

Dagstuhl Seminar on Naming and Addressing for Next Generation Internetworks

Thomas R. Henderson, Andrei Gurtov, Lars Eggert, and Christian Dannewitz

ABSTRACT

The design of naming and addressing for data networks is a fundamental architectural consideration, and several current or anticipated problems in the Internet – including mobility dynamics, forwarding table growth in the core routers, and security – point out possible limitations with naming and addressing schemes in use today. A seminar on the topic of naming and addressing for next generation internetworks was held at the Schloß Dagstuhl from October 29 to November 1, 2006. Researchers from different fields discussed their views and recent results pertaining to naming and addressing problems. Over twenty talks covered topics such as routing, naming components, APIs, mobility, delay-tolerant architectures, flat routing and deployment issues. This article briefly summarizes the seminar presentations and discussions.

Categories and Subject Descriptors

C.2.1 [Computer Communication Networks]: Network Architecture and Design.

Keywords

Network architecture, scalability, mobility, heterogeneity, extensibility, naming, addressing.

1. INTRODUCTION

The basic principles of the original Internet architecture include end-to-end addressing, global routability and a single namespace of IP addresses that serve simultaneously as locators and host identifiers. A second hierarchical namespace of fully qualified domain names was later added and the domain name system (DNS) was developed to map between such names and addresses. Over time, a number of additional namespaces have emerged, many of which include partial domain names and are also served by the DNS. The DNS uses caching to improve scalability and is consequently not optimized for rapid updating of records.

However, due to the growing trend towards mobility of users, terminals or even whole networks, dynamic naming structures are gradually replacing the static mechanisms of the traditional Internet. Meanwhile, the emergence of network address translation (NAT) has clouded the end-to-end significance of IP addresses. Moreover, during the evolution of the Internet from a small research network to a worldwide information exchange, a growing

This seminar was supported by the International Conference and Research Center for Computer Science, Schloß Dagstuhl in Wadern, Germany, <http://www.dagstuhl.de>.

Dagstuhl Seminar Proceedings 06441
Naming and Addressing for Next-Generation Internetworks
<http://drops.dagstuhl.de/opus/volltexte/2007/1129>

diversity of commercial, social, ethnic, and governmental interests has led to increasingly conflicting requirements among the competing stakeholders. These conflicts create tensions that the original Internet architecture struggles to withstand. Clark *et al.* refer to this development as “tussles in cyberspace” [1].

This evolution has prompted research into different internetworking architectures, such as *FARA*, *Plutarch*, *Triad*, *i3*, *SNF*, *TurfNet*, *DOA*, *IPNL* and *4+4*, among others. At the core of these next-generation network architectures are naming and addressing frameworks that are significantly more flexible, expressive, and comprehensive than the Internet’s hierarchical IP address space. These naming frameworks are key components that enable advanced internetworking capabilities, such as multi-homed mobility, dynamic composition of networks, or delay and disruption-tolerant communication.

The naming architectures of these new internetworking architectures frequently have a more formal framework for naming than do current networks. Many naming architectures provide *dynamic bindings* between the levels of names and objects. With dynamic bindings at multiple levels, names of objects can become location independent and support different types of mobility, e.g., nodes or services. Some naming architectures also support the notion of *indirection* or *delegation*.

To discuss these issues and to advance this field of research, a seminar on naming and addressing for next generation internetworks was held at the Schloß Dagstuhl from October 29 to November 1, 2006 [2]. Researchers from different fields discussed their views and recent results pertaining to naming and addressing problems related to the seminar topic. This article briefly reviews their presentations and discussions, as well as the research questions identified and debated at the seminar.

2. SEMINAR TOPICS

The seminar brought together a diverse community of researchers from academia and industry, with different research interests including network theory, mobile networks, interdomain routing, networking in challenging environments, privacy, deployment, and peer-to-peer aspects of networking.

