Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 4, 2018

P. Pillay-Esnault, Ed.
Huawei
D. Farinacci
lispers.net
C. Jacquenet
Orange
U. Chunduri, Ed.
Huawei
T. Herbert
Quantonium
D. Lake
Dell
M. Menth
U of Tuebingen
D. Raychaudhuri
Rutgers University
D. Meyer
July 3, 2017

Use Cases for Identity Enabled Networks
draft-padma-ideas-use-cases-01

Abstract

This document describes few deployment use cases for Identity enabled networks using a GeneRic Identity Services infrastructure.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 4, 2018.

Table of Contents

1. Introduction

The problem statement document for IDentity Enabled networkS (IDEAS) advocates a standardized Identity and mapping services infrastructure, secured access to that infrastructure with data integrity, different Identity(IDy) services and allowing for different Locator/Identifier data-plane solutions [IDEAS-PS].

This infrastructure is called GeneRic Identity Services (GRIDS). The GRIDS infrastructure (GRIDS-ARCH-TBD) is envisioned to support functionalities such as traditional mapping and resolution of Identifiers (IDfs), as well as secured Identity registration, subscription, privacy for Identity/Identifier data, discovery of Identifiers, updates to the system with data integrity. In addition, it is designed to allow flexible deployment with different scopes (local sub-domains/global).

This memo describes and focuses on few deployment use cases for Identity-based protocols that will benefit from such an infrastructure. The definition of and the need for Identity in IDEAS are described in a companion document [IDEAS-IDENTITY].

2. Specification of Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

3. Definition of Terms

This document makes use of the terms that have been defined in the problem statement draft of IDEAS [IDEAS-PS]. They are included here for reader's convenience. In case of any discrepancies between the two drafts, the problem statement draft overrides.

- Entity : An entity is a communication endpoint. It can be a device, a node, or a virtual machine (VM), that needs to be identified. An entity may have one or multiple identifiers (long-lived or ephemeral) simultaneously. An entity is reached by resolving one or more of its long-lived identifiers to its current locator(s).

- Identifier (IDf): denotes information to unambiguously identify an entity within a given scope (e.g. HIP HIT, LISP EID). There is no constraint on the format, obfuscation or routability of an IDy. The IDy may or may not be present in the packet whose format is defined by IDf-based protocols (HIP/LISP).

- Identifier-based (IDf-based): When an entity is only reachable through one or more communication access then a protocol or a solution is said to be IDf-based, if it uses an ID-LOC decoupling and a mapping system (MS) as base components of the architecture. Examples of IDf-based protocols are HIP, LISP and ILA.

- Identity (IDy): the essence of "being" of a specific entity. An identity is not to be confused with an identifier: while an identifier may be used to refer to an entity, an identifier's lifecycle is not necessarily tied to the lifecycle of the entity it is referencing. On the other hand, the identity's lifecycle is inherently tied to the lifecycle of the entity itself.

- IDentity Enabled Networks (IDEAS): IDEAS are networks that support the identifier/locator decoupling. Reaching an entity is achieved by the resolution of identifier(s) to locator(s).

- GeneRic Identity Services (GRIDS): GRIDS is a set of services to manage the lifecycle of IDy/IDfs, to map and resolve Identifiers and locators, and to associate metadata (META) with entities and entity collections. It is a distributed infrastructure that stores the IDy/IDf, the associated LOC(s), and metadata (META) in the form of tuples (IDy/IDf, LOC, and META).

- Locator (LOC): denotes information that is topology-dependent and which is used to forward packets to a given entity attached to a network (IPv4/IPv6 Address). An entity can be reached using one or multiple locators; these locators may have a limited validity lifetime.

- Metadata (META): is data associated with an Identity. The metadata may contain information such as the nature of the entity for example.

- User Equipment (UE): A user equipment is an entity per definition in [IDEAS-PS]

4. Use Cases for IDEAS/GRIDS

There are several benefits of various Locator/Identifier separator protocols that they bring to IP networking. These include but not limited to mobility with session continuity, global routable address space reduction, traffic engineering to name a few. The goal of this section is not to re-list all those deployment use cases of existing IDf/LOC protocols but specify what and how IDEAS/GRIDS can enhance the existing use cases and solve new problems in different areas.

