An Overlays-based Flat Layered Naming Architecture for the Internet

Joelle Quaini-Sousa, Marcio Augusto L Silva, Fernando Redígolo, Tereza Cristina Carvalho, Wilson Ruggiero
LARC-PCS, University of São Paulo
São Paulo, Brazil
{joelle, msilva, fernando, carvalho, wilson}@larc.usp.br

Hélio Crestana Guardia
GSDR-DC, Federal University of São Carlos
helio@dc.ufscar.br

Börje Ohlman
Ericsson Research, Stockholm, Sweden
borje.ohlman@ericsson.com

***Abstract* — A multitude of overlay network projects are emerging today providing the Internet with advanced features its architecture did not naturally support (routing, naming, addressing, mobility and security). Numerous complex scenarios faced today such as mobile devices, multi-homing equipments and middleboxes were not envisioned on TCP/IP conception and are already posing a set of convergence challenges. Our objective is to investigate and compare several examples of these advanced network layers deployed on the Internet and propose a flat layered naming architecture for the Internet.**

Keywords –Overlay Networks, Naming Architectures, Network Protoco and, Mobility Computing.

I. Introduction

The massive widespread of the Internet interconnecting millions of computers and new technologies posed a set of challenges to this architecture and evidenced its shambles. Due to some early design decisions the Internet presents shortcomings in areas such as routing, naming, addressing, security and mobility [2]. This fact, allied to the static and conservative nature of the (commercial) Internet, makes these changes a challenging task, despite being feasible from a pure technological standpoint. The introduction of new functionalities and technologies in the Internet has become more and more difficult [3] as the slow paced adoption of IPv6 illustrates).

In the reference model for network architecture from Saltzer [1], some elements can be identified, such as Services (functionalities that allow users to perform operations) and its Users, the nodes and Network Attachment Points (i.e. address) [4] and Path (the link between two communicant nodes). In the current Internet model, the distinction between a node identifier and the network attachment point identifier (i.e. IP address) is blurred, leading to a “semantic overload” in the IP.

Here, a clear containment relationship emerges: several services used by several users for a given node, which, on its turn can have several network attachment points (e.g., multiple network interface cards in a multihomed server). Regrettably, a serious problem emerges in this scenario as it is not possible to granularly identify these entities. One can not differentiate, for instance, a node with three different network attachment points from a three nodes with a single attachment point each.

Even worse, this model only considered static network attachments. With newer network access technologies, such as IEEE 802.11 and GPRS (General Packet Radio Service) [5] [6] as well as mobile computing, a node can dynamically attach to different networks choosing the one most suitable in a given moment. This network mobility causes changes in the IP addressing, disrupting communication already established with other nodes, despite the immutability on the nodes identity.

Among the various inadequacies, two of them are worthy of further examination: routing and naming. The routing front is dealt by addressing and forward decision protocols, like IPv6, MPLS/GMPLS and Multicast (several protocol proposals). The naming aspect, on other hand, seems to receive significantly less attention, being even subject of debate when cited as a complete different function from addressing.

This paper is divided as follows: section one presents an Introduction that covers the main TCP/IP limitations and explains fundamental concepts for the architecture. Section two comprises a discussion of Application Specific Networks (ASNE). The third section illustrates the naming characteristics in the current Internet and is followed by the fourth section where the Flat Naming Architecture itself and a simple proof of concept are presented to investigate the feasibility of our proposal in conjunction with the DNS. Finally, the conclusion in the last section ends the paper and point out some open questions.

II. Application Specific Networks

One of most well established and thought-trough methods of implement new functionalities is by the use of application specific networks (ASNE) [7]. An ASNE is a network formed by the aggregation of a number of nodes that communicates through a substrate network. It has its most common form of implementation as an Overlay Network. We have used two main criteria to classify ASNEs: the (OSI) layer where the ASNE is located and its purpose (or functionality).

Regarding the layer, the ASNEs can be classified as underlays (whose functionalities are implemented below the network layer) or as overlay (whose functionalities are implemented above the network layer). As an underlay the Network Pointer project [8] could be cited and as an overlay HIP [9]. Hi3 [10] and i3 [11] projects are the best well-known examples. Regarding the purposes, it could be classified as a naming, security, mobility or routing overlay.

