IAB workshop at Prague
Position Paper
Intended status: Informational
Expires: March 25, 2011

W. Chun
ETRI, Korea
February 10, 2011

Domain-Insulated Autonomous Network Architecture (DIANA)

Status of this Memo

This document is submitted as position paper for the workshop held in Prague, March 25, 2011.

Abstract

Although the Internet achieved tremendous success in the last decades, several limitations are newly emerged due to recent technical advancement and challenging requirements of diverse applications. For overcoming these limitations, the future Internet can be redesigned in the "clean slate" basis. This paper proposes a novel architecture that considers both horizontal and vertical separations called, Domain-Insulated Autonomous Network architecture (DIANA). In this architecture, a network is composed of a set of basic building blocks, called "domains," which can be defined recursively by topological, administrational, or operational grouping of communicating entities. In this paper, the design principles of DIANA and its distinguishing features are described. Then a typical scenario of communication on DIANA is described along with some use cases of DIANA

Table of Contents

1. Introduction

As a design philosophy of the future Internet architecture, this paper proposes several distinguishing principles: First, it proposes simple but powerful abstraction where all entities, once plugged into a network bus, can communicate with each other by location-independent globally unique ID. Second, the underlying network is composed of heterogeneous but autonomous network segments which are collaborating with one another. The network segments, called domains, can be composed using parent-child or peer relations to build lager networks. Finally, on the hierarchically structured domains, a scalable and efficient routing and forwarding mechanism is provided. This paper is organized as follows: Section 2 describes the design principles according to which Domain-Insulated Autonomous Network Architecture (DIANA) is designed. DIANA's potential use cases are given in Section 3. Major contributions and future work are stated in the conclusion.

2. Design Principles of DIANA

As the traditional Internet was based on several principles, DIANA is also developed under the design principles as follows.

2.1. ID based bus abstraction

The Internet is often viewed as a network of heterogeneous networks. Diversity of technologies and logical groupings of communicating entities are conceptualized in a simple and consistent access interfaces. That means applications communicating through the Internet must not be aware of the structure of underlying networks. In DIANA, any entity, once plugged into the network, can communicate with other entities regardless of their locations just as hardware boards plugged into the back-plane bus of a computer system can communicate with one another. For this purpose, we introduce global unique ID by which every communicating entity can be identified. This abstraction serves as a pervasive narrow waist for diverse communicating entities, such as host, device, service, applications, and contents, as IP layer has done in the traditional Internet.

2.2. Holistic approach

Since ISO/OSI reference model was proposed, most of the communication network architectures were explained according to the layer model. The layer model is suitable for describing a communicating system but it only describes functional decomposition within a system. As the Internet is extending and getting bigger, it must be controlled and managed by horizontally decomposed regions, such as subnet, regional networks, and autonomous systems, and logical groups, such as social and peer-to-peer networks. In the current architecture, those horizontal decompositions have not been defined by architectural principles but by individual applications, such as routing and management.

A network as a whole can be segmented according to underlying technologies, control and management regions and logical groups. Segments of the network may compose a larger component, or a segment can be composed of a collection of smaller segments. Each segment is accessed by well-defined interfaces, either access interfaces or peering interfaces. In actual implementation, those interfaces might be implemented on one of the node in the domain, which is often called a "gateway". In the sense that communications is carried out through the gateways along the path, series of gateways may form a "chain of responsibility".

2.3. Domains as autonomous Entities

While communicating systems in the Internet are modelled in vertical layer architecture, network structure such as “site”, “regional”, “metro” networks are often used in network architectures. Traditionally, communicating systems are modelled in two different ways: via either structured protocols (OSI and Internet) or collection of nodes interconnected through reference points (3GPP). The current trends in communications such as P2P, cloud and grid computing are leading a paradigm shift, where a set of collaborating communication systems provides specific services while its internal details are hidden from the service users. This principle, often referred as the “black box principle,” which has been applied in the traditional vertical layer architecture, can also be applied in segments of a network.