This section details specific deployment use cases of IDEAS based upon GRIDS infrastructure.

4.1. Privacy with Access Control

For privacy or security reasons, an entity may only want a designated list of authorized entities to look up its locators and access it. To this end, the entity can specify an access control list in the GRIDS and let the GRIDS enforce the access control when entering the locator resolution phase. For example, when Alice wants to communicate with Bob, she has to look for the Bob's IDf in the GRIDS to get his locator. The GRIDS verifies Alice's IDf and checks the IDf against Bob's access control list. If Alice is authorized, the GRIDS returns Bob's up-to-date locators; otherwise the GRIDS returns an error or a honeypot LOC to analyze further (Figure 1).

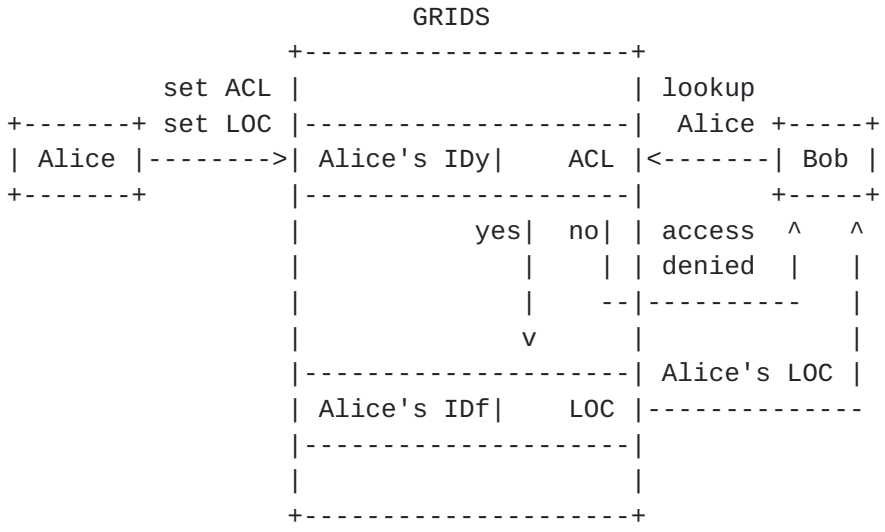

Figure 1: Authorizing access to LOC information

An access restriction policy can also be possible at the GRIDS level for a group/class of devices identified through the metadata of the entity. A specific example is restricting access for location resolution of a vehicular entity except for the authorized manufacturer or dealers.

The GRIDS Infrastructure supports the ability to verify policies associated to an Identity through an Identifier used in the data plane. This reverse lookup provision can then trigger the execution of various actions (e.g., discarding traffic or redirecting) on the

data plane nodes (routers/gateways/firewalls), according to the policies defined in IDEAS.

Receiver of the data traffic can also control anonymization through a specific policy specified by the GRIDS logic with its Identity. For example sender of the data traffic can be provided with one of the receiver's ephemeral Identifiers to use for communication, during LOC resolution response (here sender assumes to be doing LOC resolution through receiver's long-lived Identifier). Mechanisms like this can provide anonymization of the data traffic from eavesdroppers/outside observers.

4.2. Identity Services with global GRIDS

IDf-based protocols currently rely upon a customized mapping infrastructure and they do not interoperate with each other. Therefore, it would be difficult to deploy multiple IDf-based protocols in the same network due to the overhead generated by the operation of various mapping functions or it is difficult to deploy a global scope with greater flexibility. Hence, only one ID protocol is usually deployed even though there exists the need to deploy specific solutions to address various contexts.

