Towards a Secure Network Virtualization Architecture for the Future Internet

Pedro Martinez-Julia[1], Antonio F. Skarmeta[1], and Alex Galis[2]

[1] Department of Communication and Information Engineering,
University of Murcia, 30100, Murcia, Spain
{pedromj,skarmeta}@um.es
[2] Department of Electrical and Electronic Engineering, University College London,
Torrington Place, London WC1E 7JE, United Kingdom
a.galis@ee.ucl.ac.uk

Abstract. In this paper we discuss the Global Virtualization Architecture (GVA) that enables communications between network entities according to the way they refer to each other rather than understanding the constraints of particular networks. Our approach is to instantiate a virtual network that is based on identities of network entities and their demands on security and network capabilities. An entity may be physical e.g. a human user, a device, or any *thing*, or abstract, such as a computer program, service, group, or role. Each entity is identified by a set of attributes so that connections can be 1 to 1, 1 to many, or many to many. We call this a Virtual Group Network (VGN). VGNs are independent of location and device, and their properties may change with time as entities move.

Keywords: Network, Virtualization, Architecture, Security, Future Internet.

1 Introduction

The original Internet model overloads the use of IP addresses as both identifiers and locators. This adds undesirable complexity to the solutions targeting unconsidered scenarios, such as mobility and multi-homing. Such model is also based on host-based end-to-end protocols, meaning that all end systems need to implement the same protocol stack to communicate each other. This significantly reduces scalability as well as the possibility to assume the growing heterogeneity of devices, specially small and low-power smart objects that do not have the same capabilities as high-end servers.

We propose to resolve those problems by separating the network functions into different, uncoupled, and hierarchical scopes. We call it the Global Virtualization Architecture (GVA). On it, access networks have their own protocols and locator spaces, which are specialized for connecting devices to the network. The global transit network, to which access networks are connected, also has a specialized protocol and locator space to deliver information units (packets or messages)

A. Galis and A. Gavras (Eds.): FIA 2013, LNCS 7858, pp. 141–152, 2013.

among access networks. For example, resource-constrained devices may rely into special gateways that implement the network functions they miss in order to allow them interact with any other network entity.

In the GVA model, global communications are implemented by using identifiers that represent network entities, thus achieving the decoupling of location and identification. This permits to provide additional and specialized functionality in the form of network virtualization. Moreover, this approach separates the data and control plane into specialized virtual infrastructures with different network properties. Thus, specific traffic types may be assigned to their own virtual network. This enables, for example, the to treat mobility management exchanges in a different way than data exchanges.

The remainder of this paper is organized as follows. In Section 2 we describe some outstanding architecture proposals for the Future Internet. Then, in Section 3 we analyse the proposals to determine their strengths and weaknesses. In Section 4 we describe how we designed GVA to overcome the limitations shown by analyzed architectures. In Section 5 we discuss the viability of the proposed architecture. Finally, in Section 6 we conclude the paper.

2 Current Architecture Proposals

In this section we give a brief description of the approaches that may resolve the aforementioned problems found in the current Internet. We selected solutions which propose a reasonably complete architecture, separate effectively node identifiers and locators, and resolve the scalability problems.

In TRIAD [6], the authors propose an IPv4 NAT-based architecture in which FQDNs are used as identities, being mapped directly to next hops. Here, routing uses the Name-Based Routing Protocol (NBRP [9]). TRIAD needs to use resolution to reach objects outside their home realm with related scaling problems.

From the solutions that strictly separates identifiers and locators, the Host Identity Protocol (HIP) [26] achieves it by introducing cryptographic host identifiers forming a new global name space as a new intermediate layer between the IP and transport layers. On the other hand, the Locator/Identifier Separation Protocol (LISP) [8] provides routing scalability through a mapping and encapsulation solution to separate global identifiers used by end nodes in their access networks and global locators used to traverse the global transit network (the Internet).

