INTERNET-DRAFT
Intended Status: Informational
Expires: August 2018

T. Herbert
Quantonium
K. Bogineni
Verizon

February 1, 2018

# Identifier Locator Addressing for Mobile User-Plane
# draft-herbert-ila-mobile-00

Abstract

This document discusses the applicability of Identifier Locator Addressing (ILA) to the user-plane of mobile networks. ILA allows a means to implement network overlays without the overhead, complexities, or anchor points associated with encapsulation. This solution facilitates highly efficient packet forwarding and provides low latency and scalability in mobile networks. ILA can be used in conjunction with techniques such as network slices and Network Function Virtualization to achieve optimal service based forwarding.

Table of Contents

## 1 Introduction

In mobile networks, mobility management systems provide connectivity while mobile nodes move around. A control-plane system signals movements of a mobile node, and a user-plane establishes tunnels between mobile nodes and anchor nodes over IP based backhaul and core networks.

This document discusses the applicability of Identifier Locator Addressing (ILA) to those mobile networks. ILA is a form of identifier/locator split where identity and location of a node are disassociated in IP addresses. ILA nodes transform destination addresses of packets by overwriting part of the address with a locator. The locator provides the topological address for forwarding a packet towards its destination. Before a packet is delivered to the end destination, the destination address is reverted to its original value.

An ILA mobile user-plane implementation needs both data plane and control plane components.

The data plane includes the ILA transformation processing as well as handling to maintain conformance with IP protocols. The ILA data plane is described in [ILA].

The control plane's primary function is to maintain a mapping database that is shared amongst ILA nodes. The mapping database contains entries for the mobile nodes in the network, and the number of mapping entries is expected scale into the billions. In order to scale, a two level hierarchy of ILA nodes is defined by ILA routers and ILA forwarding nodes.

ILA routers maintain a full set of ILA mappings. Routers may be replicated for redundancy and load balancing. The mapping system may also be sharded, so that each router is responsible for a shard. Routers use a protocol to synchronize the mappings for each shard.

ILA forwarding nodes perform a reverse ILA transformations to restore the destination address in packets before delivery. A forwarding node can also maintain a cache of ILA mappings to perform transformations on intra domain traffic as an optimization to avoid having to forward packets through ILA routers. Forwarding nodes are typically located close to the mobile nodes. The ILA Mapping Protocol [ILAMP] is used between forwarding nodes and ILA routers to manage the cache.

## 2 Conventions and Terminology

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].

ILA related terms are defined in [ILA].

## 3 Motivation

Emerging applications such as VR, AR, and autonomous vehicle communication require very low latency, high bandwidth, and high reliability. For mobile devices, these requirements must not only be met when the device is stationary, but also across handover during mobile events. Mobility needs to be a seamless operation where IP addresses and connections are maintained. In a second dimension, the number of connected mobile devices, including a large contingent of IoT devices, is expected to grow by several orders of magnitude within a few years as enabling technologies such as 5G are deployed.

The convergence of mobile networks and datacenter networks is also pertinent. Simple physics (i.e. speed of light) dictates that very low latency for applications (order of less than five milliseconds) can only be achieved by placing application servers in close proximity to clients and minimizing the number of network hops. An emerging trend is for providers to house datacenters within the network to run applications. For similar reasons, many providers are integrating multi-tenant cloud services directly in their mobile networks. The upshot is that mobile networks need to support the convergence of mobile devices, datacenter virtualization, and cloud. A single solution framework for all of this is desirable.

Current mobile architecture is hitting the limits of scalability and performance. In particular, anchor points used in 3GPP have become single points of failure, bottlenecks, and lead to sub-optimal triangular routing. The anchor model is also inflexible in attempts to leverage services of the transport network such as network slices and Network Function Virtualization. The control plane to manage millions of GTP tunnels is complex and difficult to scale. GTP-U is narrowly defined for a particular use case, which makes it difficult to leverage for other use cases. The use of any in-network tunneling, including GTP, raises issues of overhead, MTU and fragmentation, security, and other complexities.

ILA is a proposed alternative to GTP-U and encapsulation. It does not require anchors and simplifies both the data plane and control plane. ILA has zero wire overhead so there are no issues around MTU and fragmentation. Its use is transparent to the network, and it is

compatible with existing hardware and commonly deployed protocol optimizations. ILA is a general network overlay protocol can be used to meet the requirements of use cases in a converged network. User Plane Functions (UPF) with ILA are lightweight and stateless such that they can be brought up quickly as needed.

## 4 Reference topology

Figure 1 shows an example topology of ILA in a mobile network.