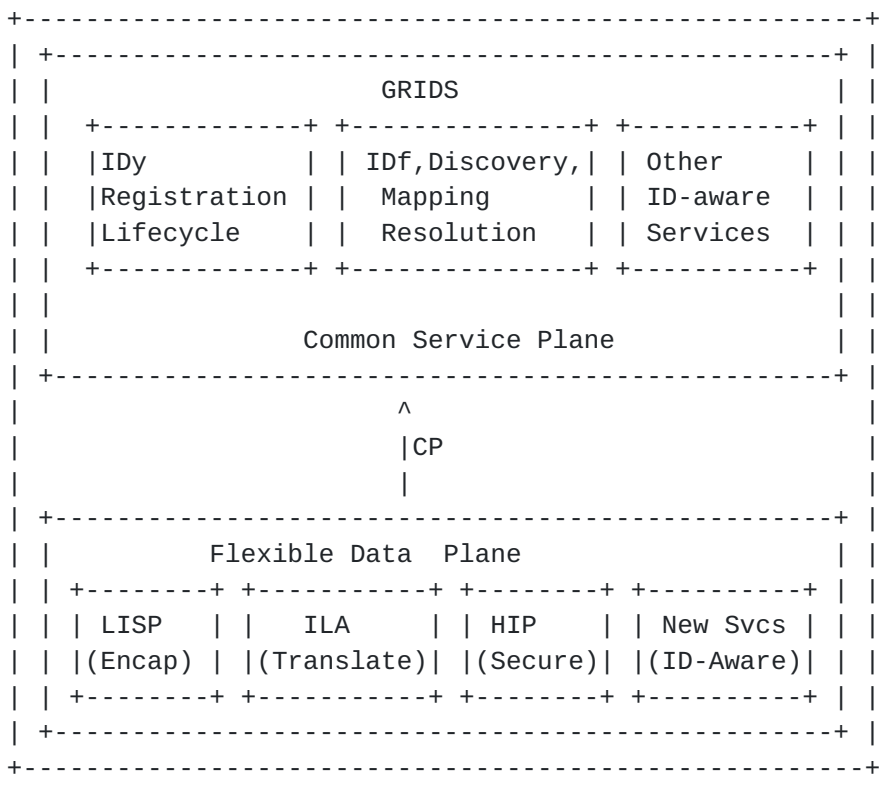

Figure 2: The GRIDS Structure

The above diagram Figure 2 shows a unified mapping system, besides the legacy mapping and resolution services. Other services Identity can potentially bring to the IDf/LOC protocols (access restrictions for resolution, discovery with metadata or grouping services) are also supported by the mapping system. Access security and registration services improve overall security of the GRIDS as these can provide authentication of the entity through a singular Identity. This process can also establish the entity type through metadata which is needed for some of the IDy services.

Control Plane (CP) options i.e., usage of existing LOC/IDf protocols or a new GRIDS-CP or usage of both, to realize services provided by GRIDS is discussed in [IDEAS-PS].

4.3. Identity Services with local GRIDS

The IoT landscape is comprised of a large variety of devices and protocols. IoT solutions cover a wide-range of protocols at varying stages of maturity, the latter may be at Layer-2/Layer-3 and usually within a single domain.

While the number of IoT use-cases is vast and diverse, there is a number of commonalities in network communications between classes of application. Some of the categories below can benefit from local GRIDS instance which can provide local authentication and security features.

Low-compute capability end devices generating small amounts of traffic with a long/short period, terminating on a central application service with no-loop/loop latency dependencies. These devices usually offload processing to a central service to which they connect. The central service may be an application hosting environment or a distributed/devolved solution (such as an Edge or FOG solution). In both cases, local GRIDS instances could be deployed for this functionality. Such devices will benefit from offloading the secure registration and access restriction policies placed at GRIDS.

Local GRIDS instance with standardized interfaces can also be deployed in enterprise networks to provide security and location based services. Here the policies applied on the Identity/Identifier can be more flexible with proprietary rules structure. These can be added on top of the minimal and standardized framework and interfaces to meet enterprise needs.

4.4. Mobility

One of the key benefits of any ID-Based protocols is its ability to provide mobility. Decoupling Identifier from location provides inherent support for mobility with session continuity and hence this document doesn't describe basic mobility use case further. But it is important to note availability of a global GRIDS Infrastructure would enable such mobility for entities connected to various IP networks securely. Other key attributes such as mapping and ID services system scalability, support of low latency and secure query (GRIDS-ARCH-TBD) are some of the aspects which can determine the adoption and deployment of this essential service.