The Heterogeneity Inclusion and Mobility Adaptation through Locator ID Separation (HIMALIS) [18, 23, 25] also follows the separation of identifiers and locators but, in contrast with LISP and HIP, it provides a higher level of routing scalability while providing an independent node identification scheme, which is different from IP and supports sensor networks (e.g. the Internet of Things). It differentiates between local locators (LLOC) and global locators (GLOC), which are correspondingly used in access networks and the global transit network. It therefore includes a new naming scheme for generating host names and IDs.

Mobile Oriented Future Internet (MOFI) [17, 24] is an internetworking architecture that follows the separation of identifiers and locators principle while

efficiently supporting envisioned mobile oriented and network diversity environment. MOFI basically pursues ID-based communication. In MOFI, a host is the basic entity that should be identified for communication but can be extended to other objects such as user, contents, etc. Locator for routing in each network is separated from the host ID.

Overlay networks are built on top of other architectures to overcome their limitations. Chord [37] is a decentralised lookup service for mapping keys to nodes, much like a churn-tolerant DHT (Distributed Hash Table), which can be used for routing on flat identifiers (labels). Chord has a fairly good performance and has many improvements, such as LPRS [39]. However Chord has problems to recover from its ring partitioning and lacks security. The architecture proposed by Routing on Flat Labels (ROFL) [5] builds on Chord to provide its benefits to the whole network, removing the necessity of hierarchical addresses in the Internet.

The clean-slate architectures jettison the current Internet. MILSA [29] and Enhanced-MILSA (EMLISA) [28] explore a novel, clean-slate architecture for the Internet, based on the principle of separating identifiers and locators but with other capabilities. They distinguish between realms (organisational areas) and zones (network connectivity areas). MILSA relies on a stable layout of the RZBS (Realm-Zone Bridging Servers), which must be pre-configured. DNS is used for resolving RZBS names (but not for other nodes). Both solutions however lack specific security mechanisms.

The Information-Centric Networking (ICN) raises the role of information to the middle of communications. The EU-funded projects 4WARD [4] and its successor, SAIL [7], are defining an ICN architectural paradigm called Network of Information (NetInf) [3] that extends the concept of identifier/locator split with another level of indirection and decouples self-certifiable objects from their storage location(s). Another EU-funded project, PURSUIT [38], also approach the ICN but proposing a publish/subscribe view where information consumers subscribe to the information they want and information providers "publish" it. Finally, Content Centric Networking (CCN) [16] is an architecture built on named data. It has no notion of host at its lowest level – a packet *address* names content, not location. However, it preserves the design decisions that make TCP/IP simple, robust and scalable.

3 Analysis of Capabilities

In this section we analyse the strengths and weaknesses of the proposals discussed in the previous section, which are HIP, LISP, HIMALIS, MOFI, ROFL, EMILSA, NetInf, PURSUIT, and CCN. For each of them we evaluated its strengths/weaknesses in several aspects by using the following parameters:

- how much architected support for policies they have [A];
- how scalable they are [B];
- how independent they are of the DNS scheme and IP layout [C];