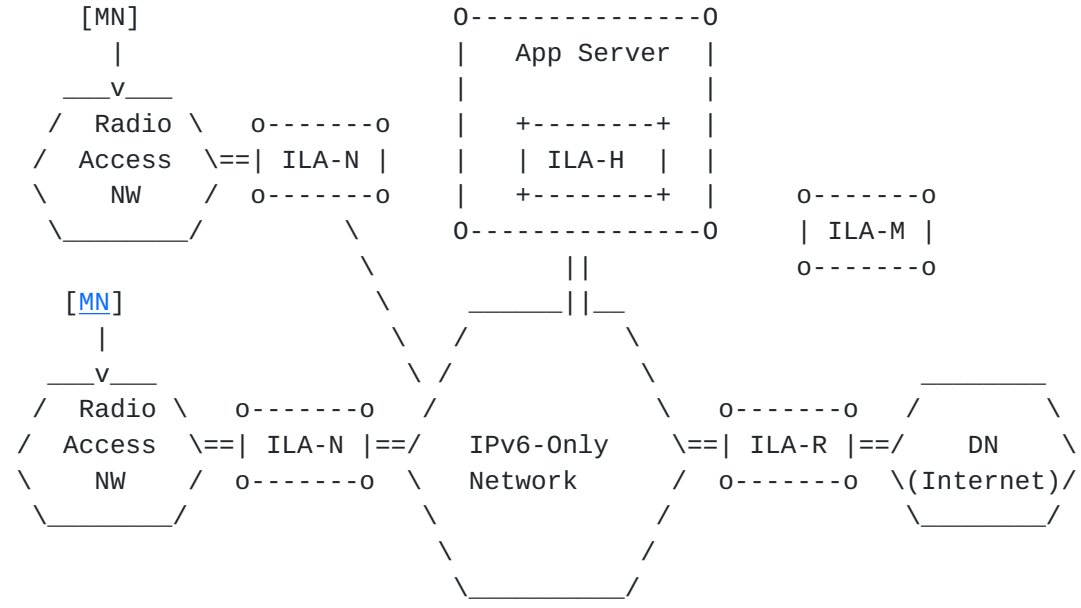


Figure 1: Mobile User-plane with ILA

There are four types of functional nodes in the ILA architecture:

- ILA routers (ILA-R)
- ILA forwarding nodes (ILA-N)
- ILA hosts (ILA-H)
- ILA managment (ILA-M)

### 4.1 ILA routers (ILA-R)

ILA routers are deployed within the network infrastructure and collectively contain a mapping database of all identifier to locator mappings in an ILA domain. The database may be sharded across the identifier space by some number of ILA routers for scalability. ILA

routers may also be replicated for scalability and availability.

ILA routers provide two main functions: ILA forwarding and mapping resolution. An ILA router may perform one both of these functions at the same time. If a router performs both functions it may send ILA redirects.

#### 4.1.1 Forwarding routers

Forwarding routers perform ILA transformations when packets enter an ILA domain. A destination address of a packet that is a SIR address is transformed to an ILA address. The process is that the router performs a lookup on the destination address in a mapping table and a locator is returned. The locator is written into the destination address of the packet (typically the high order sixty-four bits are overwritten with a locator).

In the case of a sharded database, the high order bits of the identifier indicate the shard number. This is included in a routing prefix so that the packet is routed to an ILA router that contains the database for the indicated shard.

#### 4.1.2 Mapping resolution

An ILA router that is performing mapping resolution will respond to mapping requests from ILA forwarding nodes or ILA hosts (these are described below). The mapping request protocol allows the caller to request the locator for an identifier address.

### 4.2 ILA forwarding nodes (ILA-N)

ILA forwarding nodes are deployed in the network infrastructure towards the edges to provide ILA transformations for end devices. ILA forwarding nodes have two functions: ILA to SIR address transformation and ILA forwarding. As indicated in the reference topology, forwarding nodes may be deployed near the point of device attachment (e.g. base station, eNodeB) of mobile nodes.

#### 4.2.1 ILA to SIR address transformation

In the path towards the end devices, forwarding nodes perform ILA to SIR address transformation. That is, they perform a reverse ILA transformation in order to restore the original addresses in packet. Forwarding the packet on to the destination is done based on the SIR address. For instance, an eNodeB may map a SIR address to a layer 2 address of the attached device that has the SIR address. Note that this functionality is required somewhere in the path between the ILA node that writes a locator into an address and the ultimate

destination device (e.g. a UE). It is not recommended that this functionality is implemented on end user devices.

When a node migrates its point of attachment from one ILA-N to another, the local mapping on the old ILA-N is removed. If an ILA addressed packet is received by an ILA-N for which there is no local mapping, then the packet is forwarded back into the network with a destination SIR address. The packet should be forwarded through an ILA router that can perform the tansformation for the new ILA-N. A "negative" mapping with timeout may also be set in an ILA-N to ensure that ILA-N is able to infer the SIR address (e.g. would be needed with non-local identifiers).