Some use cases specific to mobile networks, where the need for both Local and Global GRIDS are specified in the Appendix.

4.5. Discovery of devices

One of the services of the common mapping and ID services infrastructure is to allow discovery of the entities. Results of such discovery can be used to apply additional application services. Most of these services need secured and authenticated registration to the GRIDS (to verify for example the entity 'type' it is claiming) with an immutable and unique Identity [IDEAS-IDENTITY] by the GRIDS

provider. Some of these use cases require geographical co-ordinates too be part of LOC. Examples:

a. A specific mobile entity type as defined in the metadata of the IDy (if the entity allows it to be searchable) in a particular geographical area. This operation should give the long-lived IDf of the entities in that area, thus enabling communication among those entities.

b. Another example of discovery is asset tracking with extremely low cost services like bike sharing service. There is a need to track the exact location of individual bikes with low cost in long time span. In this case, there must be a module with a secured and long-lived identifier of the bike and which can set up communications whenever needed. The information associated with the identifier should be maintained with efficient resolution capability for tracking the labeled asset [ASSET1] [ASSET2].

4.6. Ad-hoc Networks Mobility across Hetnet

Today, the cost of reestablishing both the session and any security association is prohibitive in real-time applications with mobility. The local and proximity services on the GRIDS could be leveraged to provide the session continuity and security features.

Furthermore, there are opportunities for a dynamic and adhoc network to be formed between groups of vehicles on the highway, and these networks should be able to quickly peer along the edge with different networks encountered during mobility. Essential components of publicly sharable Metadata associated with entity's IDy/IDf at GRIDS would enable these peerings. But the specifics of how this can be used (TBD) and can only be expanded after IDEAS architecture is evolved.

5. Security Considerations

While access restriction policies can protect the entities from unwanted communication from unauthorized entities, how policy is specified, located and shared may cause new security threats. This aspect has to be considered during requirements and architecture.

6. IANA Considerations

This document has no actions for IANA.

7. Contributors

The following individuals (by first name alphabetical order) have contributed to this document:

Shreyasee Mukherjee

Parishad Karimi

Da Bin

Liu Binyang.

Jim Guichard

Albert Cabellos`

This present document is based on an extract of the first version of the draft. The authors and their affiliations on the original document are: D. Farinacci (lispers.net), D. Meyer (Brocade), D. Lake (Cisco Systems), T. Herbert (Facebook), M. Menth (University of Tuebingen), Dipenkar Raychaudhuri (Rutgers University), Julius Mueller (ATT) and Padma Pillay-Esnault (Huawei).

8. Acknowledgments

The authors would like to thank Kevin Smith, Alex Clemm and Yingzhen Qu for their review and input on this document.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

9.2. Informative References

[ASSET1] OFO, , <http://www.ofo.so/>.

[ASSET2] ITU, , <http://tracknet.net/1>.

[IDEAS-IDENTITY]
Chunduri, U., Clemm, A., and M. Menth, "Identity Use Cases in IDEAS", July 2017, <https://datatracker.ietf.org/doc/draft-ccm-ideas-identity-use-cases>.

[IDEAS-PS]
Pillay-Esnault, P., Boucadair, M., Jacquenet, C., Fioccola, G., and A. Nennker, "Problem Statement for Identity Enabled Networks", March 2017, <https://datatracker.ietf.org/doc/draft-padma-ideas-problem-statement/>.

[LISP-PRED]
Farinacci, D. and P. Pillay-Esnault, "LISP Predictive RLOCs", November 2016, <https://datatracker.ietf.org/doc/draft-farinacci-lisp-predictive-rlocs/>.

Appendix A. Appendix

This section contains few more use cases GRIDS or Identity can bring to IDEAS. Some of the use cases would be integrated based on the TBD architecture of IDEAS.