Some previous works are based on horizontal network structuring principles; “Turf” in [3] and “context” in [4] are building blocks for the structured networks. “DIF (Distributed IPC Functions)” in [5], “compartments” in [6], and strata in [7] adapt similar concept in the sense that networks are composed of collaborating components, each of which covers a certain range of the network. In DIANA, the segment of the network can be modelled as a “domain” that consists of physical/logical nodes and/or child domains recursively, where nodes and child domains are collectively performing distributed functionality.

2.4. Union by federation

The current Internet unified different underlying network types by the same network and transport protocols (TCP/UDP and IP). The current Internet model is based on a single set of global IP addresses and homogeneous best-effort datagram delivery. This universal IP layer has been scaffold for Internet to grow to this global scale. Its fixed semantic, however, is not suitable for newly emerging networks, such as mobile and sensor networks. In the future Internet, a single universal protocol cannot support various subnet technologies and service requirements. Consequently, a new architecture that allows inter-operability among various heterogeneous networks without mandating “one-protocol-fits-all” is required. Thus, new architecture must support loose union of autonomous domains rather than unification by a single protocol.

2.5. Routing by ID

In the proposed architecture, every entity that communicates through the Internet must be assigned a “globally unique ID”. The primary role of ID is to identify an entity unambiguously. Referencing entities based on their name or ID has been widely used in computing systems, but it requires a resolution mechanism where a location- independent ID must be mapped into a location-based address. To circumvent scalability issues of this resolution, the current Internet adapted address as a function of their location in the network. This leads to serious problems in mobility and multi-homing.

DIANA proposes a routing mechanism based only on location-independent ID. As shown in Fig.3, source and destination IDs are carried in the invariant part of a packet. Whenever a packet arrives at the entry gateway of a domain, the gateway looks up its route cache with the destination ID as the key to find the exit gateway or the final destination. In contrast to the current Internet where all routing entries have to be stored in the routing table a priori, DIANA requires only relevant entries installed in the route cache.

2.1.6. Path discovery using routing hints

The first packet of a communication, for which a routing entry is not yet installed in the route cache, triggers the “path discovery” procedure. The path discovery procedure is the procedure for finding a path to a destination ID and storing the result in a route cache of each gateway along the path. In the sense that the routing path discovery is initiated by the first packet, this is called "reactive routing" in contrast to the proactive routing of the current Internet. In the path discovery phase, the location information about where the ID has been registered to may be used as routing hints.

3. Overview of DIANA

Based on the design principles listed in the above section, we introduces a new architecture called DIANA. In this section, we define relevant concepts and terminologies, building blocks and their federation rules. We also provide an example of typical scenario on DIANA.

3.1. Concepts and Terminology

- ID as a location-independent name

An ID in DIANA is defined to be location-independent, globally unique, and flat structured. An ID is assigned to every communicating entity to unambiguously identify the entity from others. Since ID is location independent and no structure is assumed for locating the entity, it can be viewed as a name rather than address. The IDs assigned to endpoints of communication are called EID (Endpoint ID). They are carried in the packet header. Thus, we assume ID has fixed- sized and machine-processable format as defined in UUID, GUID [8] or self-certifying names as in [9].

- Locator as a routing hint

In DIANA, one of the most distinguishing features is a locator which is defined as a hierarchical concatenation of domain names, starting from the top-level domain to the domain where the ID is actually registered. DIANA models networks in the form of a tree-structured domain hierarchy and all communicating entities with ID must be registered in one or more domains. Registration can be viewed as binding a given ID with one or more locators. Those locators will be used as routing hints in the path discovery phase. In order to keep routing independent from locating process, locators are referenced only in the path discovery phase. Identical copies of content or replicated servers may be assigned a single ID with multiple locators, and also in multi-homed domains, an ID may be bound to multiple locators.

- Address in a domain specific format

In DIANA, address is defined by domain specific semantics. Since domains are defined in various criteria and its operation is autonomous, addresses have meaning within a specific domain.

