

Identifier Locator Addressing with IPv6

(Presenting in LISP WG)

draft-herbert-nvo3-ila-04

Tom Herbert <tom@herbertland.com>

Overview

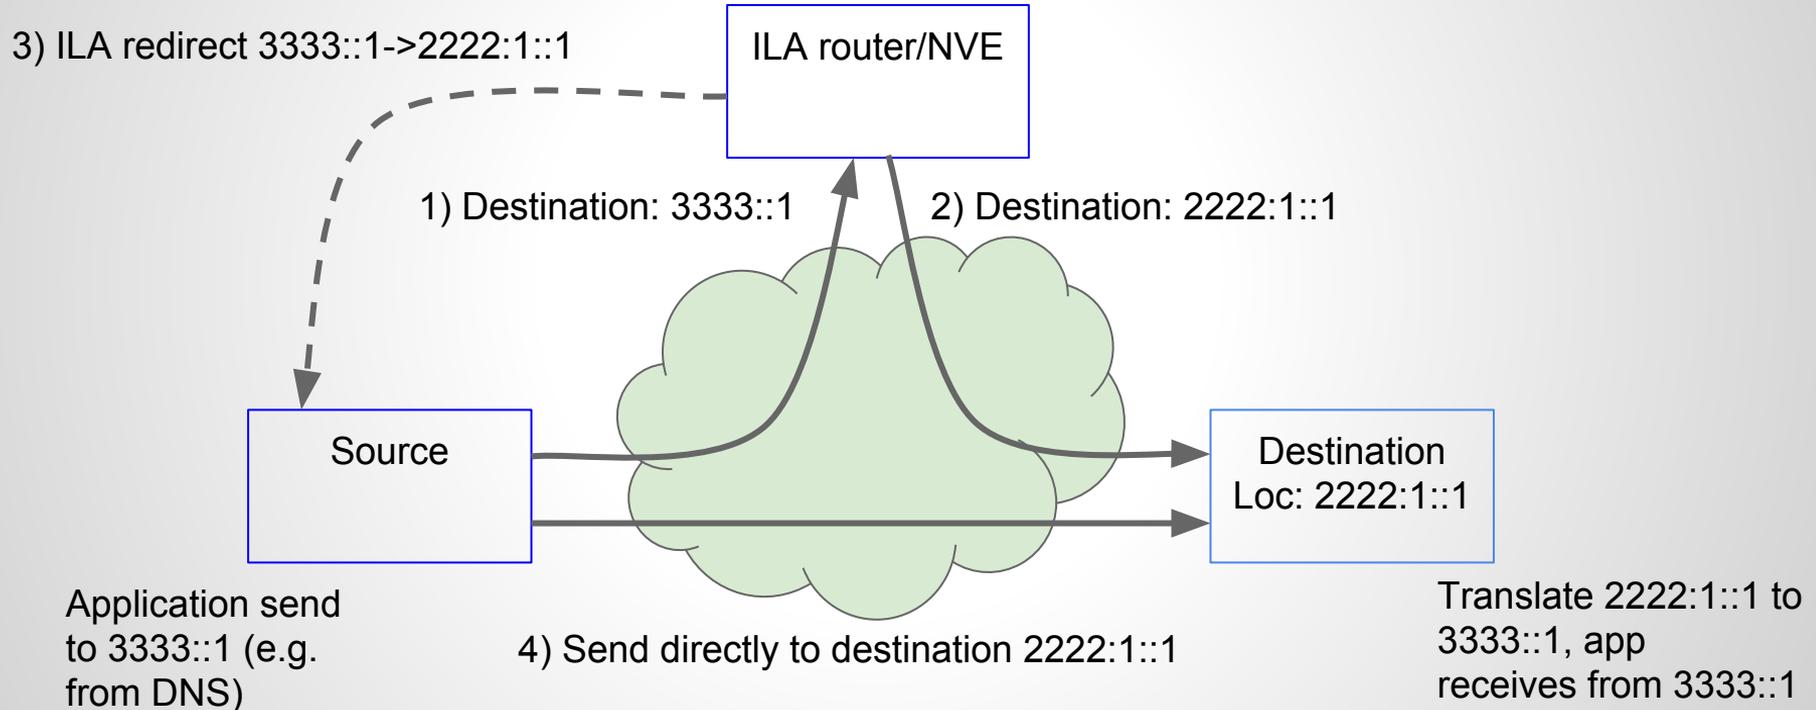
- ID/locator split without encapsulation
- Split IPv6 address into two parts
 - Locator: “where” of an entity
 - Identifier: “who” of the entity
- Applications see identifier (SIR) addresses
- How to send
 - Lookup identifier in hash table
 - Overwrite SIR prefix with found locator
 - Csum neutral translation to keep L4 csums correct

Address split



- **Locator**
 - 64 bits identifier of physical hosts
 - Routable
 - Not used as connection endpoint
- **Identifier**
 - 64 bit logical endpoint address of virtual node
 - Routable to an translator (NVE)
 - Used as connection endpoint
 - Typed to allow different modes

Flow example



Checksum neutral mapping



- Like RFC6296
- Good csum on wire without needing to access L4 headers
- Use low order 16 bits in identifier as a checksum adjustment value (SIR->ILA)
 - $\text{csum-adjust} += \text{csum_diff}(\text{SIR-prefix} - \text{locator})$
- Reverse operation going ILA->SIR

ILA use cases

USE CASE	DESCRIPTION	SCALING # NODES	RATE/SEC OF MAP UPDATE
DC task virtualization*	Assign every task an IP address	10's of millions	1000's
DC virtualization	Assign "everything" and IP address	Up to 10's of billions	Millions
Multi tenant virtualization	VNID + Vaddr (VMs)	10's of millions	1000's
5G mobility	Every UE has identifier	Billions maybe more	Millions

* Currently being deployed

Advantages of ILA

- Not encapsulation
 - No on the wire overhead
 - No MTU, UDP checksum, or other tunneling issues
- No changes to transport layer
 - Checksum neutral translation
 - Application, DNS only deal with untranslated globally router addresses
- Open source implementation
 - Linux host side implementation
 - ILA router in XDP, VPP

Challenges for ILA

- ICMP

- Hosts can get ICMP errors for packets with ILA destinations
- If host is ILA-aware attempt reverse translation there
- If ILA router is in return path of ICMP error, reverse there as in NAT (RFC5508)

- Multicast

- Can't modify destination, modifying source address would be problematic at non-ILA receiver
- Conclusion: ILA not appropriate with multicast

Status

- Asking int-area to take up draft as WG item
- Deploying for task virtualization @FB
- ILA router development
 - BPF/XDP program
 - Control plane
 - BGP initially
 - Resolution/redirect protocol for ILA hosts
 - IDEAS

Relationship to LISP

- Possibly complementary, logically a compressed LISP?
- Control plane (identifier to locator mappings) seem like something to be common
 - Is LISP control plane extensible
 - How does this relates to IDEAS

Thank you!