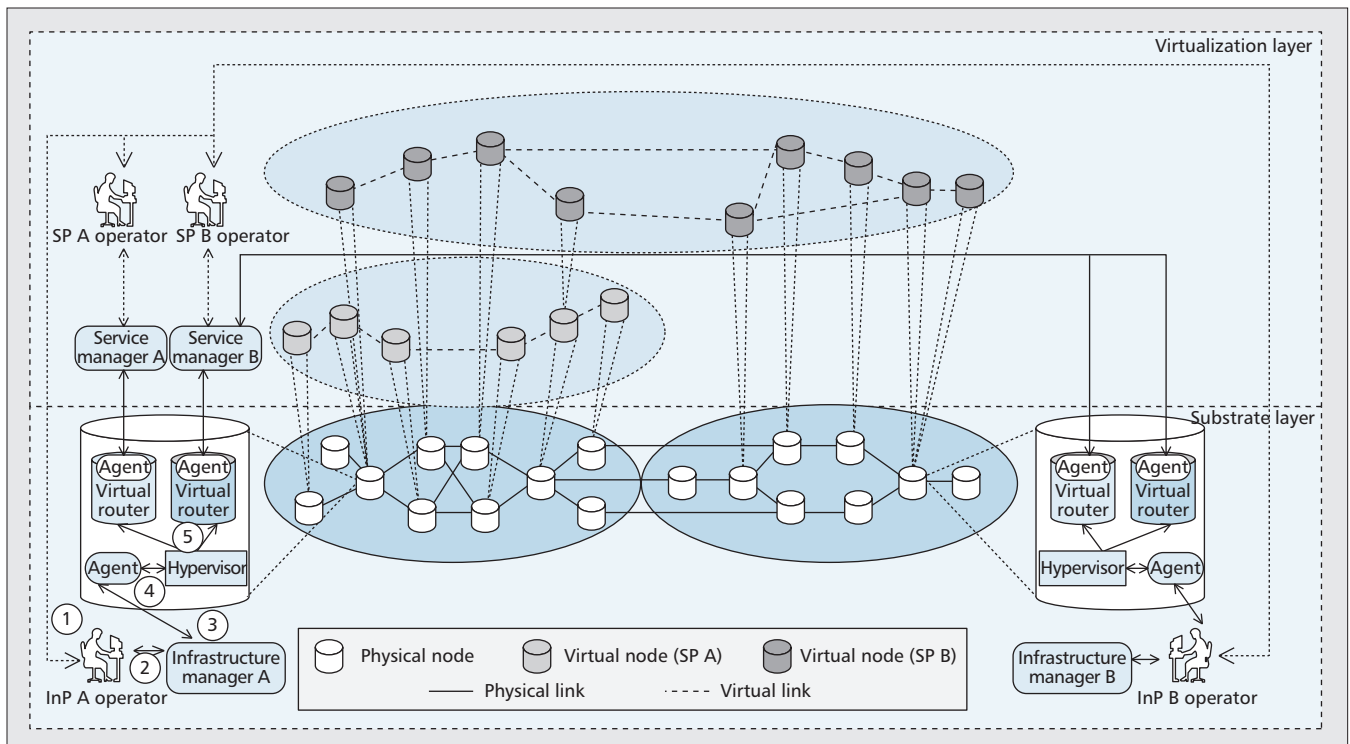


Figure 1. Management of NVEs.

tualization describing the entities, relationships, and management operations typically found in NVEs. Management operations in NVEs can be classified into InP management and SP management operations. InP management includes, for example, provisioning of virtual networks and monitoring of physical resources. SP management, in turn, deals with the operation of virtual networks and providing services to end users. Figure 1 depicts a conceptual management model for network virtualization considering the relationship between InPs and SPs.

InPs offer their physical nodes to host virtual ones owned by SPs. Several virtual nodes can be created and coexist in an isolated way inside a physical one. Physical nodes are controlled by an *infrastructure manager*, which, using a management protocol, exchanges messages with the agent located at each physical node. Once new virtual nodes are created, they are managed by the SP to which they belong. In this respect, a *service manager* communicates with the agents associated with each virtual node to collect information and enforce management actions. SPs can lease resources from different InPs to build their virtual networks. These InPs can be located in different administrative domains (or autonomous systems); thus, some level of coordination among different infrastructure managers is required.

A virtual node is hosted on a physical one. Virtual node placement can be done either manually by the SP operator or automatically by the InP using an *embedding* algorithm. Once the physical node is selected, the SP operator requests a virtual node creation to the InP operator (step 1, Fig. 1) that, using the infrastructure manager, instantiates the requested node (step 2, Fig. 1). Each physical node has a hypervisor

that allows the creation of virtual nodes. When the agent of the physical node receives a request from the infrastructure manager (step 3, Fig. 1), it contacts the hypervisor (step 4, Fig. 1), which then performs the requested action, that is, virtual node creation in the physical node (step 5, Fig. 1). In general, hypervisors provide application programming interfaces (APIs) that allow external programs to call internal operations such as virtual node creation, initialization, and removal, and to run scripts to perform fine-grained configuration of the virtual resources.

