INTERNET-DRAFT T. Herbert
Intended Status: Experimental Quantonium
Expires: August 2018

February 3, 2018

Identifier groups
draft-herbert-idgroups-00

Abstract

This draft describes a means to create logical identifier groups to manage identifiers in a mapping system for identifier-locator protocols. An identifier group consists of identifiers that have similar properties in the context of the mapping system. Identifier groups facilitate bulk operations on the mapping system that would affect multiple identifiers. A primary use case for this is to facilitate mobility of devices that are associated with possibly thousands or even millions of identifiers.

Status of this Memo

Copyright and License Notice

Table of Contents

1 Introduction

This document describes identifier groups for identifier-locator mapping systems.

Identifier-locator protocols include the concept of identifiers as a type of node addressing. Identifiers are logical endpoints of communications and only differ from canonical addresses in that they are not topological. A node may be assigned multiple ephemeral identifiers so that they be can used to create different source addresses for different communications to benefit privacy and anonymity. It is expected that individual end devices may have thousands of active ephemeral identifiers; a device that connects backend subnets could have millions of associated identifiers.

An identifier-group is an group of identifiers within a mapping system that share some common properties. A grouping is arbitrary, the given application or mapping system may create identifier groups as needed. An identifier may belong to multiple groups, however when an operation is performed it must be clear as to which group applicable properties are be derived from. Groups may also be hierarchical such that groups may be members of other groups and thus inherit properties from their parent groups.

A primary application of identifier groups is mobility where a device has a number of identifiers associated with it. When such a device moves in the network and is assigned a new locator, all of the identifiers associated with the device assume the new locator also. Identifier groups provide a level of indirection so that the locator can be set for all of the associated identifiers for the device in a single operation on the mapping system.

2 Characteristics of identifiers

This section list some salient properties of identifiers that are relevant to a mapping system and privacy.

2.1 Identifier addresses

Identifier addresses are full IP addresses that are either an identifier or contain an identifier as part of the address. Identifier addresses are used by endpoints to achieve communications. In order to reach the end host where the node indicated by an identifier resides, somewhere in the path an identifier-locator operation is performed and the packet is typically modified (either by encapsulation or address translation) to reach the correct node. At the destination node, a reverse operation is done to restore the originally sent packet before presenting the packet to the end node

or application.

Identifier addresses should have the following properties:

2.2 Desired properties

- o They are composed of a global routing prefix and a suffix that is internal to an orgnization. This is the same property for IP addresses [RFC3513].

- o The registry and organization of an address can be determined by the network prefix. This is true for any global address.

- o The organizational bits in the address should have minimal hierarchy to prevent inferences. It might be reasonable to have an internal prefix that divides identifiers based on broad geographic regions, but detailed information such as location, department in an enterprise, or device type should not be encoded in a globally visible address.

- o Given two identifier addresses and no other information, the desired properties of correlating them are:

  - o It can be inferred if they belong the same organization and registry. This is true for any two global IP addresses.

  - o It may be inferred that they belong to the same broad grouping, such as a geographic region, if the information is encoded in the organizational bits of the address.

  - o No other correlation can be established. For example, it cannot be inferred that the IP addresses address the same device, the IP addresses reside in the same subnet or department, or that the nodes for the two addresses have any geographic proximity to one another.

2.2 Policy mechanisms for identifiers

Other than a globally routable network prefix, identifier addresses require no hierarchy since they are not topological. Therefore all or most of the organizational bits in a publicly visible address form a flat, non-hierarchical space. To create identifier addresses with the properties listed above, the bits in this space are pseudo-randomly assigned to form addresses.

While the routing requirements are satisfied by the identifier-locator protocols and mapping system, the lack of internal hierarchy in addresses is a potential disruption for network deployments that

rely on address hierarchy to implement policy. For instance, an enterprise might implement a firewall rule base on destination network prefix that prevents the engineering department from talking to human resources.

In order to apply such policies and still maintain the properties to prevent inference, a firewall could create rules based on identifier groups. So when a packet arrives at the firewall, the mapping system may be consulted and information for a group is returned. A policy decision, i.e. forward or drop, may be made per this information.

In the example above, identifier groups might be created for engineering and human resources. The policy is expressed that members of the engineering group are not allowed to send to members human resources group. Since the groups are not encoded in the addresses there is no means for an external party to infer which packets belong to engineering and which belong to human resources. This is a privacy benefit compared to common method of encoding the department in the address hierarchy. An additional benefit is that such groupings are arbitrarily flexible and are not constrained by the need to format information into addresses (address prefixes for instance). Since the addresses don't contain group information, group membership can be changed for an address without requiring the node to change its address.

3 Structure of identifier groups

Identifier groups can form a hierarchical structure within a mapping system domain. The diagram below illustrates a hierarchy containing two levels of groups and six identifier mapping entries at the leaves.

```
                                  +-------+
                                  |       |
                                  | Group |
                                  |       |
                                  +---+---+
                                      |
        +------------+          +-----+-----+          +------------+
        | Identifier |---+      |           |      +---| Identifier |
        +------------+   |  +---+---+   +---+---+  |   +------------+
        +------------+   |  |       |   |       |  |   +------------+
        | Identifier |---+--| Group |   | Group |--+---| Identifier |
        +------------+   |  |       |   |       |  |   +------------+
        +------------+   |  +-------+   +-------+  |   +------------+
        | Identifier |---+                         +---| Identifier |
        +------------+                                 +------------+
```

The diagram below provides an explicit example of using an identifier

group hierarchy for mobility.