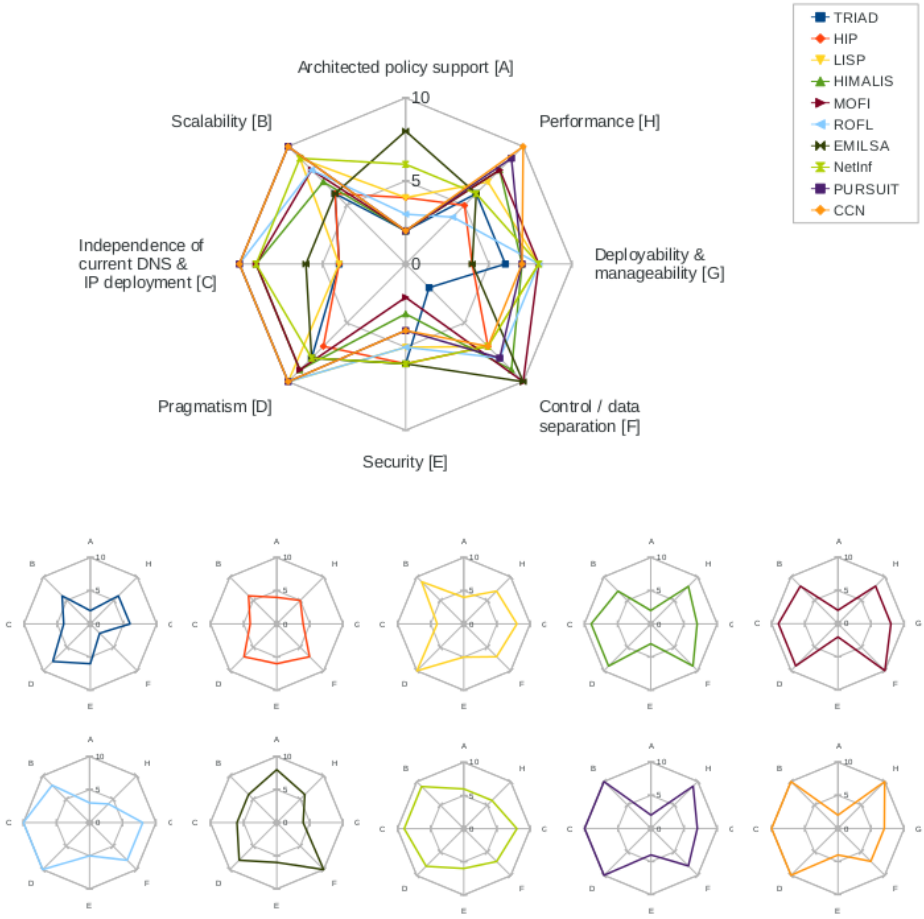

Fig. 1. Results of the strengths analysis of current proposals for the Future Internet. Each aspect has been assigned a value from 0 to 10 for each architecture in the top figure. The bottom figures represent each architecture by its own. It is clear to see that the biggest and less sharp area the better the architecture is.

- how pragmatic they are, as opposed to purely theoretical approaches [D];
- how secure they are [E];
- how much separation they manage to do between control and data flows [F];
- how deployable and manageable they are [G];
- how well they perform [H]

The results of evaluating these parameters are summarised in Figure 1. It shows how all approaches are lacking on security and many of them also lack in policy support. Specifically, even though TRIAD is a fairly complete solution, based on IPv4 and, in principle, quite deployable and scalable, it lacks however an explicit

policy framework, and is too dependent on IP addresses being topologically aggregated, and also on nodes following closely the DNS hierarchy.

This deficiency is addressed in HIP and LISP, as also seen in their figures, but they also lack in security, policy support, and independence of IP. Both HIMALIS and MOFI have similar features, as depicted in their diagrams, and share with HIP and LISP the lack on security and policy support. As a totally different architecture, ROFL improves in the independence of IP, pragmatism, and scalability. It also improves in security but neglects the performance and policy support. As it is also a totally different architecture, EMILSA have a good score in policy management, as well as in achieving a good separation between control and data. Its main drawbacks are the need for pre-configuration (placement, population, and management of directory services) and the little security services supported by the architecture. Even though the final three architectures are information-centric, there is a clear separation from NetInf and PURSUIT-/CCN. NetInf has some advancements in security and policy support, being one of the most complete architectures evaluated here. In contrast, PURSUIT and CCN, which are very similar, have many improvements in performance, scalability, pragmatism, independence of IP but, like many other architectures commented above, they lack in security and policy support.