#### 4.2.2 ILA forwarding

A forwarding node may perform ILA transformation and forward packets directly to peer ILA nodes in the same domain. The mappings for this are maintained in a working set cache in each ILA-N. As a cache there must be methods to populate, evict, and timeout entries. A cache is considered an optimization, so the system should be functional without its use (e.g. if the cache has no entries). The possibility of Denial of Service attack (DOS) on a cache being populated by unmanaged outside events, in this case mobile devices sending packets to arbitrary destinations, must be considered in the cache design.

If a packet is received by an ILA forwarding node from a downstream node that is destined to another node in the same ILA domain for which there is no existing cache entry, then:

- The packet is forwarded by address. The SIR address plus shard identifier prefix will route the packet to a forwarding ILA router which will perform ILA transformation of the packet to reach its destination.

- An ILA router may return an ILA redirect to inform the forwarding node of a direct ILA mapping.

- If the forwarding node gets a mapping from an ILA router, then subsequent packets for the destination can be directly sent using the mapping. Note that a forwarding node does not hold packets that are pending mapping resolution.

### 4.3 ILA hosts (ILA-H)

ILA host are forwarding nodes that are embedded in end servers to provide ILA transformation. Since an ILA host is integrated with the host stack sourcing packets, there are opportunities for optimizing processing.

ILA is not recommended to run on end user devices, however there may be servers or other end devices that are in the provider network that might benefit from participating in ILA (this is illustrated in the reference topology above). A server that implements ILA forwarding can directly send to ILA peers in the same domain to avoid triangular routing.

### 4.4 ILA management (ILA-M)

The ILA management node provides the interface between the ILA infrastructure and mobile management of a network. Similar to ILA-Rs there may be multiple ILA-Ms in the network and they can be replicated for redundancy and load distribution. Data managed by ILA-Ms needs to be synchronized across ILA-Ms. It is conceivable that the set of ILA-Ms could be split into shards serving different geographic area in order to localize data. ILA-Ms may be co-located with ILA-Rs so that there is a fast path between them.

The management nodes are responsible for:

- Receiving notifications from the session management in the mobile network. Notifications of interest include: when mobile nodes attach to the network, are removed from the network, or change their point of attachment in the network (i.e. they move).

- Managing identifier groups. Identifier groups are sets of identifiers (nodes) that share common properties [ILAGRPS]. In a mobile network, identifier groups are used to represent all the identifiers assigned to a mobile node. Each mobile node will have its own identifier group.

- Writing identifier locator mappings into the ILA mapping database. The written content is based on the information provide by session management.

- Changing the mapping table when a locator for an identifier, group, or mobile node changes. A locator for a device changes when its point of attachment changes.

- Creating identifiers for attached devices. Identifiers may be persistent so that each time a device attaches it gets the same identifier.

- Registering ILA-Rs, ILA-Ns and their locators. ILA-Ms coordinate the operation of ILA nodes in the network.

## 5 Data plane operation

ILA performs transformations on IPv6 addresses of packets in flight. A SIR to ILA address transformation overwrites the destination address with a locator address for forwarding over a network. An ILA to SIR address transformation restores an IP address to its original contents. The transformations are always paired so that a SIR to ILA address transformation is always undone before delivery. End hosts and applications only see SIR addresses. Effectively, ILA is a mechanism to implement transparent network overlays. Note the process is specifically called a "transformation" as opposed to "translation" which distinguishes ILA from NAT. NAT translations are not undone before reception and NAT is not transparent to the end points.

### 5.1 SIR to ILA transformation

SIR to ILA address transformations may be performed by ILA routers, ILA forwarding nodes, and ILA hosts.

SIR to ILA transformation is done by a lookup on the destination address in a mapping table. On an ILA router this table contains all the entries for the shard the router serves. On a forwarding node or host, the table is a cache of entries. If a corresponding entry is found, then a locator is returned. The locator is written into the destination address.

If checksum neutral mapping is being used to preserve transport layer checksums, then that is indicated in the mapping entry. Checksum neutral mangles the low order sixteen bits of the identifier portion of the address. The checksum difference between the SIR prefix and the locator is added into to the low order sixteen bits of the identifier.

If an ILA router does not find a match on the destination address in its table then the packet is dropped as having no route to host.

If an ILA forwarding node or host does not find a match on the destination address, then it forwards the packet unchanged. The packet may encounter an ILA router that performs the transformation.

In response to forwarding a packet, a router might send an ILA redirect to an ILA forwarding node. A redirect informs a node of an ILA mapping that may be cached to avoid triangular routing when forwarding subsequent packets. The destination of a redirect is the upstream forwarding node of the source of packet. An ILA router can determine this by performing an ILA lookup on the source address of the packet being forwarded. This assumes that the source is a SIR address for the ILA domain and that the use of ILA is symmetric so that the lookup reveals the correct forwarding node; this needs to be accounted for in network design.