- Domains as basic building blocks

As layers have been serving as building blocks for constructing communicating systems, domains are the basic building blocks for structuring networks. In DIANA, domains can be defined by physical coverage of networks, such as local area networks, site and regional networks, and core networks or by administration and operational purposes. Domains may be aggregated into a larger domain by federating domains with either parent-child or peer-to-peer relations. Since domains are units of autonomous components, internal details are hidden from outside and domains are accessible only by well-defined interfaces, to which capabilities of the domains are exposed.

- Planes consisted of collaborating peers

The traditional communication architectures often introduced planes for defining related functional components and their interactions. Similarly DIANA defined planes as collaborating domains in peer relations.

3.2. Definition of Domain

As in Fig. 1, the abstract view of a domain can be defined as a collection of cooperative child domains, and domain-specific functional components. The functionality of a domain is provided via “access interface” as a service while the domain may collaborate with other domain via “peering interface” using agreed protocols. The collection of domains for providing specific functions forms planes such as service (transport in case of communication systems), control, management, and inference plane. Those planes can be extended toward outside of the domain via peering interface with appropriate translations.

Initially a domain may be instantiated by administrator or the parent domain using bootstrapping functions of the domain. While bootstrapping, the domain initialize itself and setup the access channel with its parent domain. And the domain might configure itself by explicit command via the access interface or autonomous neighbour discovery procedures. In the configuration phase, a domain may initialize one or more child domains as required in recursive manner. Once initialized and configured, a domain is assumed to have at least one access channel from its parent.

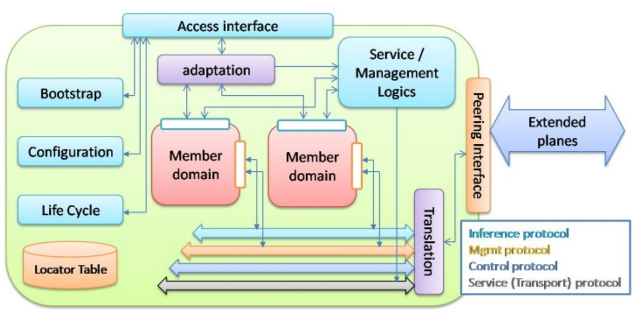

Fig. 1 Abstract View of a Domain

3.3. Composition rules

DIANA assumes that a network is composed of domains organized in parent-child or peer-to-peer relations. More specifically, more than one domain can be composed into a larger domain. As shown in Fig. 1, a domain exposes two interfaces: the access interface for vertical composition in parent-child relation and the peering interface for horizontal composition in peer-to-peer relation.

As shown in Fig. 2, vertically composed domains can be depicted in tree structured relations. The leaf domain that has no child domain is called an “atomic domain” and root of the domain tree is called the “top-level domain”. The locator of a given domain is represented by a path name starting from the domain itself to the top-level domain separated by “@”. As an example, in Fig. 2, the locator of domain A can be specified as “A@D3@D2@D1”.

As defined in the ID based bus abstraction, every communicating entity is assigned a globally unique and location-independent ID which has to be registered to one or more domains. Then, the ID is bound to the locator of the registering domain. By definition, all registered ID in a given domain are assumed to be visible from the parent domain. Thus, conceptually, all registered ID can be visible in top-level domain.

Horizontal composition of domains in peer-to-peer relation takes place through peering interfaces. Peering domains have to share common information via agreed interactions, which is said to be "protocols". Capabilities achieved by collaborating domains form planes, such as transport (or service), control and management, and inference plane. Protocols governing planes within a given domain are defined in domain-specific ways and planes can be extended to external domains by translation of those protocols. For example, the control plane of a local domain that uses local IP addresses can be extended to control plane of global domain by "Network Address Translation (NAT)".