4 Global Virtualization Architecture

The Global Virtualization Architecture (GVA) is designed to overcome the limitations of the current Internet. As shown in Figure 2, it is based on three main functional blocks: 1) A Connectivity Virtualization Layer (CVL) that abstracts the complexity of the underlying networks into different virtualized networks with different properties, like using specialized routing or specific mechanisms for information delivery; 2) An Application and Service Virtualization Layer (ASVL) that manages and offers the Virtual Group Network (VGN) functionality to applications and services; and 3) The Information Infrastructure (I2) that provides the control, management, and security functions as a separated plane to manage the network paths to forward information, the security properties of communications, and support the mobility capabilities.

To offer the benefits of the architecture to the users we defined the VGN concept as a mechanism to build a specific view of the whole network with specific properties and for a specific set of entities. The VGN properties define the security aspects of the communications, the network constraints, etc. Connected entities may be real entities, such as people, software, machines, things, etc. or abstract, such as groups or roles, each represented by its identity with its set of attributes. We emphasise the differentiation of identity and identifier. GVA meets with the ITU-T definition of identity in its X.1250 [12] and the ITU-T definitions for Future Networks [13–15] recommendations.

That said, an identity is the representation of an entity in the form of one or more information elements (attributes) which allow an entity or group of entities to be sufficiently distinguished within context. On the other hand, an identifier

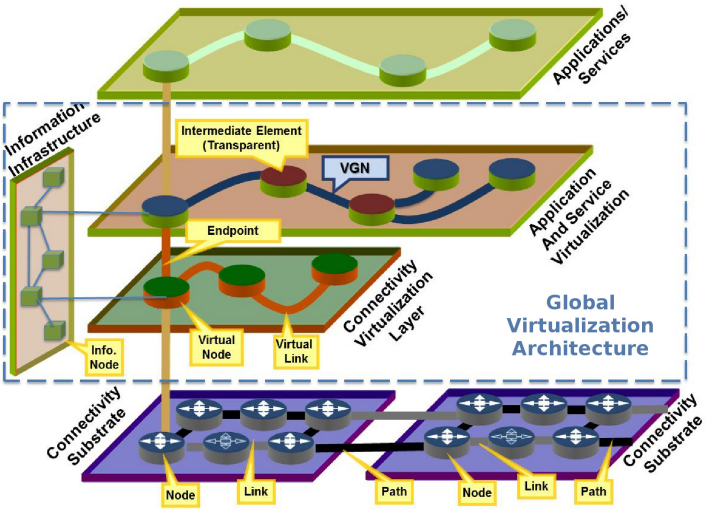

Fig. 2. Architecture Overview

is a piece of data that uniquely identifies something in a particular domain. In a general sense, the ASVL and subsequent communications use identifiers to determine the endpoints of the communication participants and to obtain information from a digital identity if permitted by policies. Nevertheless, they are not used to unambiguously associate an identity to an object over time and space, but rather just at a certain moment and communication event.

We build the CVL on the separation of end-point identifiers (EPID) from network locators (LOC). An EPID is a persistent session-based ID for each network node (entity), which ensures ID-based communication. LOC may be the IP address to which the end-point is attached but it is open to a different address space. End-to-end communication between two end-points will only use their EPIDs, whereas data packets will be delivered to an end point by using the associated LOCs, possibly through one or more transit networks (or even different architectures).

This approach permits GVA to be instantiated on top of any architecture with a smooth evolution path. It abstracts from the network topology to create particular network views for each communication using optimal intermediate elements to transfer information from source to destination. A special entity inside I2, called the Domain Trusted Entity (DTE) [20,21], has been introduced to each domain and the instances of all domains are joined through the overlay network to build a global trusted infrastructure [22] capable of negotiating the properties of communications, including security.

The VGN model, in conjunction with the I2, supports security and privacy by design, because the information about an entity, which is identified by its identity, will only be given to the allowed entities, as well as cross-layer interactions to enable end-to-end communications including identification of the endpoints.

The identity-based control plane of the I2 provides inter-administrative domain support the federation and building on a trust model that supports end-points and their attributes and identification. Requirements for machine-to-machine communication in self-managing actuator networks, like smart grids, will be covered.