Similar to virtual nodes, virtual links are created through agents located at physical network nodes. Before contacting the infrastructure manager, the SP operator specifies the desired characteristics of the virtual links. The characteristics of a virtual link include source node, destination node, and bandwidth. Accordingly, the infrastructure manager contacts the agents at the physical nodes hosting the source and destination virtual nodes to create the virtual link. Virtual links can be created by configuring Ethernet VLANs between the physical nodes hosting the virtual ones. Multiprotocol label switching (MPLS) label switched paths (LSPs) and generic routing encapsulation (GRE) tunnels are other candidates to establish virtual links. To complete the creation of a virtual link, virtual network interfaces belonging to source and destination virtual nodes need to be bound to their respective physical network interfaces.

The “owner” of a virtual network is usually a human operator or an entity other than the owner of the physical substrate. The management of a virtual network must not affect the management of the substrate and other virtual networks. Isolated management views have to be provided to different human operators at both the sub-

We classify NVE management projects in terms of management targets, management functions, and management approaches, which combined provide a holistic understanding on how management is currently tackled in modern NVEs.

strate and virtualization layers. The isolation at the management plane is dependent on the isolation at the data plane and control plane, which is provided by enabling technologies such as virtual LAN (VLAN), MPLS, and hypervisors. In the next section we present projects that reflect the state of the art in the management of NVEs.

INITIATIVES FOR THE MANAGEMENT OF NETWORK VIRTUALIZATION ENVIRONMENTS

There are several initiatives related to the management of NVEs varying in depth and range. In next subsections, we classify and compare different projects according to the following qualitative criteria: management targets, management functions, and management approaches. We begin by discussing the criteria used to categorize the projects and then describe each proposal.

CHARACTERISTICS OF MANAGEMENT SOLUTIONS FOR NVEs

We classify NVE management projects in terms of management targets, management functions, and management approaches, which combined provide a holistic understanding of how management is currently tackled in modern NVEs. These criteria are commonly used to organize network management problems in general and are applicable to NVEs as well. Although other criteria (e.g., management lifecycle, management organization [3]) could be used to classify the projects, we focus on the ones that are most significant from the technical point of view since they lie at the core of any NVE management system and are critical at this stage of virtual network design.

Management Targets — A management target refers to a managed component of an NVE. Managed components can belong to different layers of an NVE (i.e., physical, virtual, application). Single targets can be combined into more complex ones (e.g., virtual networks), demanding additional management efforts. Here, we classify management targets as node management, link management, and network management, as described below.

Node management: Node management deals with the operation of virtual and physical nodes of an NVE, including the initial creation of virtual nodes on the substrate and node migration.

Link management: Link management addresses specific aspects related to the configuration and operation of physical and virtual links, such as virtual link isolation and flow scheduling.

Network management: Network management encompasses not only a single node or link of the NVE, but an entire virtual network, including virtual networks that span multiple physical networks.

Structuring management activities according to their target (e.g., node, link, and internetwork) can help NVE operators effectively identify and delegate management tasks (e.g., providing

isolation among multiple virtual links) based on their target. Next, we enumerate the main management functions that must be supported to realize the virtualization management model presented earlier.

Management Functions — To discuss how network management has been tackled by network virtualization projects, we identify here the main management functions that are essential in any NVE management solution. These functions (i.e., provisioning, monitoring, and interfacing) are already key in traditional networks, but gain more importance in NVEs. Provisioning allows SPs to instantiate and use virtual networks. Monitoring is used to support several other management tasks, such as fault management and billing. Interfacing defines how management applications communicate with NVE resources and enabling interoperability.

Provisioning: Resource provisioning in the context of network virtualization consists of defining the mapping of virtual network resources (e.g., nodes, links) to their physical counterparts and giving SP operators access to their virtual networks.

Monitoring: Monitoring large NVEs involves gathering updated status of physical resources and their associated virtual networks. Filtering, correlation, aggregation, and compressing of monitoring information from different sources are required to reduce management overhead.

Interfacing: Appropriate management interfaces are required for InPs and SPs to respectively access, operate, maintain, and administer the physical and virtual nodes and links. Physical network devices must present a uniform management interface, allowing virtual nodes and links located on heterogeneous physical nodes and links to be part of the same virtual network and easily manageable by the SP. The functionalities that a management interface must support include registration, creation, removal, copy, initialization, shutdown, and migration of virtual nodes and links; configuration of individual attributes of virtual resources, such as CPU and memory capacity of virtual nodes, bandwidth of virtual links, and routing tables; retrieval of status variables; and notification support.

Other management functions such as reconfiguration, migration, and scaling are also important and needed for the overall management of NVEs. However, we limit our study to basic management tasks (i.e., provisioning and monitoring) required to enable any NVE. In the following, we discuss the main management approaches employed in NVEs.

