NIST National Institute of Standards and Technology

INFORMATION TECHNOLOGY LABORATORY

Enhanced Efficiency of Mapping Distribution Protocols

Kotikapaludi Sriram, Young-Tak Kim, and Doug Montgomery

National Institute of Standards and Technology

Contact: ksriram@nist.gov

LISP WG Meeting, IETF 77, Anaheim, CA, March 2010

(Acknowledgement: Thanks to Dino Farinacci and Lixia Zhang for helpful comments and suggestions.)

Background

- This work was originally presented in RRG in July 2007 at the Dublin IETF Meeting.
- Detailed updated document is at:

http://www.antd.nist.gov/~ksriram/NGRA_map_mgmt.pdf

- Darrel already discussed (on Monday) the possibility of holes in EIDs and presented ways to tailor mapping response to cope with them in LISP
- Here we take a more detailed view of the problem space and provide enhanced solution

Real-World Example with a PI Address

Announced in BGP-4:	
129.6.0.0/17	Origin: AS49
129.6.0.128/17	Origin: AS49
129.6.112.0/24	Origin: AS10886

EID to Locator Mapping:	
EID:	ETR (equivalent)
129.6.0.0/18	ETR49
129.6.64.0/19	ETR49
129.6.96.0/20	ETR49
129.6.112.0/24	ETR10886
129.6.113.0/24	ETR49
129.6.114.0/23	ETR49
129.6.116.0/22	ETR49
129.6.120.0/21	ETR49
129.6.128.0/17	ETR49