Finally, GVA will inherit proven concepts developed by Identity Management, as well as from locator-identifier and publish-subscribe approaches. New networking scenarios can be instantiated on top of the existing Internet, while providing spotlight examples of new architectural approaches that illustrate the openness to the evolving heterogeneity of the infrastructure and the magnitude of endpoints to validate the approach.

4.1 GVA Features

As a consequence of this design, the GVA operation framework provides a dynamic connectivity model with support for nomadicity and variable reliability through the concept of VGN. It provides a virtual structure within the network allowing arbitrary devices to connect and reconnect to a session. At the same time, VGNs can opt to have a certified level of trust, where the trust level may vary. This solution complements trusted devices and trusted services to provide an overall controllable level of trust for transactions and communication.

GVA supports sustained connectivity for moving endpoints by introducing new protocols and extensions to existing ones. They handle the setting up and tearing down of connections, as well as mobility and nomadicity. General methods known from existing networks will be the starting point to those new protocols developed to manage the complexity of moving endpoints with possibly multiple interfaces and contexts.

The approach to network management and control does not involve network addresses but the identities behind communications. This way, each VGN has specific parameters negotiated by the network and based on *who* is using the network and *what* it can provide. In the same way, when analysing network traffic, it can not be associated to a specific entity, although if permitted by policies, the network can reveal *to whom* messages pertain.

GVA provides an extensible framework through APIs for transport, discovery, name resolution, session control, identity management, security and privacy management will complement the operational framework of GVA and will make it open for new applications and services.

4.2 Functional Blocks

As introduced above and shown by Figure 3, GVA is based on the definition of the CVL as a lower layer functional block, the ASVL as an upper layer functional block, and the I2 as a side plane functional block. Below we detail each component.

The CVL abstracts the specific mechanisms and shape of the underlying networks, such as IPv6, HIMALIS, MOFI, CCN, PURSUIT, etc., to offer a specific

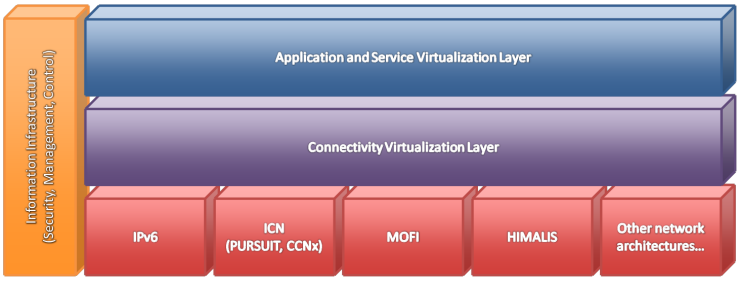

Fig. 3. Functional Blocks of the Architecture

network view that meets with the specific requirements of future communications. The construction of this layer is heavily based on the overlay network concept to support integration of multiple underlying network architectures. It can be seen as a raw arena on which any approach is able to grow without the rigid limitations of current infrastructures, which can also be modelled in different ways to accomplish different communication objectives while being instantiated on top of disjoint and heterogeneous underlying network architectures. In addition, this component permits the underlying network architectures to cooperate and thus offer the best combined functionality to the upper layers and therefore to the users.

Beyond using identifiers, GVA proposes to use attributes that are attached to entities (digital identities). These attributes, such as belonging to a specific group or a specific role within a company, become the handles with which connectivity is established, going far beyond existing approaches, such as the identifier/locator split. ASVL is intended to abstract data communications and thus facilitate to upper layer elements and application clients the access the functionality provided by GVA in the form of VGNs, as described above.

At the identity level, ASVL uses the eXtensible Resource Identifier (XRI) and eXtensible Resource Descriptor Sequence (XRDS) [33]. XRI is used to build the universal identifiers, which are related to identities, resources, and may also be related to VGNs and context identifiers. XRDS is, in turn, used to describe the resources of the endpoints owned by each entity. Thus, the ASVL includes a dynamic but consistent identifier scheme that permits the interaction with other identity management technologies and architectures like SAML [35], Shibboleth [36], and OpenID [32].