Management Approaches — Management solutions in NVEs vary in how managers and agents are organized. Some solutions rely on a centralized node responsible for performing all management tasks, while other systems allow multiple distributed nodes to share the task of managing the infrastructure. Management systems employed in NVEs can also have different levels of automation. Autonomic management helps reduce human intervention and allows dynamic adaptation to changes in the network. Policy-based management assists InP administra-

Characteristic	Management target			Management function			Management approach		
	Node	Link	Network	Provisioning	Monitoring	Interfacing	Centralized	Distributed	Autonomic
4WARD	✓	✓	✓	✓	✓	✓		✓	✓
AUTOI	✓	✓	✓	✓	✓	✓		✓	✓
FEDERICA	✓	✓	✓	✓	✓	✓	✓		
ProtoGENI	✓	✓	✓	✓	✓	✓	✓		
UCLP		✓	✓	✓	✓	✓		✓	
VNARMS	✓	✓	✓	✓				✓	✓
OpenFlow/ FlowVisor	✓			✓			✓		

Table 1. Comparison of virtualization management proposals.

tors in handling the inherent complexity of an NVE and automate resource configuration according to high-level business goals. These approaches are discussed next.

Centralized and distributed management: In the centralized NVE management approach, a single management station located at the InP (respectively the SP) is responsible for overseeing the management of the InP (respectively SP) network resources. On the other hand, in distributed NVE management, multiple nodes work in a cooperative fashion to accomplish management tasks.

Autonomic and policy-based management: Autonomic management allows the NVE to manage itself according to the current state of the network. Autonomic management solutions typically rely on high-level policies that are general rules defined to govern the functioning of the underlying network devices. In NVEs, policies are also used by InPs to enforce isolation among virtual networks by controlling access permissions for each SP.

Understanding such management approaches can help NVE administrators evaluate the trade-off between the size of the NVE and the complexity of the solution required to manage it. In the next subsection, we discuss major projects related to the management of NVEs according to the presented criteria.

NETWORK VIRTUALIZATION PROJECTS

In this survey, we have selected seven projects representative of recent work on network virtualization at the basic research, applied research, and testbed deployment levels. The surveyed projects are among the first approaches to consider management as a first-class requirement in their solutions. They also represent consolidated efforts at advanced (or already completed) development stages. In the following, we discuss these projects based on the criteria identified in the previous section and provide a summary of the discussion in Table 1.

4WARD (VNet) — The FP7 4WARD project [4] defines a network virtualization framework called VNet designed to manage multiple virtual networks hosted on a shared infrastructure.

Resource provisioning in VNet includes discovery, embedding, and instantiation. In the discovery phase, the VNet provider generates a list of candidate resources to host the virtual network. The embedding process employs a greedy algorithm to define the mapping of virtual resources to the physical network. The instantiation phase consists in actually reserving the selected virtual resources. Figure 2 depicts the 4WARD management model along with the relationships among participating entities.

VNet agents (or probes) are placed at the physical nodes to provide updated information about physical and virtual resources to the InP. The collected information is used for different purposes, including resource discovery and self-organization of the virtual networks. VNet relies on a *situation awareness framework* that aggregates monitoring information and hides unnecessary details to ensure scalability and efficiency of the monitoring process. VNet offers a management interface based on XML-RPC called the virtualization management interface (VMI) that defines a set of management operations, including creation, termination, and concatenation of virtual resources. VNet implements a distributed and autonomic management approach, referred to as in-network management (INM). In INM, self-managing entities (SEs) embedded inside the network are responsible for the autonomic operation of the physical infrastructure.

AUTOI — The Autonomic Internet (AUTOI) initiative [5] aims to develop autonomic management solutions for future Internet. Management functions are performed by distributed autonomic management systems (AMs). Different AMs can cooperate with one another in order to build end-to-end services. The AUTOI virtualization plane is responsible for the provisioning and operation of virtual networks.

Lattice is the monitoring framework for AUTOI. Lattice relies on the concept of monitoring probes and data sources. Data sources group monitored information collected by probes and send it to interested consumers following a previously defined communication model, such as publish/subscribe or IP multicast.

FEDERICA assumes that a centralized network operation center (NOC) entity performs all administrative management actions in the infrastructure, such as resource discovery, provisioning, and user control. In addition, FEDERICA offers a slice management tool to facilitate resource management.