### 5.2 ILA to SIR transformation

Transformed packets are forwarded to an ILA-N or ILA-H based on normal routing of the packet with a locator in the its destination address (upper sixty-four bits). When a node receives the packet it first performs performs an ILA to SIR address transformation by mapping the received locator (one local to the node) to a SIR address. If checksum neutral mapping has been done, the lower sixteen bits in the identifier must be fixed up. This is done by subtracting the checksum difference of the SIR address and locator from the low order bits of the identifier (the opposite operation of setting the checksum neutral bits).

After transforming a destination back to SIR address, a lookup is performed on the identifier to determine if it is local (that is it refers to a node that is downstream of the ILA node). If the node is local, it is forwarded downstream using normal mechanisms of the network. If the node is not local, the SIR addressed packet is forwarded back into the network. The packet should traverse an ILA router that can transform its destination to the correct locator and possibly send a redirect towards the source.

### 5.3 Data path efficiency

There basic operations of ILA address transformation, either SIR to ILA or ILA to SIR, are:

1) Read destination address from a packet.

2) Lookup all or are part of the destination address in table. This is a fixed length lookup.

3) Overwrite all or part of the destination address with a locator value returned from the lookup.

4) Fix the checksum neutral mapping bits in the identifier.

5) Forward the resultant packet.

The computationally intensive operations in this path are the lookup and checksum neutral processing.

The lookup operation is on a fixed length key so a simple hash table can be used. It is also amenable for use with a hardware TCAM. On an ILA host, an ILA mapping may be cached with a connection context so that a lookup does not to be performed for every packet sent on the connection.

Checksum neutral processing entails 1's complement arithmetic over sixty-four or 128 bit values. In the case that the full 128 bit identifier address is a one-to-one mapping with a locator address, then the checksum computation is constant for a mapping and can be precomputed and saved with the mapping.

### 5.4 Alternative data path use cases

ILA supports multicast encoding, virtual networking modes, and IPv4/IPv6 translation. These require different processing, and in the case of IPv4/IPv6 translation the size of the packet increases. However, these alternative cases should not fundamentally increase the cost of the lookups since instructions for alternative processing can be returned by a lookup.

### 5.5 Locator chaning

ILA allows multiple locator transformations to effectively implement hop-by-hop source routing. This can be used to deliberately have a packet visit some set of nodes. This might also be used in the case where two domains are exchange ILA mappings, but only share locators that are ingress points in their network and not final locators of a node. This would be done to protect user location from being exposed.

### 5.6 ICMP handling

A packet whose destination address is an ILA address may generate an ICMP error. In this case the ICMP data will contain an IPv6 header whose destination is an ILA address. If a sender receives an ICMP error with an ILA address as the destination of the original packet, it won't recognize the destination address as one that it sent to and this may leak information about internal nodes of the network. To prevent this from happening, upstream ILA-Ns of ILA-Hs of an end node can filter ICMP packets. When an ICMP packet is received by these nodes, an ILA destination address can be transformed back to a SIR address by performing a reverse lookup.

## 6 Control plane

This section describes the ILA control plane for the mobile user-plane.

The ILA control plane is separate from the control plane of the mobile network. An interface between the session management of the network and the control plane is needed to get device information and point of attachment. The intent is that the interface is well compartmentalized to minimize the amount of specialization needed to adapt ILA for use in different access technologies.

### 6.1 ILA router mapping database

There are a number of options to use for implementing the ILA mapping system and router protocol amongst ILA-Rs. The mapping database must be able to scale and provide fast converge when mobile nodes move within the network.

#### 6.1.1 ILA with BGP

A traditional routing protocol could be used for route dissemination. [BGPILA] defines multiprotocol extensions to BGP for distributing ILA mappings.

#### 6.1.2 Key/value store

A mapping database is logically a simple key/value store where the lookup key is fixed length (sixty-four or 128 bytes). This characteristic affords the possibility of using a key/value database in lieu of traditional routing protocols.

The idea of the key/value database is that each shard is a distributed database instance with some number of replicas. When a write is done in the database, the change is propagated throughout all of the replicas for the shard using the standard database replication mechanisms. Mapping information is written to the database using common database API and requires authenticated write permissions. Each ILA router can read the database for the associated shard to perform its function.

The database is assumed to be (mostly) persistent and recoverable if database nodes are lost. The selection of an ILA router shard and shard instance is idempotent and stateless per packet, so that shards and shard replicas can be dynamically added or removed.

### 6.2 ILA Mapping Protocol

The ILA Mapping Protocol [ILAMP] is used between ILA forwarding nodes and ILA mapping resolution routers. The purpose of the protocol is to populate and maintain the ILA mapping cache in forwarding nodes.

ILA forwarding nodes can use a pull model (request/response), push model (pub/sub), or redirects to populate the mapping table. ILAMP runs over TCP which provides reliability, statefulness implied by established connections, allows use of HTTP and RESTful APIs, and standard security in the form of TLS.