In this paper, this key concept (the notion of ASNE) is explored. The establishment of an analogy between the data and metadata concepts, and its semantic counterparts applied to networks is essential. In order to demonstrate the feasibility of such expectation, it is our final aim to propose a novel naming overlay for the Internet. Taking the risk to sound pretentious at the first sight, we firmly believe that some of intrinsic characteristics of the proposed architecture can be, based on the previous considerations, realistically suitable for a slow but widespread adoption

A. Routing versus Non-Routing Overlays

The previous categorization is a consequence of a rather extensive analysis of proposed ASNEs we have conducted as part of our work [12]. Two main conclusions can be drawn from such compilation. The first is the ubiquity of overlays relative to the use underlays, rendering these to an almost niche application role. The second is the distinct separation of overlays in two great functional packs: routing and “non-routing” overlays. In a “non-routing” overlay, the primary concern is the provision of a designed functionality (e.g. naming resolution service) by establishing control channels for the metadata (e.g. sending the name and address of an object and, sometimes the route to get to it) separate from the data (e.g. the actual content that is referred by the triple: address, name and route) [13].

This way, it is left for the lower layers and devices of the network (e.g. IPv4, IPv6 or MPLS routers in the Internet) the task of actually moving the data from one point to another. In this form of ASNE, new functionalities and protocols can be transparently and incrementally introduced with minimal, if any, disruption from the substrate network standpoint, being subject to modification only the participating nodes.

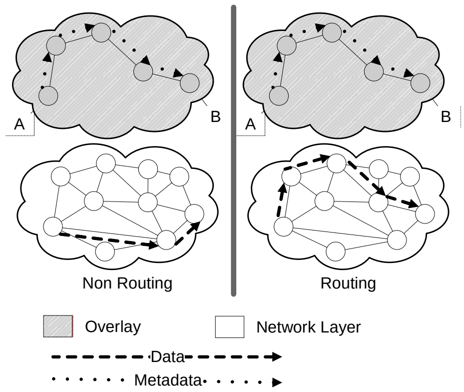

Figure 1 – Routing and “Non-Routing”overlays

On a routing overlay, on other hand, there is data and metadata flowing directly through the hosts and devices that compose the overlay network, augmenting or even replacing the current routing (data movement). Figure 1 illustrates the above explanation.

In our work we have analyzed several ASNE proposals [14] [10] [11] [8] [13] regarding their purpose, i.e., the functionalities they bring to the current Internet architecture.

Naming overlays deal with providing to its substrate network infrastructure special capabilities to identify and resolve its objects and are the target subject of our work.

Routing overlays intend to replace or complement the network substrate routing mechanism with enhanced functionalities, including overlayed unicast/multicast forwarding techniques, source routing or loose source routing. Security overlays provide additional capabilities to prevent the basis network from suffering security attacks or to remedy already existent attacks. At last, mobility overlays intend to support mobility to a pair of communicating nodes, allowing one or both entities to move away from their domains, changing their network attachment point during an already established connection while maintaining the semantic connection initially established (respectively called single jump or double jump support).

B. Overlays Benefits andDisadvantages in a Naming Infrasctructure

The main benefit of using overlays for enhancing the Internet with features previously not natively supported is the possibility of an incremental deployment: there is no need to deploy new equipment, or modify existing software/hardware/protocols. Also, they do not have to be deployed at every node. However, there are some disadvantages of using overlays. With their extra "layer" of indirection/virtualization, they add complexity and overhead to the current model, usually with unclear security properties. Also, they sometimes do not scale as intended.

In this sense, the most promising resolution method adopted to mitigate the expected overhead due to the addition of our incremental naming overlay, that require to be flat, is through DHTs. This fact alone is responsible for the rising of many questions, like the type of DHT to be used, its resolution algorithm and, more importantly the economical viability of the maintenance (or rewards) of the participating nodes of a DHT.