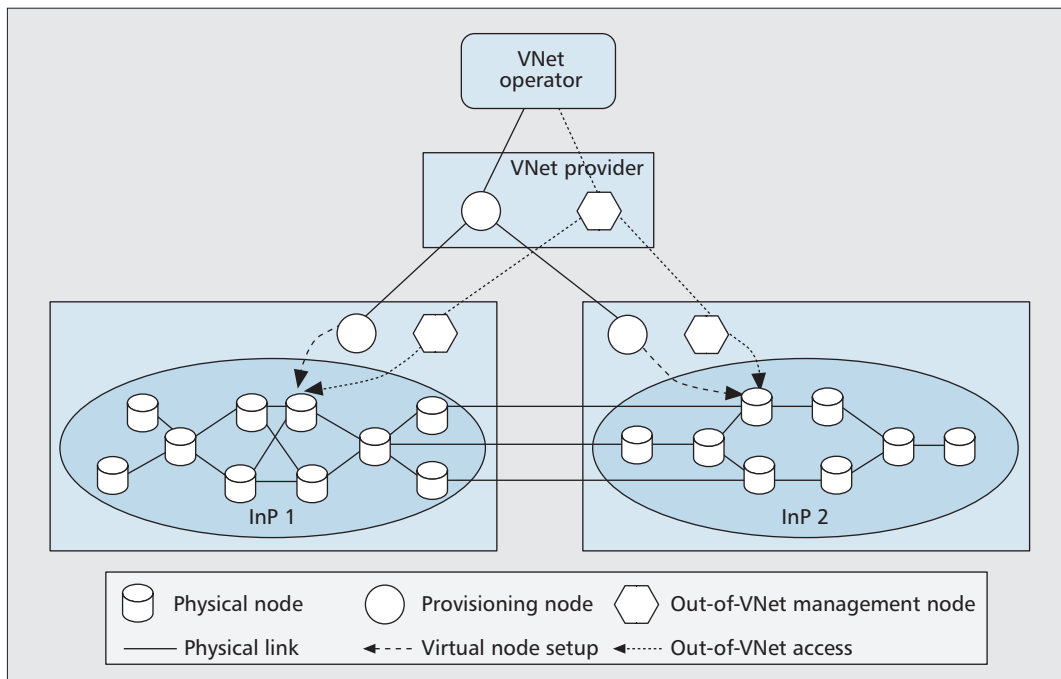


Figure 2. 4WARD VNet provisioning scenario.

Each AUTOI domain is managed by one AMS running a control loop. AMSs can cooperate with one another in order to build end-to-end services. AMSs interact with the virtualization plane through two well defined interfaces: the virtualization system programmability interface (vSPI) and virtualization component programmability interface (vCPI). vSPI provides a macro view of the virtual resources to the AUTOI orchestration plane, which, in turn, uses vCPI to build and manage virtual networks. vCPI defines basic primitives for management of virtual nodes (registerVM, startVM, shutdownVM, migrateVM, unregisterVM) and virtual links (instantiateLink, removeLink, modifyLink).

FEDERICA — Federated E-Infrastructure Dedicated to European Researchers Innovating in Computing Network Architectures (FEDERICA) [6] focuses on building a large-scale networking infrastructure to enable experimentation of new Internet protocols and architectures.

FEDERICA assumes that a centralized network operation center (NOC) entity performs all administrative management actions in the infrastructure, such as resource discovery, provisioning, and user control. In addition, FEDERICA offers a slice management tool to facilitate resource management. The latter allows the NOC operator to create slices (i.e., aggregation of virtual network resources), add virtual resources to a slice, and export slices to the SPs. SPs, in turn, can perform configurations on their assigned slices without affecting other SPs. Monitoring of physical nodes in FEDERICA is performed mainly through Simple Network Management Protocol (SNMP). FEDERICA relies on the VMWare remote command line interface (RCLI) to monitor virtual nodes.

In FEDERICA, when a researcher requests a slice, she/he contacts the FEDERICA NOC,

which creates appropriate (i.e., public, private, management) interfaces on a virtual server to allow the user to access her/his slices. The NOC then creates credentials and sets the expiration date and time of the slice. The NOC also defines the mapping of the slice on the corresponding physical machines and links, and creates VLANs to complete the slice creation.

ProtoGENI — ProtoGENI [7] is a deployed prototype of the Global Environment for Network Innovations (GENI). In ProtoGENI, researchers can create *slices* composed of *slivers*. Slivers are instances of virtual computing and networking resources. The main management entities in ProtoGENI are the *clearinghouse*, *slice authorities*, and *component managers*. The clearinghouse is the central management point in ProtoGENI, responsible for registering and tracking all slices, users, and component managers, and enabling the exchange of root certificates between ProtoGENI members. Slice authorities are the entry points for researchers to request slices from several component managers belonging to the participants of the ProtoGENI federation. Component managers control resource provisioning inside a member of the ProtoGENI federation. Figure 3 illustrates the main entities of ProtoGENI.

To obtain a slice, a researcher needs to register it at the level of a slice authority and get a corresponding credential. The credential allows one to create slivers using component managers belonging to the ProtoGENI federation. Then the researcher contacts and requests tickets from component managers. Tickets are special credentials guaranteeing that requested resources will be bound to a given slice. Both slice authorities and component managers implement an XML-RPC server and provide APIs for managing slices and slivers, respectively.