The protocol is composed of message primitives:

- Map request: Sent by an ILA-N or ILA-H to an ILA-R to request mapping information for an IPv6 address.

- Map information: Sent by an ILA-R to an ILA-N or ILA-H and provides mappings. A map information message can be sent in response to a map request, when mappings are pushed in pub/sub, or a mapping is being advertised by ILA redirect. The reason the mapping information was sent in included in a message.

- Subscribe/unsubscribe: Sent by an ILA-N or ILA-H to an ILA-R. "Subscribe" requests mapping notifications for the listed identifiers. Notifications are sent when a mapping entry for an identifier changes. "Unsubscribe" requests that notifications for the listed identifiers stop.

- Locator unreachable: sent by an ILA-R to an ILA-N or ILA-H to indicate that another ILA-N is no longer reachable so all cache entries using that ILA-N or ILA-H should be evicted.

### 6.3 Address assignment

Mobile nodes are assigned addresses that serve as identifiers. A node may be assigned singleton addresses or a network prefix. Privacy is an important consideration in address assignment.

#### 6.3.1 Singleton address assignment

DHCPv6 or static address configuration can be used to assign singleton addresses to a node. These addresses have no topological component and are not meaningfully aggregable for routing, so an entry in the ILA mapping table would be created for each address. Nodes may be assigned thousand of addresses or even millions of IPv6 addresses. Given the large IPv6 address space there are few concerns about address depletion, however to the mapping system each address is represented in a identifier to locator mapping. Scaling this needs to be carefully considered. Sharding, replication, and caching on forwarding nodes are meant to provide scalability.

#### 6.3.2 Network prefix assignment

A node may be assigned a /64 address via SLAAC as is common in many provider networks. In this scenario, the low order sixty-four bits contains IIDs arbitrarily assigned by a device for its own purposes; these bits cannot be used as an identifier in identifier/locator split.

To support /64 prefix assignment with ILA, the ILA identifier can be encoded in the the upper sixty-four bits of an address and the lower

sixty-four bits are ignored by ILA. Since only a subset of bits are available, a level of indirection can be used so that ILA transforms the upper sixty four bits to contain both a locator and an index into a locator (ILA-N) specific table. The entry in the table provides the original sixty-four bit prefix so that ILA to SIR address transformation can be done.

## 7 ILA in 5G networks

The section describes applying ILA for use in a 5G network. ILA is instantiated as a function in the 5G services architecture described in [3GPPTS].

### 7.1 Architecture

Figure 2 depicts the use of ILA in a 5G reference point architecture. ILA is logically a network function and ILA interfaces to the 5G control plane via service based interfaces.

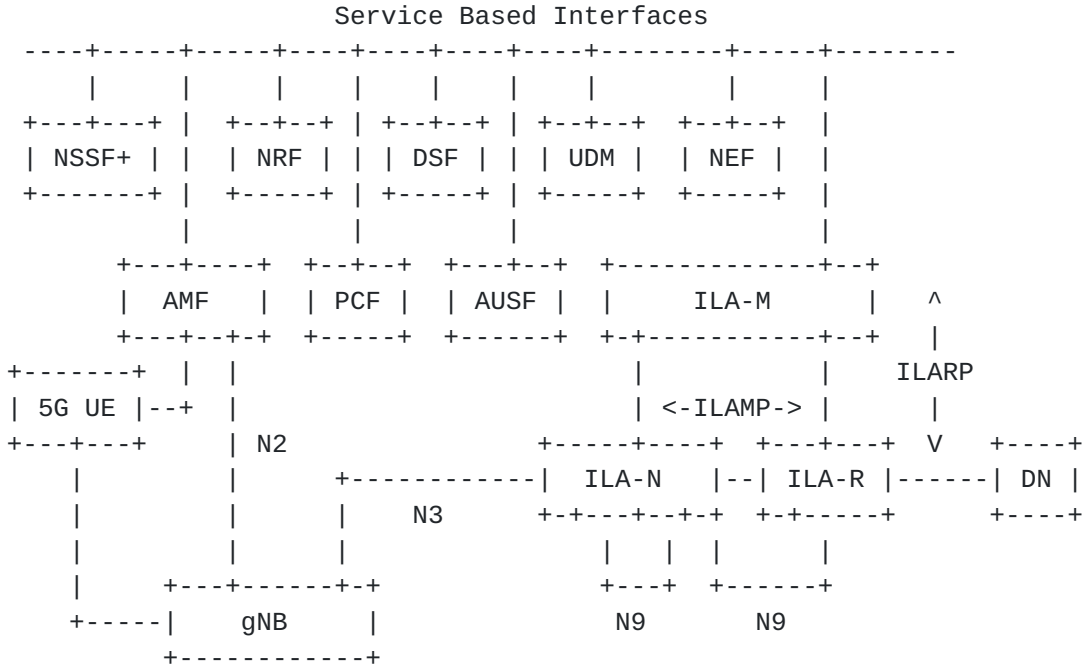


Figure 2: ILA in 5G reference point architecture

ILA is used over the N9 interface. SIR to ILA address transformations in the downlink from the data network are done by an ILA-R. Transformations for intra domain traffic can be done by an ILA-N close to the gNB or by an ILA-R in the case of a cache miss.