In this scenario, we consider a bus has an onboard WIFI network. There are two UEs attached to the WIFI, where both have been assigned three identifiers.

```
                                 +----------+
                                 |  WIFI    |
                                 |  Bus     |
                                 |  Locator |
                                 +-----+----+
                                       |
        +------------+           +-----+-----+            +------------+
        | Identifier |--+        |           |        +---| Identifier |
        +------------   |  +-----+---+   +---+-----+  |   +------------+
        +------------+  |  | UE      |   | UE      |  |   +------------+
        | Identifier |--+--| Locator |   | Locator |--+---| Identifier |
        +------------+  |  |         |   |         |  |   +------------+
        +------------+  |  +---------+   +---------+  |   +------------+
        | Identifier |--+                             +---| Identifier |
        +------------+                                    +------------+
```

In this hierarchy, each UE has an associated group that contains all the identifiers for the UE. The WIFI device has an associated group that contains the groups for the attached UE devices. With this structure, each identifier has two locator mappings. The first one maps the identifier to the WIFI device in the bus. The second maps the identifier to the UE attached to the WIFI network.

When a packet from an external network is sent to one of the identifiers, the mapping system is consulted to retrieve the top level locator to forward the packet. This locator will direct the packet to the WIFI router on the bus. At the bus WIFI router, the second level locator mapping for the identifier is consulted to determine the locator of the UE that has the identifier. The resultant locator is used to forward the packet to the appropriate UE device. At the UE, the identifier is used to deliver the packet to the appropriate application.

As the bus moves through a mobile network, the locator for the WIFI changes so effectively the top level locator for all the identifiers for all the UEs within the bus also must be changed. Identifier groups allow this to be done in one operation on the mapping system. When passengers disembark and leave the range of the WIFI, the group membership of the UE is disassociated from the WIFI bus group. The UE may attach to another network so that the locator or group membership for the UE would be set appropriately.

Note that in the above example, an identifier group hierarchy is used

to create a locator hierarchy. That is, multiple identifier locator operations are performed to get packets to destination. This is expected to be common in identifier-locator deployments. It is analogous to a packet going through a routing hierarchy where at each level the information applied became progressively more specific to the final destination (i.e. at each layer the prefix match is longer).

4 Interfaces

The mapping system interface is logically divided into the management interface and the query interface.

4.1 Management interface

The management interface is used to create and manipulate mapping entries and identifier groups.

The allowed operations on the management interface are:

- Create groups
- Set properties of a group, such as a locator or membership in another group in a group hierarchy
- Change properties of a group
- Create identifier mapping entries
- Set identifier mapping properties such as locator or group membership
- Change identifier mapping properties
- Delete an identifier mapping entry
- Remove all members from a group
- Delete all identifier mappings in a group
- Delete a group (that has no members)

Note that there is no public interface defined that will return all the members of a group. This is intended to limit visibility to this sensitive information.

4.2 Query interface

The query interface is used by devices that require identifier to locator mappings. This interface is read-only.

The basic operations in the query interface are:

- Lookup locator for an identifier. In the case that a group hierarchy is present, the lookup request includes an indication as to which level in the hierarchy is applicable.

- Lookup group information by group identifier. This is needed if the entry returned in a mapping entry indicates a group in a level of indirection. The internal structure for mapping entries which are members of the same group may reference a single group structure.

- Request notifications of mapping entry changes if the mapping system supports pub/sub model. This includes notifications that a group membership has changed.

- Request notifications of group changes. For example, if the locator for an identifier group changes.

5 Security Considerations

Access to mappings of group identifier to member identifiers MUST be strictly controlled. If this information is compromised, then privacy and anonymity of users could be undermined. In the case that the group identifiers refer to a single device, such as a UE in a mobile network, breach of the mapping from group identifier to identifiers may be sufficient to compromise individual user identities. Note that these concerns are not specific to identifier-locator mapping systems, but in any scenario where address assignment is done for devices.

The management interface should provide very strong authorization and employ encryption when communicating with the mapping system. The mapping system should enable security mechanisms associated with databases that contains sensitive information.

The query interface is always read-only, however this should also have strong access authorization methods for security and privacy.

A distributed identifier-locator mapping system should be deployed within a single administratively controlled domain. Low level information that potentially contains PII (Personally Identifiable Information) or specific location information should never be shared between administrative domains. It is conceivable that two networks could share a high level identifier-locator mapping system distinct

from their internal systems to support cross domain identifier-locator mappings. In this case, a locator hierarchy would be employed so as not to reveal any detailed information or PII. Specifically, identifier group information that refers specific devices and end locators for specific devices should not be visible.

6 IANA Considerations

7 References

7.1 Normative References

7.2 Informative References

Author's Address

Tom Herbert
Quantonium
Santa Clara, CA
USA

Email: tom@quantonium.net