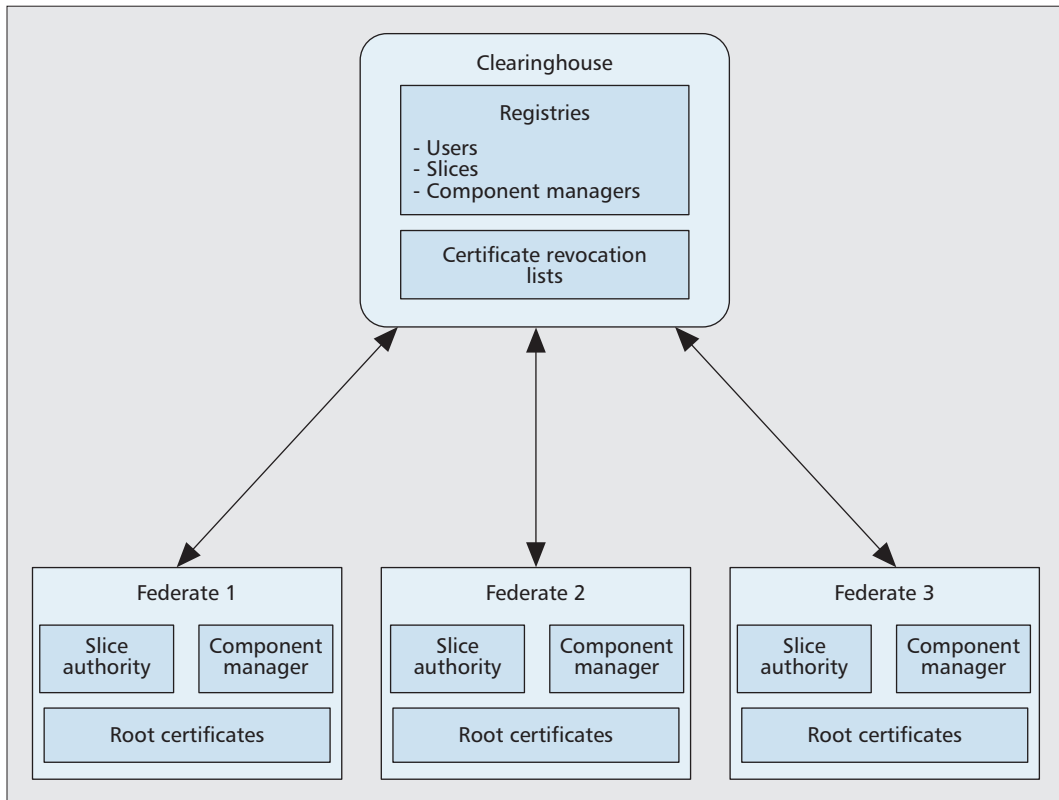


Figure 3. ProtoGENI main entities.

With FlowVisor, a switch can be properly sliced and allocated to different users. One of the main issues that FlowVisor has to deal with is managing isolation among multiple slices.

UCLP — User Controlled Lightpaths (UCLP) [8] is a management system for provisioning and controlling optical networks across multiple domains. UCLP is based on a service-oriented architecture (SOA) and allows end users to establish interdomain *lightpaths* on demand. Lightpaths can be created, destroyed, advertised, leased, and concatenated using distributed *Web services*. UCLP is structured in three main layers: the user access layer, service provisioning layer, and resource management layer, illustrated in Fig. 4.

The user access layer is the entry point from which human users can request and manage lightpath objects through a Web interface. Lightpath operations are implemented by a set of services defined in the service provisioning layer, which also acts as an access point for external applications (e.g., grid). The resource management layer comprises a set of resource agents responsible for communicating with technology-specific physical devices (e.g., synchronous optical network/digital hierarchy, SONET/SDH, switches). Monitoring in UCLP is mainly performed through standard SNMP.

VNARMS — The Virtual Network-Based Autonomic Network Resource Control and Management System (VNARMS) [9] relies on autonomic management to build virtual networks with performance guarantees. In each virtual network, there are two basic entities: the virtual network resource manager (VNRM), responsible for managing the virtual network by controlling a set of distributed resource agents (RAs) that communicate with individual network elements. Both entities are autonomic, and mon-

itor the managed resources to identify problems and react accordingly.

VNARMS uses the concept of a *root-VN* to abstract the physical network and create virtual networks. The root-VN can spawn multiple *child-VNs* that satisfy specific quality of service (QoS) requirements. When an SP requests a new virtual network, the VNRM of the root-VN calculates a topology based on the SLA requirements of the SP, spawns a child-VN from the root-VN, and instantiates a new VNRM for the child-VN. The RA of the root-VN also creates new RAs to manage individual virtual resources of the child-VN. Second-level virtual networks can be provisioned from a previously created child-VN in a recursive way, as illustrated in Fig. 5. VNARMS relies on differentiated services (DiffServ) for QoS enforcement.

OpenFlow — OpenFlow [10] is an abstraction layer that enables programming network switches. This is achieved through a flow-based abstraction in which the user/application determines the actions that will be performed by the switch on receiving packets belonging to a specific flow type.