2.1 Background

The first session, by Bengt Ahlgren, was a review of a seminal paper by Jerome Saltzer on the role of names and addresses in networks [3]. It became quickly apparent that many participants had different definitions for key terms, including *name*, *address*, and *identifier*. During the following discussion, it was suggested that a general axiomatic basis of networking was needed to make progress on terminology. Later in the seminar, a presentation by Martin Karsten on “An Axiomatic Basis for Communication” [4], attempted to address this point. His presentation observed that the

original Saltzer paper did not cover broadcast and anycast network names/addresses and that the interesting aspect in naming and addressing may be more in how one performs the allocations than in the naming semantics and bindings themselves.

As additional background material, Kevin Fall presented discussions from a workshop [5] conducted by the Internet Architecture Board (IAB) two weeks before the Dagstuhl seminar. The main topics discussed at that workshop were the scaling trends of core router forwarding tables and the potential of the identifier/locator split.

2.2 On What to Name

One repeating topic of several presentations was the fundamental question of “what should be named in a computer network?”

Ken Calvert briefed the participants on a new NSF FIND effort called the “Postmodern Internet Architecture Project” [6]. In this clean-slate design, the project anticipates all foreseeable policy goals affecting the network layer with an architecture not requiring deep packet inspection in the forwarding path. The design also achieves separation of concerns by isolating the basic forwarding mechanism from any kind of endpoint identifier, isolating infrastructure from hierarchical, topology-based identifiers, separating path determination from forwarding, and providing customer-provider relationships independent of topology. These assumptions have led to a novel network-layer header including generic fields for forwarding directives (“where to send packet”), motivation (“why”), accountability (“who”), knobs (“advice to network layer on how to send packet”) and dials (“advice to transport/application from below”). In this architecture, every link (not endpoint) has a globally unique identifier and nodes become anonymous interconnections of links; this practice is thought to allow abstraction without naming the aggregated entity. This architecture provoked a lot of discussion and it was noted that a basic premise of this proposed architecture is that inter-AS routing is the problematic area of the current architecture, which may not necessarily be the case.

Anthony Joseph gave a brief talk on naming when people are the endpoints. His perspective is that although there are many networks and devices, convergent services with people as the endpoints is what most users care about. A problem for many people is that they have many devices with little control over them and over how they relate to one another. Current management practice is to maintain a directory, such as an LDAP directory, to map a person to multiple devices, but it is then up to the person initiating the contact to select the appropriate endpoint. Anthony Joseph argued that networks should be designed to return control to the callees, allowing them to specify when and how they can be reached. Although a unique identifier for people is probably important, many issues remain, such as cooperation between competing parties, privacy and location information as well as routing aspects at the service layer. A basic question is: “what support do we need in the networks to support context aware, ‘intelligent’ control of endpoints?”

Börje Ohlman reflected on the layered naming architecture [7] and argued that networks ought to name attachment points, nodes, data and services. He suggested the use of a flat, cryptographic namespace for most of these, with the possible exception of network attachment points, which could perhaps be hierarchical for performance reasons [8]. He identified a basic research issue of

how to bind human-readable names to cryptographic names and suggested that cryptographic names might fundamentally change some assumptions about network architecture, because it would no longer matter from where someone obtained the data.

Teemu Koponen briefed the seminar on a data-oriented network architecture (“DONA”) [9]. The premise of this architecture is that data retrieval is the fundamental operation in networks and that network naming systems should be oriented to supporting names for data. DONA builds on two basic primitives sitting directly on IP: “find,” by which a client requests a piece of data by its name (not location), and “register,” by which a host offers to serve a particular piece of data.

2.3 Application Programming Interfaces

Application Programming Interfaces (APIs) are related to the question of what to name, because they specify the names that users must use at system interfaces. A talk by Saleem Bhatti on “What’s in a name?” focused on the fact that the sockets API continues to be the primary means by which applications use the network, and that experiments with new namespaces are difficult. He asked whether a new API for user-space applications, possibly based on domain names and service names, might provide more flexibility for stack engineering. Clearly, a challenge here is to agree on the correct abstraction, syntax and semantics for such a new API.