All this is achieved by defining context envelopes to wrap communication descriptors (identities of participants, security requirements, environment, etc.). The necessary discovery mechanisms to enable network nodes find each other and know the services they offer is also integrated into the architecture.

I2 is a vertical plane that supports the operations performed by CVL and ASVL. It is intended to abstract ASVL from network complexity by permitting the negotiation of communication parameters (including security) from the identity of their participants, in a secure and isolated manner. Creating this

infrastructure is part of our *Security by Design* approach, where security is built into the network infrastructure at the design stage, rather than trying to bolt it on afterwards, which is never an effective or reliable way of managing risks. Therefore, I2 is built with the integration of a security plane, a management plane, and a control plane.

The security plane is built on top of a separate network infrastructure to be totally independent of the underlying network and upper layer architectures. Also, this separation facilitates the fact that data and control messages/packets are separated from the security messages/packets. However, the necessary intermediate elements to build the security plane can be instantiated in the same logical elements or physical equipment as the context control plane.

Finally, communications in GVA are not bound to addresses or identifiers derived from the network attachment point but instead use special context identifiers to identify the VGN endpoint. Thus, a context enclosed in a VGN represents the information that wraps and defines a communication, including the identities of its participants (senders and receivers), the security requirements and parameters, the environment, the communication channel, and the path followed by the information in that channel.

5 Initial Viability Analysis

To get an initial view of the feasibility and viability of GVA we performed a quick study of currently existing solutions to see how they might need to be changed in order to be integrated together to obtain the functionality of some of the components and modules defined in the architecture.

As CVL is primarily built with overlay network mechanisms, we can take advantage of existing approaches like Chord [37], so it is the starting point of our research for this functional block. With this and the storage extension provided by some existing database manager, such as one based on NoSQL like CouchDB, we can have the base to build both the routing table manager and the storage hashing extension modules.

For ASVL we can obtain the functionality of the identity and service discovery by integrating an existing discovery service like those found, for instance, in the SPITFIRE project [31]. Also, the service description module could be instantiated by reusing RDF solutions [19] which provides the necessary service semantics and can be easily integrated with the previous component.

Finally, to build I2 we found many existing solutions. For the mobility and multihoming module, we start by reusing the mechanisms used in other network architectures like Mobile IP [11], HIMALIS [18], or MOFI [17]. The policy engine is provided by a XACML [27] module, like XACML-Light [10], and the topology management can be provided by reusing one of the many existing topology engines. The claims validation service can be based on the credential validation services produced in the TAS3 project, whilst the claims themselves are SAML [35] assertions with holder of key confirmation elements.

6 Conclusions and Future Work

In this paper we analyse different architecture proposals for the Future Internet to know the current gaps and challenges. Then, we introduce GVA, an architecture design that fill those gaps by means of virtualization techniques working together with the overlay network concept. This architecture resolves the previously introduced shortcomings by extensive virtualization and a robust control plane. The separation of identifiers and locators is achieved by the collaboration between CVL and I2, being offered to the user through the ASVL. The global scalability and heterogeneity of underlying architectures is provided by the separation of local and global locators, defining an identity space managed by I2 to address global communications.

We briefly discuss the feasibility of the proposed solution with an initial viability analysis that places the architecture design in terms of other existing architectures, mapping each functional block with other existing architectures, solutions, and infrastructures. Therefore, in future work, we will investigate the relations of the architecture with the main design principles for the Future Internet [30], and how we can improve it by applying them. In particular we would analyse the GVA's realization as part of Software Defined Networks [34] and Network Virtualization Functions [1]. After that we plan to continue the evaluation of the architecture by building a prototype intended to perform extensive experimentation.