OpenFlow requires a virtualization layer to allow multiple users or applications to share a switch. To this end, the FlowVisor [11] virtualization layer has been introduced. With FlowVisor, a switch can be properly *sliced* and allocated to different users. One of the main issues with which FlowVisor must deal is managing isolation among multiple slices. FlowVisor achieves isolation through a series of mechanisms. For bandwidth isolation, FlowVisor configures minimum

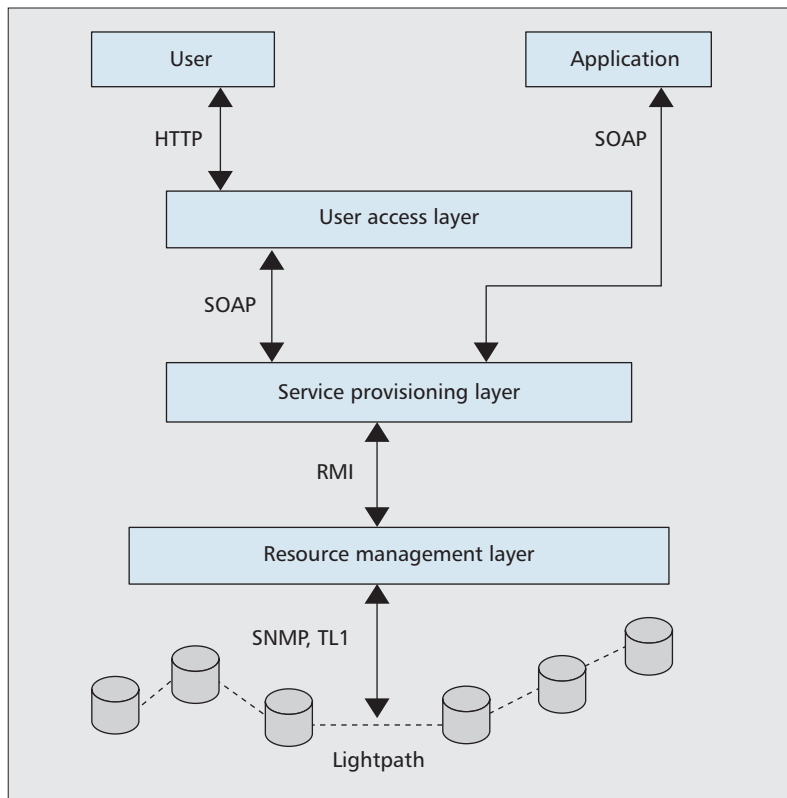


Figure 4. UCLP.

bandwidth queues for each slice sharing a part of a switch. To deal with CPU isolation, FlowVisor limits the number of control messages a user can send. Other isolation mechanisms include limiting the number of entries in the flow tables for each slice and the rewriting of control messages originated at a particular slice to prevent conflicts with other slices. OpenFlow-based switches are typically managed by a centralized controller used to create, remove, and modify flow entries.

SUMMARY

Comparing the surveyed proposals (Table 1), we found that 4WARD and AUTOI cover most of the criteria we have identified, reflecting the goal of these projects to consider management at the design stage. The other proposals represent significant achievements in the area and emphasize the trend of considering management as a first class requirement in future networks design. Another noteworthy finding is that distributed and autonomic management are considered in most solutions, except perhaps from OpenFlow-based architectures, which reflects the paradigm shift also occurring in traditional network management design, even though some functions (e.g., registry) are still performed by centralized entities.

RESEARCH PERSPECTIVES

Research on management of network virtualization is still in its infancy. There are management issues still to be uncovered and others that need further investigation. In this section, we discuss

the requirements we consider important for managing NVEs but not sufficiently explored by current solutions, and which, in our opinion, reflect future trends in managing modern NVEs. Also, considering the recent developments in cloud computing and software-defined networks (SDNs), we dedicate special attention to the issues arising in the management of virtualized cloud computing environments and SDNs.

FEDERATIONS AND SLA NEGOTIATIONS

The federation of virtualized infrastructures from multiple InPs enables access to larger-scale infrastructures. This is already happening with virtualized network testbeds allowing researchers to conduct realistic network experiments at scale, which would not have been possible otherwise. ProtoGENI is an example of federation that allows cooperation among multiple organizations. However, guaranteeing predictable performance for participating entities through service level agreement (SLA) enforcement has not been properly addressed by current solutions and remains an open issue.

INTERPLAY BETWEEN INP MANAGEMENT AND SP MANAGEMENT

This refers to the needed cooperation between the management systems of InPs and SPs to avoid/resolve conflicts and ensure overall system stability. Indeed, InPs and SPs often have conflicting management goals: the InPs want to maximize the utilization of their infrastructures and hence their revenue, and the SPs want predictable performance for their virtual networks.

STANDARD MANAGEMENT PROTOCOLS AND INFORMATION MODELS

The VR-MIB module [12] described a set of SNMP management variables for the management of physical routers with virtualization support. However, it did not progress in the Internet Engineering Task Force (IETF) standardization track, leaving the area with no SNMP-based solution. Other existing management protocols can be used instead. NETCONF, for example, would be more appropriate for configuration aspects, while NetFlow could be expanded for virtual router monitoring. There is a clear lack in this area today, which represents an interesting opportunity for research and standardization.

MANAGEMENT OF VIRTUALIZED CLOUD COMPUTING ENVIRONMENTS

Virtualization is a key enabling technology of cloud computing. In order to support a large number of customers (a.k.a. tenants), modern cloud infrastructures require that every resource (e.g., computing, storage, network) is virtualized. Open-source infrastructure as a service (IaaS) platforms such as OpenStack and CloudStack represent noteworthy developments in this respect by facilitating the development of private and hybrid clouds supporting multitenancy and advanced management capabilities. Nevertheless, several management challenges are still open, some of which are discussed below.