In our case, the prominent DHT candidate was the OpenDHT [8] that is widely distributed throughout the PlanetLab [9] nodes and presents a relatively low latency name resolution whose worst case is O(log n) for a space of n keys. In addition to that, many open questions will yet remain, like the heavy dependence on caching to mitigate the deleterious effects of the requirement of two or three queries for the naming infrastructure per connection, and possibly, a need for a certificate authority for the names in the namespace.

III. Naming in the Current Internet

The current hierarchical Internet naming infrastructure, based on the DNS, has performed its function with distinction, much beyond the intents of its original designers. However, new system requirements, such as mobility in ubiquitous computing and middleboxes [3] emerged (e.g. transparent caches, firewalls, NAT servers, content proxies and so on), violating the original Internet architecture premises in various ways.

With DNS, a global hierarchical namespace is employed to translate host-centric names to IP addresses, what clearly does not leave enough room to consistent identification of service and users. Besides, hierarchical namespaces pose a challenge to mobility as they embed intrinsic network location characteristics to nodes (or data itself) network attachment point [4], hindering their mobility tracking capability.

The incessant search for suitable fashionable domain names is a common place and is becoming a bottle neck stressing the DNS structure in complex ways [13]. Besides, social and political problems such as ICANN principles, "name squatting", trademark infringement and reverse infringement are problems stringent related to human friendly-names.

Even worse, middleboxes disrupt the original end-to-end connectivity model by getting itself involved in internal (in this context internal should be understood as confined in a single administrative realm) forward decisions to specific end-nodes (e.g. source route or loose source route activities). Finally, the algorithms that drive these decisions are arbitrary and system-specific, a fact that compromises the fundamental philosophy of the Internet of being based on open and public standards.

IV. The Flat Layered Naming Architecture

This proposed Layered Naming Architecture based on a Flat Namespace was primarily motivated by the need of an assortment of enabling technologies from the seminal work on the network architecture proposed by Saltzer [4]. Such work can be seen as an attempt to bring the OSI feature richness to a generic network infrastructure. The following objects were there defined: Data, Service and Users, Nodes, Network Attachment Points and Paths [4]. In the particular case of Data/Service/Users and nodes, three main attributes should be associated to it according to Shoch [15]: a name, an address and a route. While the name identifies the object itself, the address will identify its location on the network (which is subject to change) and a route will identify the concatenation of paths to reach the object (also subject to change).

The work of Shoch and Saltzer [4] laid the groundwork for the enumeration of the directing principles for a layered naming architecture presented by Stoica et al [11]. Since this work is mainly a well founded case for the use of a layered naming infrastructure it is the main guideline for our proposal that is basically formed by three directing principles.

The first principle will be the restriction of its scope to the univocal naming resolution, with the insertion of its components in the "metadata path" only based on persistent names (in our case represented by a 20 byte hexadecimal numbers). As a direct consequence, the second one states that changes in the routing infrastructure are not required, being this new proposed system intended to work vis-à-vis with the current naming resolution infrastructure [5].

The third directing principle aims to a diligent analysis of the current naming system proposals, followed by a judicious effort to combine existing mutual complimentary solutions and aspects of it rather than proposing and/or coding a whole new protocol or application from scratch. The rationale behind this is the fact that, while not as complete and decisive in the solution of the naming problem, each individual current proposal is mature and checked for primary inconsistencies.

In the same sense, while it implements a completely new architecture based in a novel naming architecture, it also lies on the current Internet to route the data and in the robustness and simplicity of DHT’s (Distributed Hash Table) to accommodate the new flat namespace resolution layers. Following the principles, our proposal will mainly focus on the creation of concepts and functionalities based on overlays protocols/applications.

Looking at these principles from the naming architecture perspective, the name (also called here as “reference” due to it human unfriendly nature and univocal identification), yet without any arbitrary restrictions should provide capability to either encode a chain of resolution on itself or at least be capable of providing any form of “hint” to the Resolution Reference Service (RRS) infrastructure, with the flexibility of changing of those information without disruption to network operations.