Acknowledgments. This work is partially supported by the European Commission's Seventh Framework Programme (FP7/2007-2013) project GN3 and UniverSELF project [2] under grant agreement 257513, by the Ministry of Education of Spain under the FPU program grant AP2010-0576, and by the Program for Research Groups of Excellence of the Séneca Foundation under grant 04552/GERM/06.

References

1. Network Functions Virtualisation White Paper (2012), `http://www.tid.es/es/Documents/NFV_White_PaperV2.pdf`
2. UniverSELF project (2013), `http://www.univerself-project.eu/`
3. Ahlgren, B., D'Ambrosio, M., Marchisio, M., Marsh, I., Dannewitz, C., Ohlman, B., Pentikousis, K., Strandberg, O., Rembarz, R., Vercellone, V.: Design considerations for a network of information. In: Proceedings of the 2008 ACM CoNEXT Conference, pp. 1–6. ACM, New York (2008)

4. Brunner, M., Abramowicz, H., Niebert, N., Correia, L.M.: 4WARD: A European perspective towards the future internet. IEICE Transactions on Communications E93-B(3), 442–445 (2010)
5. Caesar, M., Condie, T., Kannan, J., Lakshminarayanan, K., Stoica, I.: Rofl: Routing on flat labels. SIGCOMM Computer Communication Review 36(4), 363–374 (2006)
6. Cheriton, D.R., Gritter, M.: Triad: A scalable deployable nat-based internet architecture. Tech. rep. (2000)
7. Edwall, T., et al.: Scalable and Adaptive Internet Solutions, SAIL (2011), http://www.sail-project.eu
8. Farinacci, D., Fuller, V., Meyer, D., Lewis, D.: Locator/id separation protocol (LISP). Internet-draft, IETF (2012)
9. Gritter, M., Cheriton, D.R.: An architecture for content routing support in the internet. In: Proceedings of the Usenix Symposium on Internet Technologies and Systems (2001)
10. Gryb, O., et al.: XACML Light (2010), http://xacmllight.sourceforge.net
11. Gundavelli, S., et al.: Proxy Mobile IPv6 (2008), http://www.ietf.org/rfc/rfc5213.txt
12. ITU-T: Series X: Data Networks, Open system communications and security. Cyberspace security - Identity management. Baseline capabilities for enhancing global identity management and interoperability. Recommendation ITU-T X.1250 (2009)
13. ITU-T: Y.3001 Recommendation: "Future Network Vision - Objectives and Design Goals" (2011)
14. ITU-T: Y.3011 Recommendation: "New Framework of Network Virtualization for Future Networks" (2011)
15. ITU-T: Y.3021 Recommendation: "New Framework of Energy Saving for Future Networks" (2011)
16. Jacobson, V., Smetters, D.K., Thornton, J.D., Plass, M.F., Briggs, N.H., Braynard, R.L.: Networking named content. In: Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies (CoNEXT 2009), pp. 1–12. ACM, New York (2009)
17. Jung, H., Koh, S.J.: MOFI: Future internet architecture with address-free hosts for mobile environments. Telecommunications Review 21(2), 343–358 (2011)
18. Kafle, V.P., Inoue, M.: HIMALIS: Heterogeneity inclusion and mobility adaptation through locator id separation in new generation network. IEICE Transactions on Communications E93-B(3), 478–489 (2010)
19. Klyne, G., Carroll, J.J.: Resource Description Framework (RDF): Concepts and Abstract Syntax (2004), http://www.w3.org/TR/rdf-concepts/
20. Martinez-Julia, P., Gomez-Skarmeta, A.F.: A novel identity-based network architecture for next generation internet. Journal of Universal Computer Science 18(12), 1643–1661 (2012)
21. Martinez-Julia, P., Gomez-Skarmeta, A.F.: Using identities to achieve enhanced privacy in future content delivery networks. Computers and Electrical Engineering 38(2), 346–355 (2012)
22. Martinez-Julia, P., Gomez-Skarmeta, A.F., Girao, J., Sarma, A.: Protecting digital identities in future networks. In: Proceedings of the Future Network and Mobile Summit 2011, pp. 1–8. International Information Management Corporation (2011)
23. Martinez-Julia, P., Gomez-Skarmeta, A.F., Kafle, V.P., Inoue, M.: Secure and robust framework for id/locator mapping system. IEICE Transactions on Information and Systems E95-D(1), 108–116 (2012)