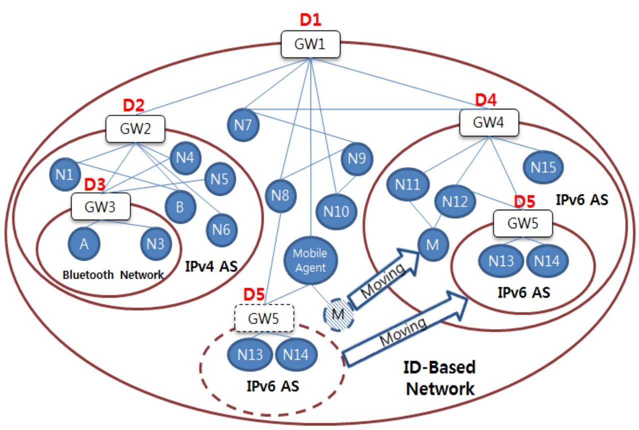

Fig. 2 An example of the DIANA network

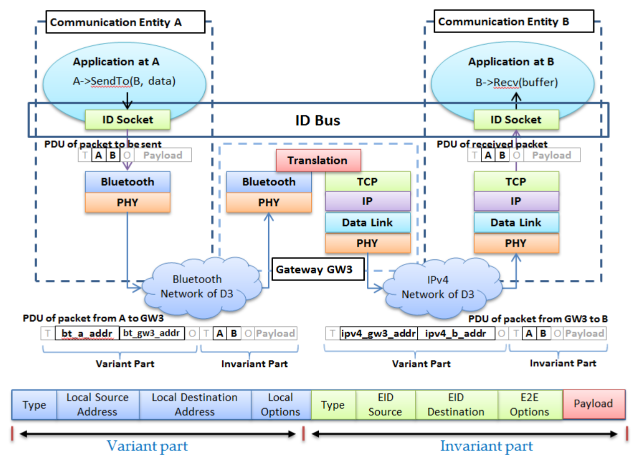

Fig. 3 A typical scenario of DIANA

3.4. A Typical scenario of communication on DIANA

The primary goals of DIANA are two-folds: One is providing the abstract ID based bus view to all communicating entities, whatever they are, either application processes, contents, servers, or even collections of servers. Note that passive entities, such as contents or data, must be registered to a domain that has an active component that can manipulate the passive components.

Fig. 3 is an example of a typical scenario of DIANA when a communication entity (CE) called A in Domain D3 wants to communicate with CE B in Domain D2 of Fig. 2. First, any CE must have an ID and must be registered in at least one domain prior to communication. For the provided scenario, A is registered to D3 and B to D2, assuming A and B are the IDs of CE A and CE B respectively.

When CE A tries to send data to CE B, CE A has to find the location of CE B. If CE A and B have communicated with each other in the past, the record, such as the route, may exist, and the existing record could be used. If not, CE A has to find the location of CE B, which is accomplished by finding the locator of CE B. The procedure for finding the locator of CE B is called the lookup propagation. The lookup propagation is a recursive procedure, in that CE A asks for the locator of CE B to its gateway; the gateway asks the gateway of its parent if it does not have the locator information for CE B; and the same process is repeated until the locator of CE B is found. Once the lookup propagation is complete, CE A will have the locator of CE B, which is “B@D2@D1” for the case in Fig. 2.

Based on the locator of CE B, CE A can determine the location of CE B. However, the routing path to CE B is still not defined. To determine the routing path, the path discovery procedure is used. The path discovery procedure builds entries about next hop information to CE B in the routing caches of all entities between CE A and CE B. Once the path is determined, then, finally CE A and CE B can communicate with each other.

An important aspect of DIANA is how the packet travels within or across domains. The ID socket of CE A creates the invariant part of a packet with the source, destination ID pair and the payload. Then the adaptation capability of the source domain (D3) encapsulates the invariant part of the packet using the domain specific protocol suit—the Bluetooth protocol for the example in Fig. 3. For an intra-domain communication, the packet will be received by the destination entity, and the communication is complete. But for an inter-domain communication, the packet will be sent to a gateway of the domain. The gateway translates the packet—replaces the variant part of the packet using a protocol suite for the domain that the packet will be traversing through. For example, in Fig. 3, the Bluetooth protocol specific variant part of the header is replaced with the IPv4 specific variant part.