Dynamic resource scaling: This refers to the ability of dynamically modifying a previous resource allocation to satisfy new SP objectives. For example, an SP may need to add more virtual nodes or increase the bandwidth of a virtual link to accommodate an increasing customer (end user) base. Cloud management systems must provide *elasticity* in order to allow rapid adaptation to changes in SPs' demands. Current solutions (e.g., Amazon EC2) provide elasticity at the virtual machine level. However, dynamic capacity adjustment of network resources (e.g., bandwidth) requires further investigation.

Application-aware resource provisioning: The main limitation of current resource allocation schemes in clouds is that the characteristics of the applications are commonly ignored [13]. For example, business-critical applications (e.g., ticket reservation, order processing) may require that virtual (service) nodes are replicated and placed on distinct physical servers. On the other hand, delay-sensitive applications benefit if virtual (service) nodes are placed in edge data centers (i.e., physically close to end users) in order to reduce response time [14]. Adaptive application-driven resource provisioning can allow multiple tenants and a large diversity of applications to efficiently share cloud infrastructures.

Energy management: Energy is a main concern in cloud data centers, accounting for a significant portion of the operational costs of the InP. Achieving energy proportionality in data centers by consolidating virtual resources into a small number of physical devices can alleviate the problem. In this respect, finding a good trade-off between energy consumption and applications' performance is a promising research direction.

Data center network management: Important network issues in data center network management include address configuration, traffic management, and flow scheduling. In modern cloud data centers, the identifier of a resource is decoupled from its physical location, which requires a management infrastructure to efficiently maintain ID/locator mappings. Also, dealing with different flow patterns typically found in data centers (short vs. long flows), flow scheduling, bandwidth allocation, and leveraging the inherent path diversity of data center networks are important challenges.

MANAGEMENT OF SDNS

Software-defined networking has recently become extremely popular as a means to program network devices and customize their behavior. In SDN, the control and forwarding functions of a networking device are decoupled. This separation of control and data planes and their implementation in software offer flexibility in controlling how network devices forward packets. Commonly, SDN architectures rely on a virtualization layer, which abstracts the underlying physical network devices and topology and provides isolation in shared environments. The virtual resources are seamlessly controlled and orchestrated for the efficient delivery of network services. Management in this dynamic environment is of paramount importance. Some of the management issues that need to be addressed include:

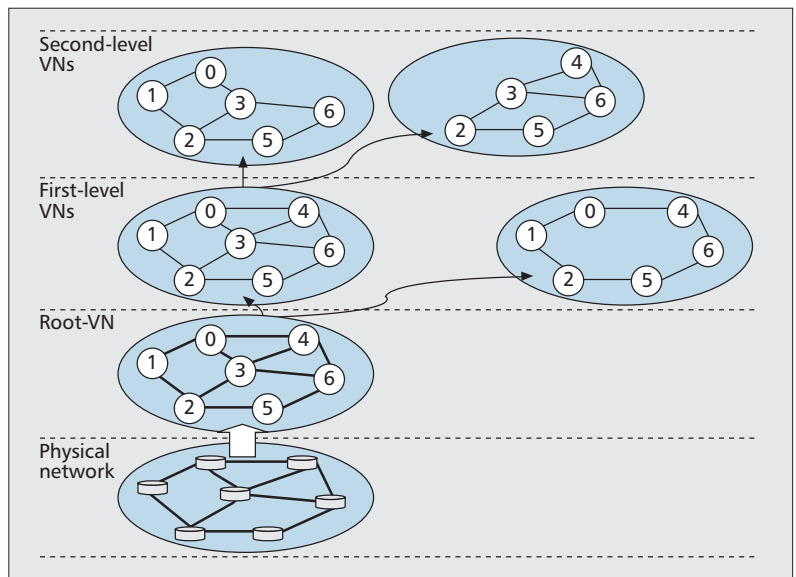


Figure 5. VNARMS.

- **Management abstractions:** Current SDN solutions require network operators to develop customized management packages using low-level instructions of a network operating system (e.g., NOX), which may be a hurdle for administrators. Providing adequate management information models, interfaces, and protocols, and advanced monitoring capabilities/tools represent opportunities to facilitate management of SDNs. The OMNI system [15] is one attempt in this direction.
- **Interoperability and management API:** SDNs can be deployed using virtualized forwarding resources from different providers using a variety of network operating systems and implementations. The provisioning of services end to end and across multiple administrative domains stresses the need for a widely accepted management API.

CONCLUSION

Network virtualization has recently gained significant importance as a viable platform, enabling the development of novel solutions to known structural problems in the Internet. However, the management of virtualized network environments raises a number of challenges yet to be addressed.