A.1. Network Simplification

Whilst the term IoT encompasses a wide range of applications ranging from short, low data-rate communications through to latency-sensitive haptic control loops, the ability to reduce network overhead through the simplification and potential removal of addresses at specific points on the network has several benefits.

A.1.1. Proximity Services and Scopes

In a system generating small amounts of data, the relative size of the addressing which can be considered to be network operator overhead compared to the size of the data payload which is customer data can be very high to the point where the addressing and control data vastly outsizes the customer data. In a traditional IT network, this is not often the case; the address and management data is very-much smaller than the payload and is thus an acceptable tax on the overall communication. Small data IoT applications require a solution to address the imbalance which now exists between the payload and the overhead.

It is possible to imagine local services within a scope that require only a small ID within a scope and can communicate with a local GRIDS to resolve its mapping and location services if it is only

communicating with local devices such as local networked sensors. Such an example could sensors be in a factory.

A.1.2. Ease of troubleshooting and Management

Trouble-shooting a complex, multi-layer network where data is handed from one encapsulation to another is complex. It is already seen in solutions such as SIP that having embedded naming structures vastly improves the ability to determine a data-path in a trouble-shooting instance. Similarly, an IDf-based solution would apply from the data producer to the consumer and can also be used to improve data-traceability thereby resulting in lower operational costs.

The presence of a GRIDS can improve on managing the data paths as they are stored in their mappings. Such improvement can be especially beneficial if GRIDS capabilities interact with a SDN-based computation logic, whose service-inferredpolicy-based decision making process may choose to redirect traffic onto alternate paths as a function of the nature of the service, the load status of the primary path, etc. Application of such dynamics can be critical for specific IoTservices, such as e-health services. Data analytics on the GRIDS may also be logging events for easier postmortem if there are issues.

A.1.3. Application of Common policies

A Mapping System can contribute to enforce a set of common policies for a given connectivity service by providing a global, consistent identification and resolution service to any participating device involved in the forwarding of the corresponding traffic.

A.2. Mobile Networks

For Seamless and global mobility of an entity, GRIDS can potentially offer various services that can include not only traditial mapping and resolution services based on Identifier, but protecting and honoring the entity's privacy restrictions on who can acces entity's LOC information for example.

Other services and Optimizations such as "late binding", in which identity of in-transit packets can be bound to a locator closer to the edge of the network to account for highly mobile vehicular scenarios can be further developed based on the GRIDS.

A.2.1. Mobility within a single provider network

Figure 4 provides an example of the topology of a single mobile carrier network. The goal of this section is not to specify specific mobile technology arechitecture but represent the need for a

standardized mapping and ID services solution, which are relevant in these domains.

UE.A, UE.B, and UE.C are devices connected to the mobile network and are themselves mobile. UE.A and UE.B are currently connected to base station Base1 and UE.C is currently connected to Base2. The base stations are connected to the radio access network (RAN) of the provider. The RAN is connected to the Internet through a transport network via providers core network (CN). Regardless of which mobile technology is used and CN is vritualized with separated control plane (5G) GRIDS instance can be deployed for all ID services.

Host1 Host1 and Host2 are hosts in the Internet that are communicating with the UEs. In this example we assume that ID functionality is handled within the network and is transparent to the UEs or hosts.