A. General Requirements

Based on the three principles for a layered naming architecture, and having in mind the typical requirements from the general concept of ubiquitous or pervasive computing a list of general requirements of a naming architecture can be drawn.

1st - Univocal and flexible identification of data/services/users, irrespective to its location on an administrative realm (domain), on a node (server) and network attachment point (IP address).

2nd - Univocal and flexible identification of a node irrespective to the services, data and users currently associated with that, to its location on an administrative realm and network attachment point.

3rd - The namespaces for identification of data/services/users and nodes must be non-hierarchical nor impose any restriction on the names, being it then semantic free namespaces.

4th - Ability to map human-readable namespaces to semantic free ones.

These requirements lead the architecture to the addition of two new flat naming layers to the current TCP/IP model and three levels of resolution. In Figure 2 it is illustrated these new resolutions naming mechanism: SID Resolution (Service Identifier) and EID Resolution (End Point Identifier).

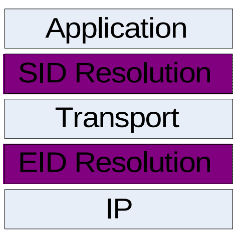

Figure 2 - The new naming Architecture layers [1]

These layers require more levels of identification, in replacement to the current single level of resolution from the Internet based on the DNS (that converts a name to its address). First, one must convert User Level Descriptor to Service Identifiers (1), then translating it to End Point Identifiers (2). Finally, it finishes by translating EIDs to their Network Attachment Point (2) or, in the Internet case to IP Address. These resolutions are better understood in the Figure Y.

Figure 3 - Architecture Reference Resolution

In order to check the feasibility of this proposal, a Java-based proof of concept was designed to accommodate the new flat naming infrastructure with the current hierarchical DNS. It is believed that an incremental adoption of this architecture, as a gradual replacement to the DNS represents the best alternative in face of the cost to establish the initial naming system setup covering the entire Internet.

B. The Proof of Concept

Our proposed Naming Architecture, in accordance with the principles and requisites supra cited, is responsible to map flat name identifiers to hierarchical names (e.g. an URI – Uniform Resource Identifier). A flat reference consists of a tag reference that points to an object (in other words, to meta-data associated with the object).

In this sense, content providers are responsible to insert an object’s metadata into the Architecture infrastructure and associate it with a tag. Consumers of the content submit these tags into the infrastructure and receive object meta-data in response. Our implementation uses the OpenDHT service from PlanetLab to store the flat objects and envisions that a Canonical Search Engine exists and translates the meaning full names (i.e. hierarchical) to the opaque references. The same as search engines (such as Google, Yahoo and so on) do today without any mandatory billing for users.

A simple resolution mechanism is illustrated in Figure 4 below. If a client that participates to the overlayed Internet (and, of course, it has configured its applications to use Scone Proxy) desires to access a web content (lines 1.1 and 1.2) it sends its requisition without regarding if it is a flat or hierarchical reference. Then, the RRS (Reference Resolution Service) resolves the references in 2.1 and 3.1 by using the OpenDHT infrastructure and by using the DNS in 2.2 and 3.2. The following steps in 4, 5 and 6 are common to both architectures.

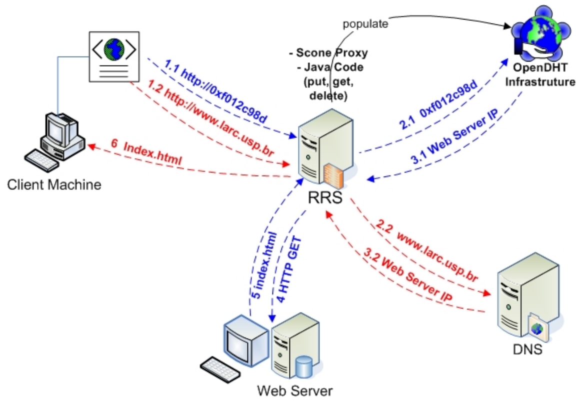

Figure 4 - Name Resolution in the Overlayed Internet

V. Conclusions

Overlay Networks provide a real opportunity for innovation as one can quickly deploy advanced and innovative capabilities by programming them to overlayed nodes of the substrate network without having to compete with the existing infrastructure. The main contribution of this naming architecture project is due to the compilation of several ideas we borrowed from the literature making available the final novel proposition.