Measurement of # Prefix Holes

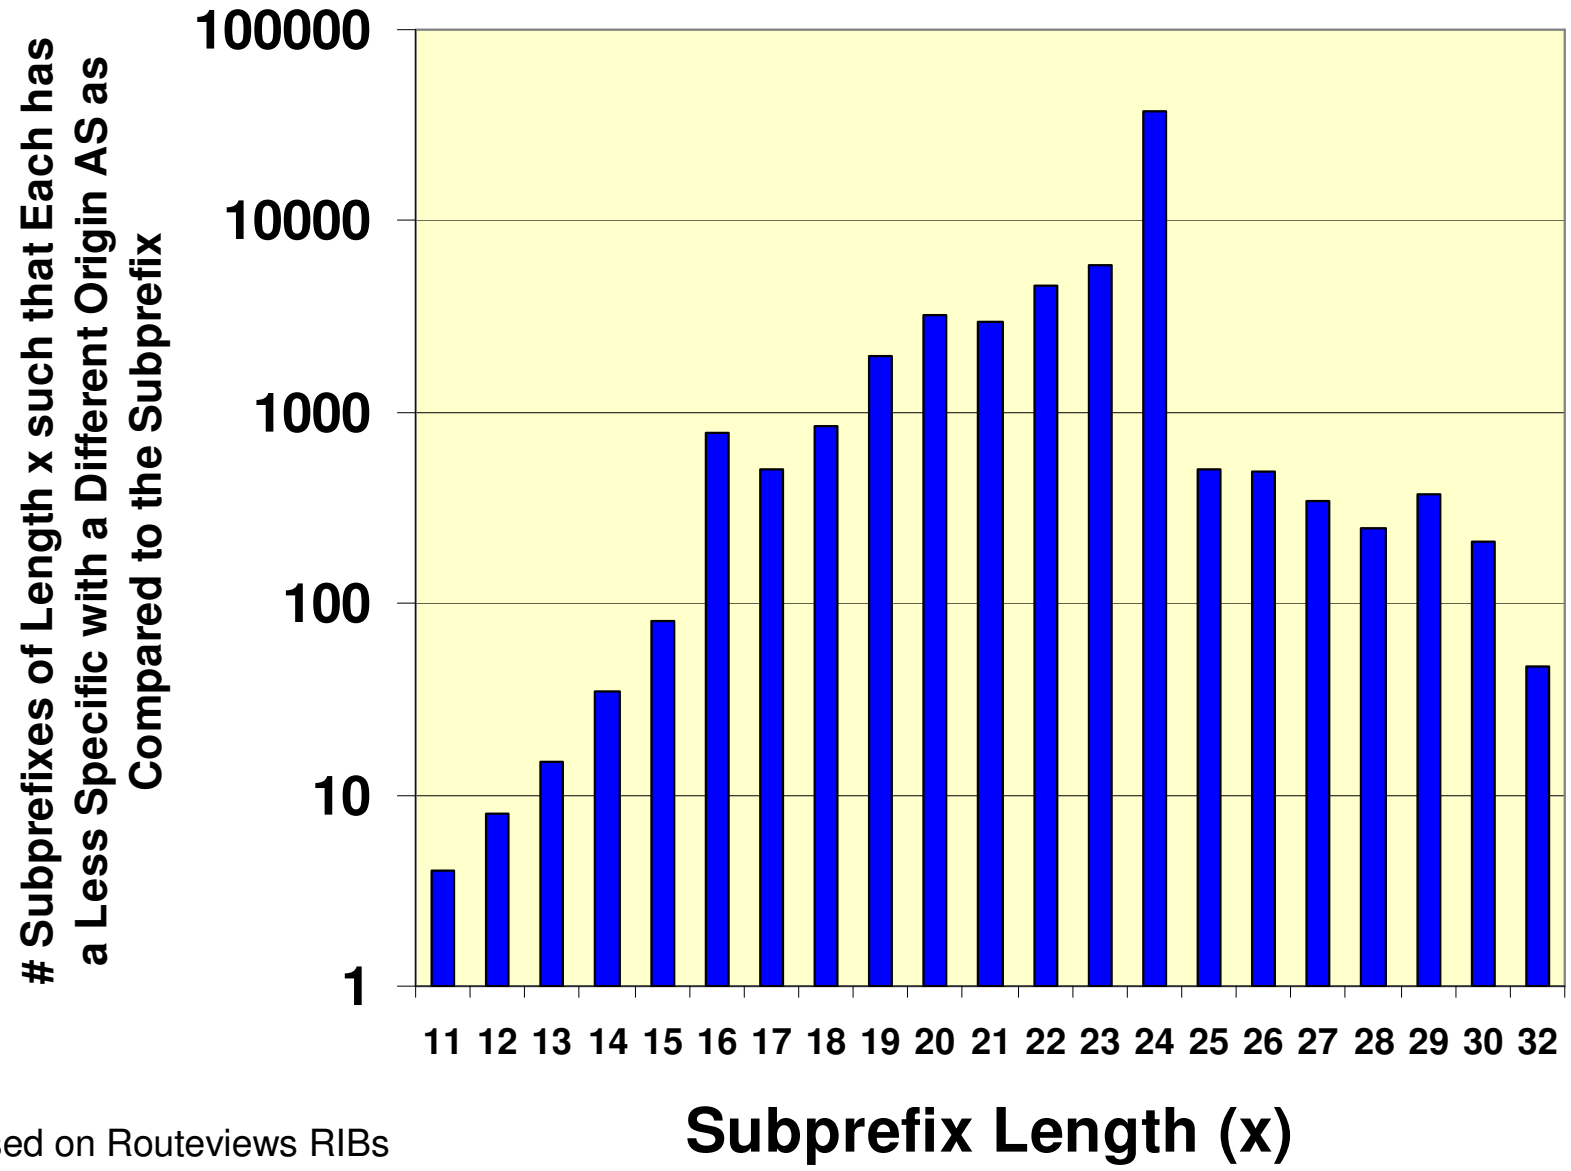

Based on Routeviews RIBs trace data – Feb 2010