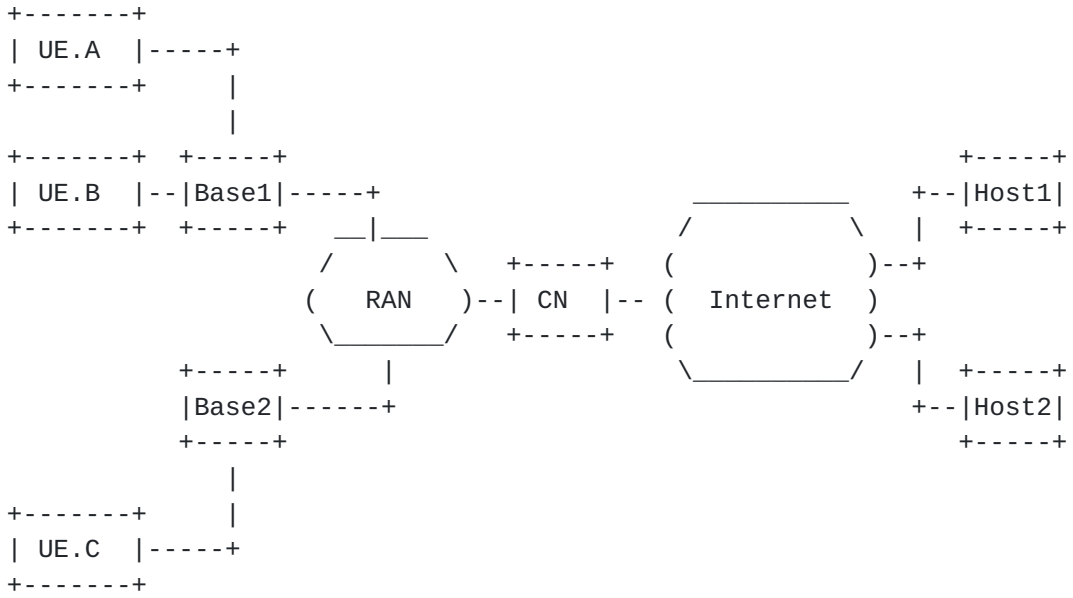

The role of the CN is to leverage the GRIDS, that maps destination addresses on ingress packets into the network to deliver packets to the destination. The GRIDS Infrastructure maintain a list of ID-to-LOC mappings. Hosts on the Internet only know the identifier of the mobile nodes. Depending on the dataplane solution packets forwarded from the Internet may be routed to a CN that is able to retrieve the mapping ID to the destination LOC to facilitate delivery.

Consider that Host1 sends a packet to UE.A. The destination address of the packet is the identifier address of UE.A. The packet is sent by Host1 and is forwarded across the Internet to the provider's

network based on the identifier address. A CN receives the packet. The destination address of a packet (identifier) is looked up in the mapping table (local GRIDS). The return result is the LOC for UE.A. The CN modifies the packet so that the destination address is the LOC for UE.A (either by encapsulation or address translation). The packet is then forwarded to Base1. At Base1 the destination is changed back to the identifier address and forwarded to the UE.

In the case that two UEs need to communicate the process is similar. Consider that UE.A sends a packet to UE.C. Again, UE.A only knows the identifier address of UE.C. UE.A sends a packet into the network and it is forwarded to a CN. The CN maps the destination address of UE.C to its locator and forwards the packet to Base2 which translates the destination back to an identifier address.

In order to reduce latency and improve performance, a mapping cache may be implemented at or near the base stations. A mapping cache would maintain a subset of mappings in the network that are being used for communications by the attached UEs to other devices in the network. A protocol would be run between the base stations and CNs to populate and invalidate entries in the cache.

Now consider that UE.B changes locations so that it is now attached to Base2 (figure 5).