The control interface into ILA is via an ILA-M that interacts with 5G network services. ILA-M uses RESTful APIs to make requests to network

services (see section 7.3). An ILA-M receives notifications when devices enter the network, leave it, or move within the network. The ILA-M writes the ILA mapping entries accordingly.

ILA-Ms communicate with other ILA-Ms, ILA-Ns, and ILA-Rs in the same ILA domain via ILA control protocols that are independent of the 5G control plane. The mapping database is shared amongst ILA-Ms and ILA-Rs via an ILA routing protocol which is denoted by ILARP in the figure (see section 6.1). ILA-Ns communicate with ILA-Rs using ILAMP (see section 6.2).

ILA could be supported on a gNB. In this case, an ILA-N would be co-resident at a gNB and ILA is be used over N3 interface in lieu GTP-U.

## 7.2 Protocol layering

Figure 3 illustrates the protocol layers of packets packets sent over various data plane interfaces in the downlink direction of data network to a mobile node. Note that this assumes the topology shown in Figure 2 where GTP-U is used over N3 and ILA is used on N9.

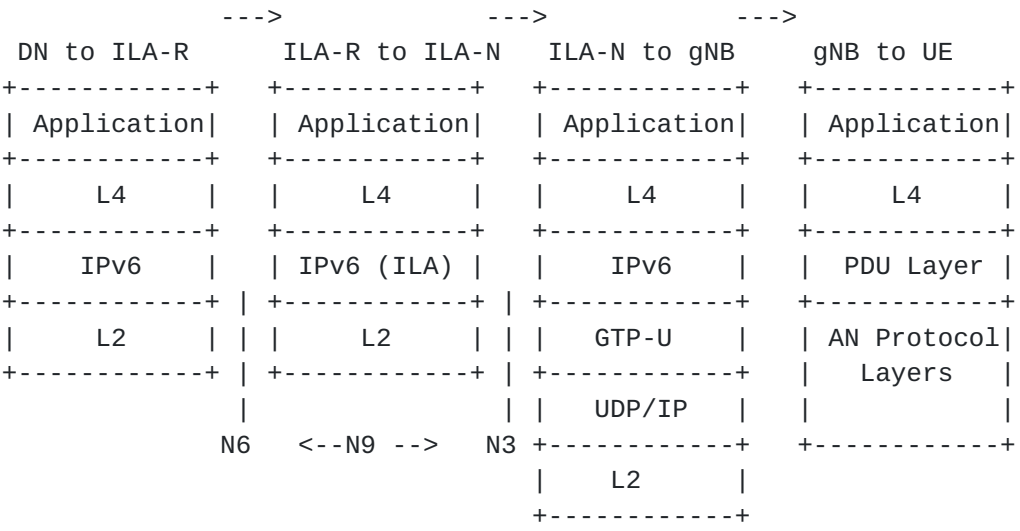


Figure 3: ILA and protocol layer in 5G

## 7.3 Control plane between ILA and network

ILA is a consumer of several 5G network services. The service operations of interest to ILA are:

- Nudm (Unified Data Management): Provides subscriber information.
- Nsmf (Service Managment Function): Provides information about PDU sessions.

- Namf (Core Access and Mobility Function): Provides notifications of mobility events.

ILA-M subscribes to notifications from network services. These notifications drive changes in the ILA mapping table. The service interfaces reference a UE by UE ID (SUPI or IMSI-Group Identifier), this is used as the key in the ILA identifier database to map UEs to addresses and identifier groups. Point of attachment is given by gNB ID, this is used as the key in the ILA locator database to map a gNB to an ILA-N and its locator.

## 7.4 ILA and network slices

Figure 4 illustrates the use of network slices with ILA.