Miika Komu also suggested that part of the emerging HIP architecture might be the use of an “endpoint descriptor” at the sockets API [10]. The endpoint name would be used in the stack above the network layer and would be resolved into a network-layer identifier (IP address) by a shim sublayer operating just above the IP layer. A benefit to this approach is to provide separation of locator and identifier so that the binding between the node and node name at the sockets API can be stronger and more direct. Some costs identified were the cost of resolution, and the costs of porting applications to use the sockets API extension.

One interesting discussion point was that in the present Internet, Google has almost become the *de facto* name resolution service and the main names that matter are the search strings. This was also a viewpoint expressed in a presentation about the Autonomic Network Architecture (ANA) project [11], where it was observed that addresses are increasingly short-lived entities in the network, and that the trend is towards naming becoming a search activity (via Google) instead of lookup-based activity. While some participants considered that Google might provide the universal naming “band-aid” for users, others remarked that there are still a lot of problems with Google as a universal naming service, including the human effort needed to sift through results, the fact that the majority of the population outside of computer science does not have much searchable data stored by Google, and that there remain issues with updating, trust models, and revocation of naming data found in Google’s databases. Christian Tschudin of the ANA project hypothesized that a future naming architecture might consist of short-lived “labels,” partly human readable but non-persistent (searchable) “attribute sets,” and underlying persistent cryptographic IDs, with multiple competing mapping schemes that integrate resource discovery with routing.

2.4 Mobility

Mobility stresses naming and addressing systems in that mobile devices may acquire new names or addresses as they change attachment points to the network, and the bindings may need to be updated. Two presentations in the seminar specifically addressed the challenges that mobility offers to naming and addressing architectures, as well as allocation schemes.

Raquel Morera described an architecture for mobile hosts based on the decoupling of names and addresses for mobile hosts [12]. She observed that existing server-based solutions for name to address resolution (DNS, SIP registries) do not scale or perform well for dynamic ad-hoc networks that may partition or suffer from intermittent connectivity. She then described a hierarchical, multi-domain system consisting of Logical Name Servers (LNS), optimized for mobile environments. Mobile nodes register with topologically relevant LNS servers that map from logical name to current address and that inherit some properties from DNS and SIP servers. The LNS server can forward the binding information to an LNS server in the mobile node's home domain. One remaining question is whether name and address decoupling is suitable for highly mobile and dynamic networks.

James Kempf raised an interesting question: "Is mobility an incompatible architectural challenge to IP Routing?" He surveyed a number of proposals for mobility in IP networks, including Mobile IP, HIP, and GPRS, and identified limitations with each of them. He observed that current solutions enforce disaggregation of routes in the middle of the network, causing scalability, reliability, and resilience problems for network operators, and poor traffic engineering consequences. He then explored whether the emerging field of network coding held any promise, and observed that it was fundamentally a new forwarding but not routing paradigm, and that it did not obviate the need for a routing protocol. Kempf concluded that prospects were gloomy for a unified architectural approach for mobility that was simultaneously simple, secure, easy to deploy, transparent to the end user, and had a clean interaction with IP routing [13].

2.5 Delay-Tolerant Architectures

Delay-Tolerant Networking (DTN) refers to a new class of networks that fundamentally operates on a store-and-forward message paradigm, rather than a best-effort packet delivery service. In the course of defining DTN architectures, researchers have had to tackle fundamental naming and addressing issues for the DTN approach, including the interconnection of disparate networks such as deep-space networks and the traditional Internet. Two presentations in the seminar were principally concerned with DTN and its impact on naming and addressing.

Jörg Ott outlined the overall DTN architecture [14], as defined by the IRTF DTN research group, for supporting asynchronous interconnection of different networks. DTNs operate as overlays on existing networks, exchanging messages (bundles) with other nodes. In DTNs, names are equivalent to addresses, are based on the Uniform Resource Identifier (URI) format, and agents send bundles to endpoint identifiers. Routing takes place on the entire name, and the architectural principle of late binding is used to avoid the need for address resolution at the sender. Ott observed that DTNs are a very green field environment for research, and

lots of elements are left to be defined including security and scalability.