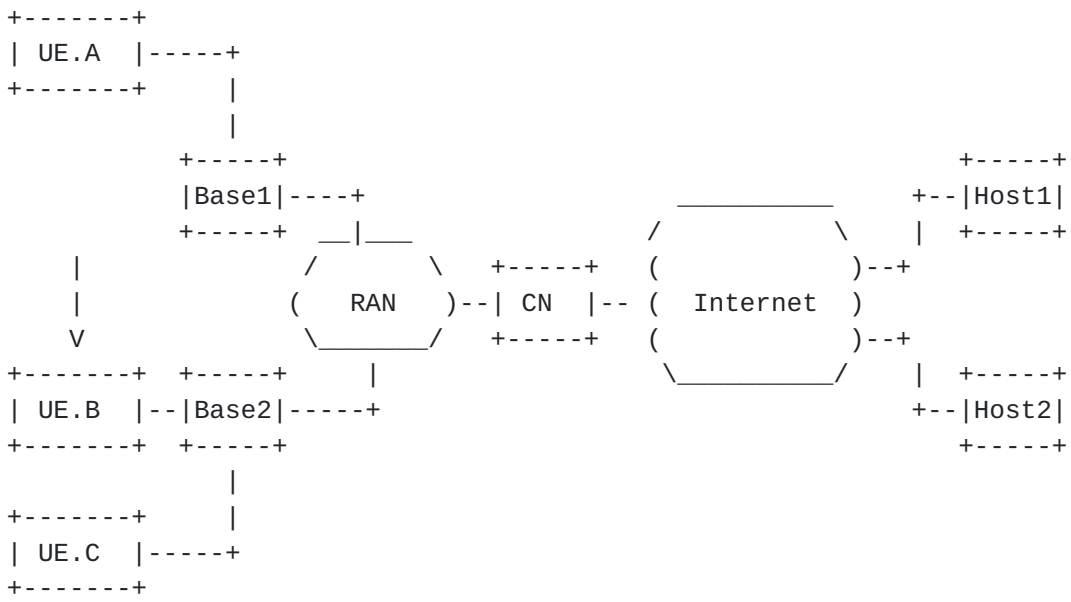

The CNs are updated to reflect UE.B's new location. As updating location is not instantaneous across all the mapping gateways in the

network, it is possible that a packet is forwarded to an old destination. Several possible solutions exist today:

1. Predictive routing. Pre-populating the (ID,LOC) tuple in the GRIDS itself so that multiple packets may get delivered ahead of the move as described in [LISP-PRED]. A predictive algorithm can be used to forward packets ahead of the move with minimum duplication.

2. Late binding. This method refers to rebinding of identifiers to addresses after a packet enters the network. This capability makes it possible for routers in the network to redirect packets whose end-point address might have changed due to fast mobility or rapid changes in radio link association or multicast group membership. This type of dynamic rerouting can help to improve network efficiency and reduce packet drops. Late binding generally requires in-network elements such as routers and base stations to be able to store data packets temporarily and access the ID to address mapping (name resolution) service.

3. Traditional Redirection. For instance, Base1 may receive packets for a UE.B after it has moved to Base2. A "care of address" could be set on Base1 for some period after the move. When Base1 receives a packet for UE.B it can map the destination address to reflect UE.B's new location and forward the packet. This doesn't need any local GRIDS instance and in some form done currently.

Some of the above additional services can be done through local GRIDS instance in the control plane of the operator network, thus not needing any anchor based solution and avoid "Traditional Redirection" for example.

A.2.2. Mobility across CNs

The role of GRIDS and applicability depends on the mobility scenarios specific to the mobile technology i.e., for intra RAN mobility local GRIDS instance can be used for an anchor free transport network. For seamless mobility across CN's of a same provider needs a global GRIDS Infrastructure as host can be located any where and the LOC of UE changes. This is one of the open issues of the current mobile networks.

A.2.3. Mobility of a Subnet

Figure 8 presents a topology where UE.A, UE.B, and UE.C are connected to a subnet and the whole subnet may itself be mobile (for instance a Wi-Fi network on a bus). To treat the subnet as an aggregate devices, the subnet identification is reflected in the identifier

address. A single mapping entry then could then represent this subnet where the identifier is the prefix for the subnet and the locator reflects the location of the subnet (Base1 in this case)