This article surveys a set of representative research projects related to network virtualization management. The surveyed projects have been analyzed from different perspectives, including their management targets, management functions, and management approach. We found that some management aspects have received or are currently receiving more attention than others. For example, resource allocation and monitoring have been extensively addressed in existing projects. Other aspects have been less so, such as autonomic and policy-based management, management federations, SLA management, dedicated management information models and proto-

We believe that the network management community should embrace this emerging area and leverage its expertise to develop a management plane for virtualized environments. Virtualized clouds and SDNs are examples of such environments, calling for novel management solutions.

cols, standard management APIs, and cooperative management in a multitenant multi-provider environment. The research issues discussed in this article are by no means exhaustive, and will be complemented by others as the research in this area progresses. In general, we believe that the network virtualization community should take advantage of the developments made by the network management community over the last three decades. In turn, we believe that the network management community should embrace this emerging area and leverage its expertise to develop a management plane for virtualized environments. Virtualized clouds and SDNs are examples of such environments, calling for novel management solutions.

REFERENCES

- [1] N. Chowdhury and R. Boutaba, "Network Virtualization: State of the Art and Research Challenges," *IEEE Commun. Mag.*, vol. 47, no. 7, pp. 20–26, 2009.
- [2] A. Khan et al., "Network Virtualization: A Hypervisor for the Internet?" *IEEE Commun. Mag.*, vol. 50, no. 1, 2012, pp. 136–43.
- [3] A. Clemm, *Network Management Fundamentals*, Cisco Press, 2006.
- [4] L. Correia et al., *Architecture and Design for the Future Internet: 4WARD Project*, Springer, 2011.
- [5] A. Galis et al., "Management Architecture and Systems for Future Internet Networks," *Towards the Future Internet — A European Research Perspective*, 2009, pp. 112–22.
- [6] P. Szegedi et al., "With Evolution for Revolution: Managing FEDERICA for Future Internet Research," *IEEE Commun. Mag.*, vol. 47, no. 7, 2009, pp. 34–39.
- [7] "ProtoGENI," available: <http://www.protogeni.net>.
- [8] R. Boutaba, W. Golab, and Y. Iraqi, "Lightpaths on Demand: A Web-Services-Based Management System," *IEEE Commun. Mag.*, vol. 42, no. 7, 2004, pp. 101–07.
- [9] M.-S. Kim and A. Leon-Garcia, "Autonomic Network Resource Management Using Virtual Network Concept," *Managing Next Generation Networks and Services, LNCS*, vol. 4773, Jan. 2007, pp. 254–64.
- [10] N. McKeown et al., "OpenFlow: Enabling Innovation in Campus Networks," *ACM Comp. Commun. Rev.*, vol. 38, no. 2, 2008, pp. 69–74.
- [11] R. Sherwood et al., "Can the Production Network Be the Testbed?," *Proc. OSDI 2010*, Oct. 2010, pp. 1–6.
- [12] E. Stelzer et al., "Virtual Router Management Information Base Using SMIv2," draft-ietf-ppvpn-vr-mib-05, June 2003.
- [13] R. Esteves et al., "Paradigm-Based Adaptive Provisioning in Virtualized Data Centers," *Proc. IM 2013*, May 2013, pp. 169–76.
- [14] "Smart Applications on Virtual Infrastructures (SAVI)," <http://www.savinetwork.ca>.
- [15] D. Mattos et al., "OMNI: OpenFlow MaNagement Infrastructure," *Proc. NOF '11*, Nov. 2011, pp. 52–56.

BIOGRAPHIES

RAFAEL PEREIRA ESTEVES (rpesteves@inf.ufrgs.br) received his M.Sc. (2009) and B.Sc. (2007) degrees from the Department of Informatics of the Federal University of Pará, Brazil. Currently, he is a Ph.D. student in computer science at the Institute of Informatics (INF) of the Federal University of Rio Grande do Sul (UFRGS), Brazil. His research interests include network and service management, network virtualization, cloud computing, and software-defined networking.

LISANDRO ZAMBENEDETTI GRANVILLE (granville@inf.ufrgs.br) received his M.Sc. and Ph.D. degrees in computer science from INF, UFRGS, in 1998 and 2001, respectively. Currently, he is a professor at INF-UFRGS. He is Co-Chair of the Network Management Research Group (NMRG) of the Internet Research Task Force (IRTF) and Chair of the Committee on Network Operations and Management (CNOM) of the IEEE Communications Society (ComSoc). He was also Technical Program Committee Co-Chair of the 12th IEEE/IFIP Network Operations and Management Symposium (NOMS 2010) and 18th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management (DSOM 2007). His research interests include network and services management, software defined networking, network virtualization, and information visualization.

RAOUF BOUTABA [F] (rboutaba@uwaterloo.ca) received M.Sc. and Ph.D. degrees in computer science from the University Pierre & Marie Curie, Paris, in 1990 and 1994, respectively. He is currently a professor of computer science at the University of Waterloo. His research interests include resource and service management in networks and distributed systems. He is the founding Editor-in-Chief of *IEEE Transactions on Network and Service Management* (2007–2010) and on the Editorial Boards of other journals. He has received several best paper awards and other recognitions such as the Premier's Research Excellence Award, and the IEEE Hal Sobol, Fred W. Ellersick, Joe LoCicero, Dan Stokesbury, and Salah Aidarous Awards. He is a fellow of the Engineering Institute of Canada.