24. Martinez-Julia, P., Skarmeta, A.F., Jung, H.Y., Koh, S.J.: Evaluating secure identification in the mobile oriented future internet (mofi) architecture. In: Proceedings of the Future Network and Mobile Summit 2012, pp. 1–8. International Information Management Corporation (2012)
25. Martinez-Julia, P., Skarmeta, A.F., Kafle, V.P.: Research and experimentation with the himalis network architecture for future internet. In: Proceedings of the Future Network and Mobile Summit 2012, pp. 1–8. International Information Management Corporation (2012)
26. Moskowitz, R., Nikander, P.: Host Identity Protocol (HIP) Architecture (2006), http://www.ietf.org/rfc/rfc4423.txt
27. OASIS XACML Technical Committee: XACML: eXtensible Access Control Markup Language (2010), http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml
28. Pan, J., Jain, R., Paul, S., Bowman, M., Xu, X., Chen, S.: Enhanced milsa architecture for naming, addressing, routing and security issues in the next generation internet. In: Proceedings of the International Conference on Communications, pp. 14–18. IEEE, Washington, DC (2009)
29. Pan, J., Paul, S., Jain, R., Bowman, M.: Milsa: A mobility and multihoming supporting identifier locator split architecture for naming in the next generation internet. In: Proceedings of the Global Communications Conference, pp. 2264–2269. IEEE, Washington, DC (2008)
30. Papadimitriou, D., Zahariadis, T., Martinez-Julia, P., Papafili, I., Morreale, V., Torelli, F., Sales, B., Demeester, P.: Design principles for the future internet architecture. In: FIA 2012, LNCS, vol. 7281, pp. 55–67. Springer, Heidelberg (2012)
31. Pfisterer, D., Romer, K., Bimschas, D., Kleine, O., Mietz, R., Truong, C., Hasemann, H., Pagel, M., Hauswirth, M., Karnstedt, M., et al.: Spitfire: Toward a semantic web of things. IEEE Communications Magazine 49(11), 40–48 (2011)
32. Recordon, D., Reed, D.: Openid 2.0: A platform for user-centric identity management. In: Proceedings of the Second ACM Workshop on Digital Identity Management, pp. 11–16. ACM, New York (2006)
33. Reed, D., Chasen, L., Tan, W.: Openid identity discovery with XRI and XRDS. In: Proceedings of the 7th Symposium on Identity and Trust on the Internet (IDtrust 2008), pp. 19–25. ACM, New York (2008)
34. Rubio-Loyola, J., Galis, A., Astorga, A., Serrat, J., Lefevre, L., Fischer, A., Paler, A., Meer, H.: Scalable service deployment on software-defined networks. IEEE Communications Magazine 49(12), 84–93 (2011)
35. Security assertion markup language (saml), http://saml.xml.org
36. Shibboleth, http://shibboleth.internet2.edu
37. Stoica, I., Morris, R., Karger, D., Kaashoek, M.F., Balakrishnan, H.: Chord: A scalable peer-to-peer lookup service for internet applications. In: Proceedings of the 2001 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, pp. 149–160. ACM, New York (2001)
38. Trossen, D., et al.: Pursuing a Pub/Sub Internet, PURSUIT (2011), http://www.fp7-pursuit.eu
39. Zhang, H., Goel, A., Govindan, R.: Incrementally improving lookup latency in distributed hash table systems. In: Proceedings of the 2003 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, pp. 114–125. ACM, New York (2003)