The Hagggle project [15], presented by James Scott, is another clean-slate redesign of a mobile node architecture. The architectural motivation is based on the observation that traditional networking stacks embed names in packets, from application to link layer, and force end nodes to be able to resolve them each step of the way. This works well for well-connected networks but causes problems in disconnected or asynchronously-connected networks. In Hagggle, the API is asynchronous (as in a DTN) and data-centric with store-and-forward forwarding paradigm. Elements of Hagggle include Data Objects (DOs), DO filters, Name Objects, and Forwarding Objects. Names and links between names are obtained from various sources, and are used as destinations for sent data. Hagggle performs routing on several names in the name graph using their respective multiple mechanisms in parallel. Scott described that the point of this architecture is to provide naming beyond traditional stack-based architectures, and instead more like directed acyclic graphs. This type of architecture is flexible for DTN but it is an open question how useful it would be on a well-connected Internet.

2.6 Routing

Anja Feldmann gave an overview of the current state of BGP routing. BGP facilitates inter-domain routing through customer-provider and peering relationships between Autonomous Systems (AS). The problems of route flapping and routing table growth endanger future Internet operation. Because of its importance in Internet routing, BGP is one of the most studied protocols. Unfortunately, BGP enhancements developed by academia are challenging to experiment with and deploy because they may pose a risk to the operations or business models of in-service networks.

Thomas Fuhrmann presented some results on the use of flat identifiers at the network layer, instead of using them in an overlay as is common in peer-to-peer networks [16]. Scalable Source Routing combines source routing with Chord-like hash-based routing to achieve scalability with small per-node forwarding tables. The architecture trades path length and proactive responses to mobility to reach its scalable routing goals. The work is interesting in that it challenges the view that network-layer identifiers need to be hierarchical to achieve global scalability.

2.7 Deployment

Deployment of novel architectures, even for experimental purposes, is a challenge, and Tom Henderson presented some perspectives on the Host Identity Protocol (HIP) experiment being carried out in the IETF and IRTF [17]. HIP has been touted as a possible solution to several architectural stresses in the Internet, including BGP routing table growth, mobility and multihoming, and security, and the HIP experiment presently is trying to transition to larger scale experiments to evaluate the benefits and costs of running HIP. Henderson suggested that HIP experimenters focus more energy on defining and conducting the experiments that will shed light on these larger deployment issues, as well as focusing on the incentives that might lead to more people joining the experiment.

At the end of our seminar, Jari Arkko and Marcelo Bagnulo presented a different perspective on the HIP experiment, with a pres-

entation entitled “The Identity/Locator Merge,” suggesting that perhaps cryptographically generated addresses (CGAs) could be used to add cryptographically-secure identifiers into IPv6 addresses, and avoid the deployment costs of a resolution system for locators and addresses.

Richard Gold presented an architecture based on “network pointers” that can be thought of as an indirection service offered by the network layer [18]. The focus of his presentation was on how to deploy such a service without affecting backward compatibility with IP-based applications. The options considered were overlays, underlays (e.g., MPLS), IP via middleboxes, and name system interposition (modified *libc*). He described an implementation of a “netbox” that intercepted packets causing the network layer to believe that every destination was a single hop away. He observed that implementing an overlay by means of an underlay was providing a form of address virtualization. He noted that a next step might be to modify *libc* to provide a name virtualization service, and suggested that such a system, while hard to realize, would provide a nice research platform for new naming architectures.

2.8 Other Topics

An overview of the Ambient Networks project [19] was presented by Rolf Winter. Ambient Networks spans several of the above topic categories (mobility, interconnection of heterogeneous realms, ID/locator split). The Ambient Networks architecture assumes that multiple, disparate “locator domains” are tied together by core networks that are statically interconnected. Nodes have a globally unique (cryptographic) identifier, but may have locators of different types. Within a locator domain, routing can be performed based on locators, but between such domains, routing consists of forwarding to a core node, where a central naming service (based on distributed hash tables) keeps the current dynamic binding between a name and a locator domain. Questions on this architecture centered around whether we can assume a flat routing scheme in the core based on distributed hash tables.