4. Conclusions

One of the major contributions of DIANA is providing an integrated framework for diverse directions of future Internet researches. As an architectural framework for diverse services and heterogeneous networks, DIANA proposes a simple but powerful “abstraction of ID based bus” where entities with ID can communicate each other regardless of their locations.

Another contribution is the introduction of domains as building blocks for constructing much bigger networks. Most of the traditional network architectures have been based on functional layering that decomposes complex communication tasks into vertical stack of protocols. Layering architecture, however, may add unnecessary complexity to the system, and may cause scalability and inter-layer dependency problems. Meanwhile, network architects often view a network as a collection of autonomous network segments. Consequently, DIANA provides architectural framework for not only system designer but also network architects.

For scalable routing based on location-independent and flat ID, DIANA proposes two steps: registration for binding ID with locators and presence for acquiring domain specific addresses. In order to separate locating process from routing, DIANA introduces locators as routing hints and path discovery procedure using the routing hints. Path discovery, one of the most distinguishing features in DIANA, is usually triggered by the first packet of communication. It utilizes locators bound to the destination ID to find the optimal path and its results are stored in the route caches along the path, which are used for subsequent packets. So, routing in DIANA is performed in a “reactive” manner in contrast to the traditional “proactive” routing. Selecting the optimal path under given administrative policies in the path discovery phase is a challenging issue for further studies.

The domain principles can be applied for logical grouping such as VPN, social network, and cloud computing. With the help of the domain principle, logical domains are no longer collections of nodes and links as in traditional networks. Rather, they are defined as collections of entities and arbitrary relations, which may be mapped into paths along the physical domains when actually required. This mapping is another challenging research topic.

Since innovative approaches have started under the context of future Internet, many bright ideas have been proposed and experimented. However, those approaches pursue their own purposes and do not share common principles. As the traditional Internet has been taking a role of communication infrastructure in global scale, the future Internet also must be global infrastructure. In this sense, architectural framework that can be shared by diverse research areas in future Internet is strongly required. We hope DIANA would be one of alternatives to provide a simple but powerful framework for structuring the networks as well as developing individual communication systems.

5. References

[1] A. Feldman, Internet Clean-Slate Design: What and Why?, ACM SIGCOMM Computer Communications Review, Vol. 37, No. 3, July 2007, pp. 59-64

[2] D.D. Clark, The Design Philosophy of the DARPA Internet Protocols, ACM SIGCOMM Computer Communications Review, Vol. 18, No. 4, Aug. 1988, pp. 106-114

[3] S. Schmid, L. Eggert, M. Brunner, J. Quittek, Towards Autonomous Network Domains, INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Mar. 2005

[4] J. Crowcroft, S. Hand, R. Mortier, T. Roscoe, A. Warfield Plutarch: An Argument for Network Pluralism, ACM SIGCOMM 2003 workshops Karlsruhe, Germany, Aug. 2003

[5] J.Day, Patterns in Network Architecture: A Return to Fundamentals. Prentice Hall, 2008.

[6] A.Keller, T.Hossmann, M.May, G.Bouabene, C.Jelger, and C.Tschudin, “A System Architecture for Evolving Protocol Stacks”, in 17th International Conference on Computer Communications and Networks (ICCCN), Aug. 2008

[7] M. Johnson, J. Huusko, T. Frantti, F. Andersen, et, Towards a New Architectural Framework - The Nth Stratum Concept, 4th International ICST Mobile Multimedia Communications Conference, July 2008

[8] I. Hakala, M. Tikkakoski, From vertical to horizontal architecture – a cross-layer implementation in a sensor network node, InterSense ’06 Proc. Of 1st International Conference on Integrated Internet Adhoc and Sensor Networks, Nice France, May 2006

[9] D.Mazi`eres, M.Kaminsky, M.Kaashoek, and E.Witchel “Separating key management from file system security”, Operating Systems Review 34 (5): 124–139, Dec. 1999

Authors’ Addresses

Woojik Chun, Tea Sang Choi
Electronics and Telecommunications Research Institute
138, Gajeong-Ro, Yusong-Gu, Deajeon, 305-700, Korea
Email: wjchun@etri.re.kr