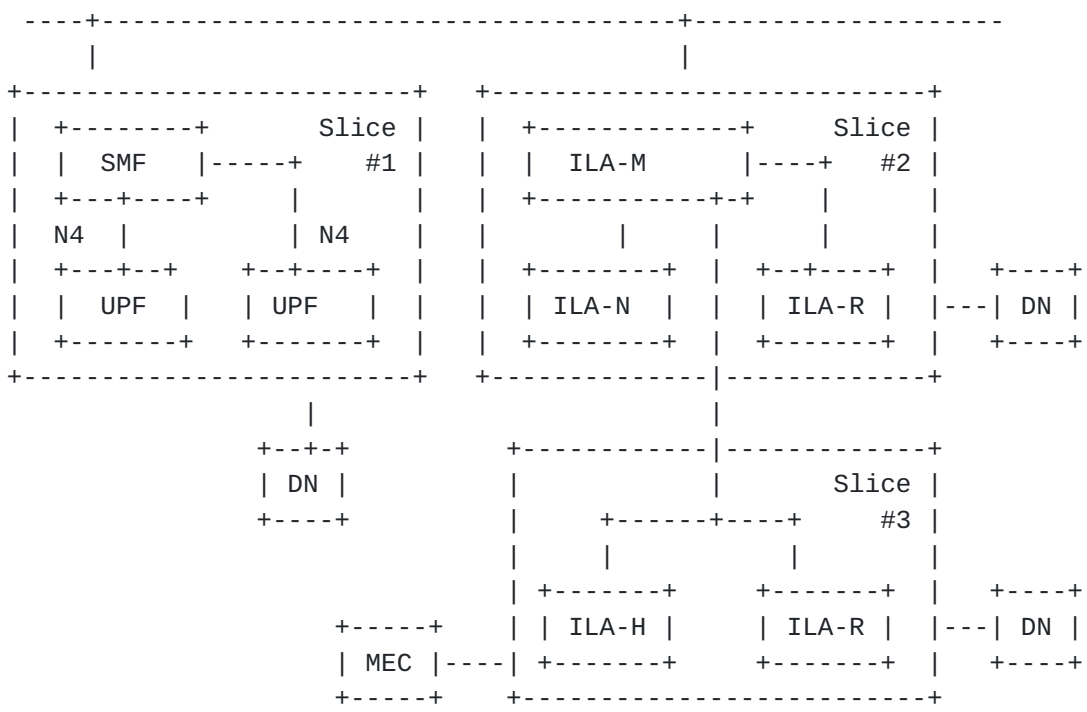


Figure 4: ILA and network slices in 5G

In this figure, slice #1 illustrates legacy use of UPFs without ILA in a slice. ILA can be deployed incrementally or in parts of the network. As demonstrated, the use of network slices can provide domain isolation for this.

Slice #2 supports ILA. Some number of ILA-Ns and ILA-Rs are deployed. ILA transformations are performed over the N9 interface. ILA-Rs would be deployed at the N6 interface to perform transformations on packets received from a data network. ILA-Ns will be deployed deeper in the

network at one side of the N3 interface. ILA-Ns may be supplemented by ILA-Rs that are deployed in the network. ILA-M manages the ILA nodes and mapping database within the slice.

Slice #3 shows another slice that supports ILA. In this scenario, the slice is for Mobile Edge Computing. The slice contains ILA-Rs and ILA-Ns, and as illustrated, it may also contain ILA_Hs that run directly on edge computing servers. Note in this example, one ILA-M, and hence one ILA domain, is shared between slice #2 and slice #3. Alternatively, the two slices could each have their own ILA-M and define separate ILA domains.

## 8 Security considerations

A mobile public infrastructure has many considerations in security as well as privacy. Fundamentally, a system must protect against misdirection for the purposes of hijacking traffic, spoofing, revealing user identities, exposing accurate geo-location, and Denial of Service attacks on the infrastructure. Security must be considered for both the data and control planes.

### 8.1 Data plane security

The ILA data plane must protect against spoofing, inadvertent leakage of sensitive information, and Denial of Service attack.

Locator addresses must be contained within an ILA domain. ILA to SIR transformations MUST be performed before allowing a packet to egress an ILA domain.

Nodes outside of an ILA domain MUST NOT be permitted to send packets into the domain that have an ILA address in either the source or destination. A stateless firewall at the domain boundary can be used to drop such packets. Note that in the ILA protocol, ILA addresses are not used in source addresses.

Section 5.6 describes the handling of ICMP with ILA to avoid leaking locators outside the ILA domain.

When a cache is employed that is populated by events from an outside party there is the possibility of Denial of Service attack. A conceptual attack on ILA-N would be for an attacker will flood its link with packets destined to random SIR addresses. The intent is to exhaust the cache memory so that legitimate traffic is blocked from using the cache and hence needs to take sub-optimal routing. The attack can also generate vast numbers of control messages to DOS the infrastructure.

It is recommended that ILA redirects, as opposed to query model or pub/sub, is used to mitigate attacks. The reasoning is:

- On a cache miss the packet is forwarded and might encounter a router that sends a redirect. The packet itself implies a request for a mapping so no additional control message are needed.
- An ILA router will send a redirect only if there is a mapping to the destination. It doesn't sent negative information. In particular, if the identifier space is reasonably sparse a random address attack will not be very effective.
- A cache entry is created only when a valid redirect is received. This can be contrasted with a query mechanism that might create state for pending resolutions.
- An inactivity timeout can used to evict cache entries. Given the incoming packet rate and a preferred inactivity timeout, a cache can be sized to absorb an attack.
- An ILA router may apply its knowledge to rate limit, prioritize, and shape the use of redirects to manage caches. For instance, an ILA router might identify "hot nodes" in the network that receive a lot of traffic and provide the most benefit when cached in forwarding nodes.