The proposal of a Layered Naming Architecture for the Internet relying on Overlay Networks to deploy a novel flat namespaces and consequently requiring new RRS systems represents a little but essential changes. As well as, the introduction of a naming facility to identify univocally services and data, irrespective to its node, would have an even acute changing effect. When joint, these two relatively simple characteristics can have a worthy effect, improving the reliability and usage of the Internet, accelerating the development and introduction of new applications and enabling mobile computing.

Obviously, sensitive performance impacts due to the changes in the way of the data and metadata are handled by overlays can be foreseen. Consequently, exhaustive investigations of optimized capabilities and introduction of caches would probably push the architecture towards a possible standardization, which, indirectly will lead to implementations with higher degrees of reliability and performance.

References

[1] H. Balakrishnan, K. Lakshminarayanan, S. Ratnasamy, S. Shenker, I. Stoica, and M. Walfish, "A Layered Naming Architecture for the Internet," presented at SIGCOMM, Portland, OR, USA, 2004.

[2] I. Stoica, "Overlay Networks: Introduction and Multicast," vol. 2006, 2003, pp. Presentation.

[3] L. Peterson, T. Anderson, D. Culler, and T. Roscoe, "A Blueprint for Introducing Disruptive Technology into the Internet," presented at First ACM Workshop on Hot Topics in Networks (HotNets-I), Princeton, NJ, USA, 2002.

[4] J. Saltzer, "RFC 1498 - On the naming and binding of network destinations," M.I.T. Laboratory for Computer Science August 1993.

[5] M. Buddhikot, G. Chandranmenon, S. Han, Y. W. Lee, S. Miller, and L. Salgarelli, "Integration of 802.11 and third-generation wireless data networks," *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies. IEEE*, vol. 1.

[6] S. Chakrabarti, "Bluetooth Scatternet Formation and Internetworking with 802.11 and GPRS," The University of British Columbia, 2004.

[7] J. Quaini-Sousa, M. A. d. L. e. Silva, F. F. Redigolo, T. C. M. Carvalho, H. C. Guardia, W. V. Ruggiero, and B. Ohlman, "Suitability of Overlays as a General-Purpose Data Communication Substrate on Gigabit Channels," presented at 18th IASTED International Conference on Parallel and Distributed Computing and Systems (PDCS 2006), Dallas, Texas, USA, 2006.

[8] C. Tschudin and R. Gold, "Network Pointers," presented at 1st ACM Workshop on Hot Topics in Networks, Princeton, NJ, 2002.

[9] P. Nikander, J. Ylitalo, and J. Wall, "Integrating Security, Mobility, and Multi-Homing in a HIP Way," presented at Network and Distributed Systems Security Symposium, San Diego, CA, USA, 2003.

[10] P. Nikander, J. Arkko, and B. Ohlman, "Host Identity Indirection Infrastructure (Hi3)," presented at The Second Swedish National Computer Networking Workshop 2004 (SNCNW2004), 2004.

[11] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana, "Internet Indirection Infrastructure," presented at ACM SIGCOMM, 2002.

[12] F. F. Redigolo, J. Quaini-Sousa, M. A. L. e. Silva, T. C. Carvalho, W. V. Ruggiero, and B. Ohlman, "Evaluating Overlays for Gigabit Channels Data Communication," presented at International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities, Orlando, Florida, USA, 2007.

[13] M. Walfisha, H. Balakrishnana, and S. Shenkerb, "Untangling the Web from DNS," presented at 1st USENIX/ACM Symposium on Networked Systems Design and Implementation (NSDI '04), San Francisco, CA, USA, 2004.

[14] R. Moskowitz and P. Niklander, "Host Identity Protocol Architecture," January 2004.

[15] J. Shoch, "Inter-Network Naming, Addressing, and Routing," presented at IEEE COMPCON, 1979.