Janne Lindqvist introduced some problems relating to privacy and security of naming systems. He then described a privacy- and security-enhancing architectural concept based on statistically unique and changing MAC addresses, IP addresses (interface identifier part), and anonymous HIP identifiers. Such an architecture would potentially be more resistant to eavesdropping and traffic analysis, but imposes some costs in terms of identifier management and protocol overhead to synchronize these changes. It was also not clear how such an architecture would interact with legacy authentication systems that desire a stable identifier.

James Sterbenz presented a view on the changing role of the Internet Protocol model entitled “On the Death of the Hourglass.” In the traditional hourglass model, IP serves as a waist of the TCP/IP protocol stack with a diverse set of protocols beneath it (Ethernet, Token Ring, Token Bus, etc.) and above it (TCP, UDP, HTTP, SMTP, SNMP, etc.). The current situation is different in that the Ethernet-family protocol is by far the most used link layer; most user applications run on top of HTTP over TCP. Furthermore, the co-existence of IPv4 and IPv6 as well as diverse architectural hacks including NATs, traffic engineering tools and caches does not match the hourglass model anymore. A discussion after the talk focused on whether HIP, with its capability to interoperate across IP versions, could become a new waist of the Internet.

3. CONCLUSIONS

The contributions and discussions echoed many similar concepts during the workshop, including cryptographic names, flat names, search as a naming mechanism, DTNs, DHTs, etc. Hence, an important question is: “Are we all working on similar things and is the packaging and terminology an issue?” It appears that a commonly accepted terminology is missing for naming and addressing. There have been several attempts to define a common terminology for addressing and naming, but the results of the workshop suggest that those approaches are not sufficient and more general evaluation of the terminology problems is needed. To start with, a clear problem description for naming and addressing issues would help. For example, it is not clear whether the problems are more about performance and efficiency or about providing new capabilities. Furthermore, additional research has to evaluate which features and characteristics the current APIs are missing that hinder users from offering services with more flexible naming.

Part of the problem with naming seems to be the lack of hard requirements, stemming, among others, from a fuzzy understanding of the naming and addressing infrastructure users. Hence, there is a clear need to better define the problem space and resulting requirements. Questions like “what is being named?”, “what are connectivity properties of the network?”, “what must it interoperate with?”, “what are the service models?”, and “what infrastructure is needed?” have to be addressed. Participants would also like to see boundaries and tradeoffs in a more analytical and rigorous manner, taking aspects like security, heterogeneity, and mobility into account.

But where are the *specific* problems in generalizing naming or in finding axiomatic approaches to naming? Routing seems to be one problem, as it is hard to separate from naming and addressing. Furthermore, there are three different groups of stakeholders – operators, terminals and users/services – and it seems to be hard to find a common ground between them. A general understanding of the specific problems is still missing, though. Mobility is a good example; it is an intensively studied aspect in the context of naming and addressing. Yet, it is not clear if the lack of a clear solution for IP mobility is a result of the architecture or a fundamental problem with mobility.

With respect to some of the perceived fundamental problems, there are nevertheless some promising, but not fully tested solutions, such as HIP. Hence, future work should focus more on deployment and implementation to gain more experience with those solutions.

Moreover, short-lived labels and multiple coexisting namespaces are promising areas that should be further investigated. If namespaces are added to the current architecture, implications on the search, resolution, security, and routing mechanisms have to be considered, resulting in additional engineering and operating costs that have to be taken into account as well.

Current approaches place lots of focus on the network layer. The strong impact that Google and search engines in general have had on the way that people think about names and networking APIs implies that we may need to be more radical in thinking beyond the network layer.

Participants agreed that a follow-up workshop in about 18 months to discuss those topics and other advances would be useful.

ACKNOWLEDGMENTS

Lars Eggert initiated the Seminar and Bengt Ahlgren, Anja Feldmann, Andrei Gurtov, and Tom Henderson served as co-organizers. We are grateful to the Dagstuhl Foundation for providing a venue and support for the Seminar.