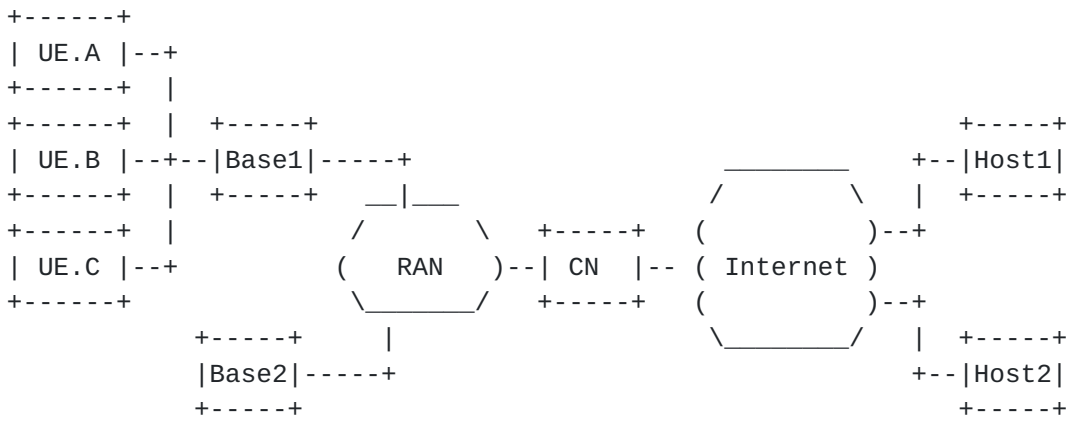

When the subnet moves the mappings to the identifier, the prefix for the subnet is updated in the CNs (figure 9). Note that a UE could also move out of its subnet and attach to the network on its own, for instance a UE might switch from using Wi-Fi to using a mobile network. This will work if a specific mapping for UEs is allowed in the mapping database, and that mapping is preferred over the aggregate mapping since it has a longer identifier prefix.

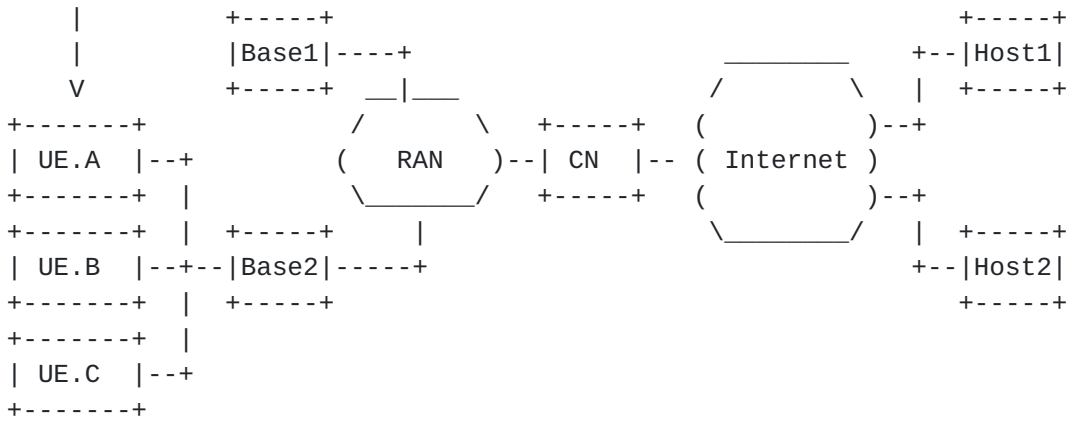

Similar to the case of a UE moving between carrier networks, a subnet can move between carrier networks if the networks share identifier locator mappings that include identifier prefixes. Also, subnet mobility can modify community management and therefore impact GRIDS

service operation: if UE.A, UE.B, and UE.C have subscribed to different services where, for example, UE.A exchanges information with UE.B but not with UE.C, whereas, UE.C exchanges information with UE.B but not UE.A, the mobile subnet may support different "closed user groups" that may be serviced by different GRIDS. When the subnet moves, other GRIDS may be invoked to make sure communication within the said closed user groups is not disrupted.

Authors' Addresses

Padma Pillay-Esnault (editor)
Huawei
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: padma.ietf@gmail.com

Dino Farinacci
lispers.net
San Jose California
USA

Email: farinacci@gmail.com

Christian Jacquenet
Orange
Rennes 35000
France

Email: christian.jacquenet@orange.com

Uma Chunduri (editor)
Huawei
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: uma.chunduri@huawei.com

Tom Herbert
Quantonium

Email: tom@herbertland.com

David Lake
Dell

Email: d.lake@surrey.ac.uk

Michael Menth
U of Tuebingen

Email: menth@uni-tuebingen.de

Dipenkar (Ray) Raychaudhuri
Rutgers University

Email: ray@winlab.rutgers.edu

Dave Meyer

Email: dmm@1-4-5.net