## 8.2 Control plane security

A mapping system contains sensitive privacy information that could be used to make inferences about user's identity or their geo-location. This information needs to be protected.

Mapping protocols must be secured to prevent an attacker from injecting mapping entries to redirect traffic to their own devices. To this end, mapping protocols for ILA are intended to use TCP. The statefulness of TCP deters spoofing of messages and allows for privacy and identity verification in the form of X.509 certificates. The control protocol includes "secure" redirects that must be authenticated to originate from a legitimate ILA router.

Mapping protocols must also be resilient to DOS attack, especially in a scenario where a cache of mappings is being employed. Such a cache might be populated in response to the activities of a third party (for instance an application sending packets to different destinations). An attack on the cache whereby an attacker attempts to fill the cache with entries to random destinations must be mitigated. The recommendation of ILA is to use "secure redirects" as a scalable

and secure means to populate a forwarding cache.

Write access to the ILA mapping database must be strictly controlled. In the ILA architecture only ILA-Ms write to the mapping database. Write access to the database should require strong credentials, validation of each operation, and encryption and authentication of operations being sent over the network.

Read access the ILA routing database should also be controlled. Devices should only access data on a "need to know" basis. For instance, ILA routers might need identifier to group mappings to perform forwarding, but they should not need to retrieve all the identifiers for a group. The latter information can be contained in the ILA-Ms.

### 8.3 Privacy in address assignment

A node may use multiple addresses to prevent inferences by third parties that break privacy. Properties of addresses to maintain strong privacy are:

- They are composed of a global routing prefix and a suffix that is internal to an organization or provider. This is the same property for IP addresses [RFC3513].
- The registry and organization of an address can be determined by the network prefix. This is true for any global address.
- The organizational bits in the address should have minimal hierarchy to prevent inference. It might be reasonable to have an internal prefix that divides identifiers based on broad geographic regions, but detailed information such as accurate location, department in an enterprise, or device type should not be encoded in a globally visible address.
- Given two addresses and no other information, the desired properties of correlating them are:
  - It can be inferred if they belong the same organization and registry. This is true for any two global IP addresses.
  - 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 (e.g. are in the same shard).
  - No other correlation can be established. For example, it cannot be inferred that the IP addresses address the same

node, the addressed nodes reside in the same subnet, rack, or department, or that the nodes for the two addresses have any geographic proximity to one another.

Ostensibly, assigning a /64 prefix to a node is good for security. The end device can create its own random addresses in the low order sixty-four bits which mitigates address scanning attacks. However, the upper sixty for bits of the address become a static identifier for the node that potentially allows DOS on the device, as well as third party correlations on addresses that deduce that different flows are sourced from the same user.

[RFC4941] recommends rotating addresses to protect privacy. In the case of sixty-four bit address assignments this would entail that a new prefix for the device is periodically requested. There is no recommendation for the frequency of address change and there is no quantitative description of the effects of periodic address change.

For maximum privacy, a different address could be used for each connection. If this were done for every connection in the network, it would create network state for each connection (note that is sort of thing already exists with stateful NAT). Scaling the mapping system to accommodate this is challenging. One alternative to be investigated is use a reversible cryptographic hash to aggregate identifiers and reduce the number of mappings needed.

# 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, <https://www.rfc-editor.org/info/rfc2119>.

[ILA] Herbert, T., and Lapukhov, P., "Identifier Locator Addressing for IPv6" draft-herbert-intarea-ila-00

[ILAMP] Herbert, T., "Identifier Locator Addressing Mapping Protocol" draft-herbert-ila-ilamp-00

## 9.2 Informative References

[RFC3513] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, DOI 10.17487/RFC3513, April 2003, <https://www.rfc-editor.org/info/rfc3513>.

[RFC4941] Kolkman, O. and R. Gieben, "DNSSEC Operational Practices", RFC 4641, DOI 10.17487/RFC4641, September 2006, <https://www.rfc-editor.org/info/rfc4641>.

[ILAGRPS] Herbert, T., "Identifier Groups in ILA", To be published

[BGPILA] Lapukhov, P., "Use of BGP for dissemination of ILA mapping information" draft-lapukhov-bgp-ila-afi-02

[3GPPTS] 3rd Generation Partnership Project (3GPP), "3GPP TS 23.502", http://www.3gpp.org/DynaReport/23-series.htm

## Authors' Addresses

Tom Herbert
Quantonium
Santa Clara, CA
USA
Email: tom@quantonium.net

Kalyani Bogineni
Verizon
One Verizon Way, Basking Ridge, NJ 07920
USA
Email: kalyani.bogineni@verizon.com