REFERENCES

- [1] David Clark, John Wroclawski, Karen Sollins and Robert Braden. Tussle in Cyberspace: Defining Tomorrow's Internet. Proc. *ACM SIGCOMM*, Pittsburgh, PA, USA, August 19-23, 2002, pp. 347-356.
- [2] Dagstuhl Seminar 0644. <http://www.dagstuhl.de/06441/>
- [3] Jerome Saltzer. On the Naming and Binding of Network Destinations. RFC 1498, August 1993.
- [4] Martin Karsten, S. Keshav and Sanjiva Prasad. An Axiomatic Basis for Communication. Proc. *Fifth Workshop on Hot Topics in Networks (HotNets V)*, Irvine, CA, USA, October 2006.
- [5] IAB Routing and Addressing Workshop. <http://www.iab.org/about/workshops/routingandaddressing/>
- [6] Bobby Bhattacharjee, Ken Calvert, Jim Griffioen, Neil Spring and James Sterbenz. Postmodern Internetwork Architecture. Technical Report ITTC-FY2006-TR-45030-01, The University of Kansas, 2335 Irving Hill Road, Lawrence, KS 66045-7612, February 2006.
- [7] Hari Balakrishnan, Karthik Lakshminarayanan, Sylvia Ratnasamy, Scott Shenker, Ion Stoica and Michael Walfish. A Layered Naming Architecture for the Internet. Proc. *ACM SIGCOMM*, Portland, Oregon, USA, August 30 - September 3, 2004, pp. 343-352.
- [8] Pekka Nikander, Jari Arkko and Börje Ohlman. Host Identity Indirection Infrastructure (Hi³). Proc. *Second Swedish National Computer Networking Workshop (SNCNW2004)*, Karlstad University, Karlstad, Sweden, November 23-24, 2004.
- [9] Data-Oriented Network Architecture. <http://radlab.cs.berkeley.edu/wiki/DONA>
- [10] Miika Komu, Sasu Tarkoma, Jaakko Kangasharju and Andrei Gurtov. Applying a Cryptographic Namespace to Applications. Proc. *ACM Workshop on the Dynamic Interconnection of Networks (DIN)*, Cologne, Germany, September 2, 2005.
- [11] Autonomic Network Architecture Project. <http://www.ana-project.org>
- [12] Anthony McAuley and Raquel Morera. LNS-SID Mobility Management in Dynamic Ad Hoc Networks. Proc. *IEEE VTC (Fall)*, Orlando, FL, USA, October 2003
- [13] Phil Roberts and James Kempf. Mobility Architecture for the Global Internet. Proc. *ACM/IEEE Workshop on Mobility in the Evolving Internet Architecture (MobiArch)*, San Francisco, CA, November 2006.
- [14] Vinton Cerf, Scott Burleigh, Robert Durst, Kevin Fall, Adrian Hooke, Keith Scott, Leigh Torgerson and Howard Weiss. Delay-Tolerant Networking Architecture. Internet Draft draft-irtf-dtnrg-arch-08 (Work in Progress), December 2006.
- [15] James Scott, Pan Hui, Jon Crowcroft, and Christophe Diot. Haggie: A Networking Architecture Designed Around Mobile Users. Proc. *IFIP Conference on Wireless On-Demand Network Systems and Services (WONS)*, Les Menuires, France, January 2006.
- [16] Thomas Fuhrmann, Pengfei Di, Kendy Kutzner, and Curt Cramer. Pushing Chord into the Underlay: Scalable Routing for Hybrid MANETs. Technical Report 2006-12, Universität Karlsruhe (TH), Fakultät für Informatik, Germany, June 2006.
- [17] Tom Henderson and Andrei Gurtov. HIP Experiment Report. Internet Draft draft-irtf-hip-experiment-03 (Work in Progress), March 2007.
- [18] C. Tschudin and R. Gold. Network Pointers. *ACM Sigcomm Comput. Commun. Rev.*, Vol. 33, No. 1, 2003, pp. 23-28.
- [19] Ambient Networks. <http://www.ambient-networks.org/>