Case 1: More-Specifics (Holes) Absent

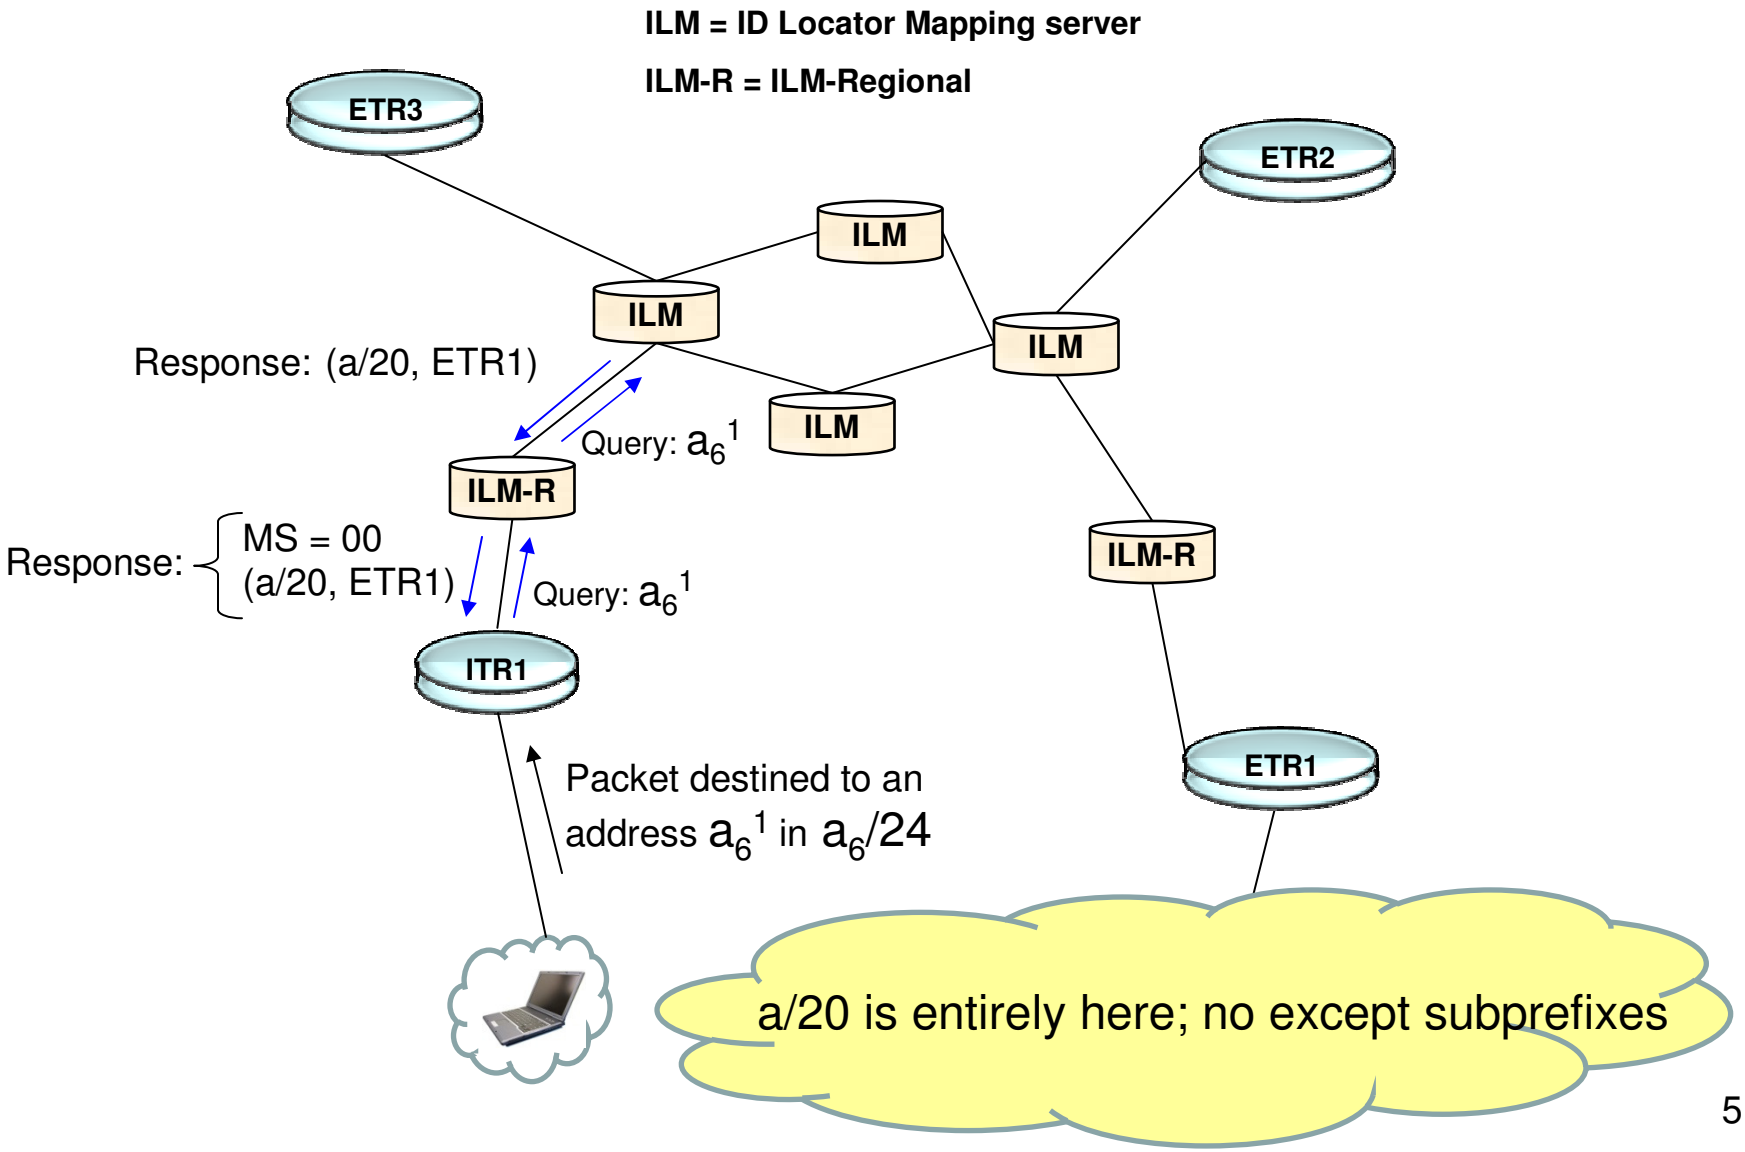

Case 2: All More-Specifics Communicated

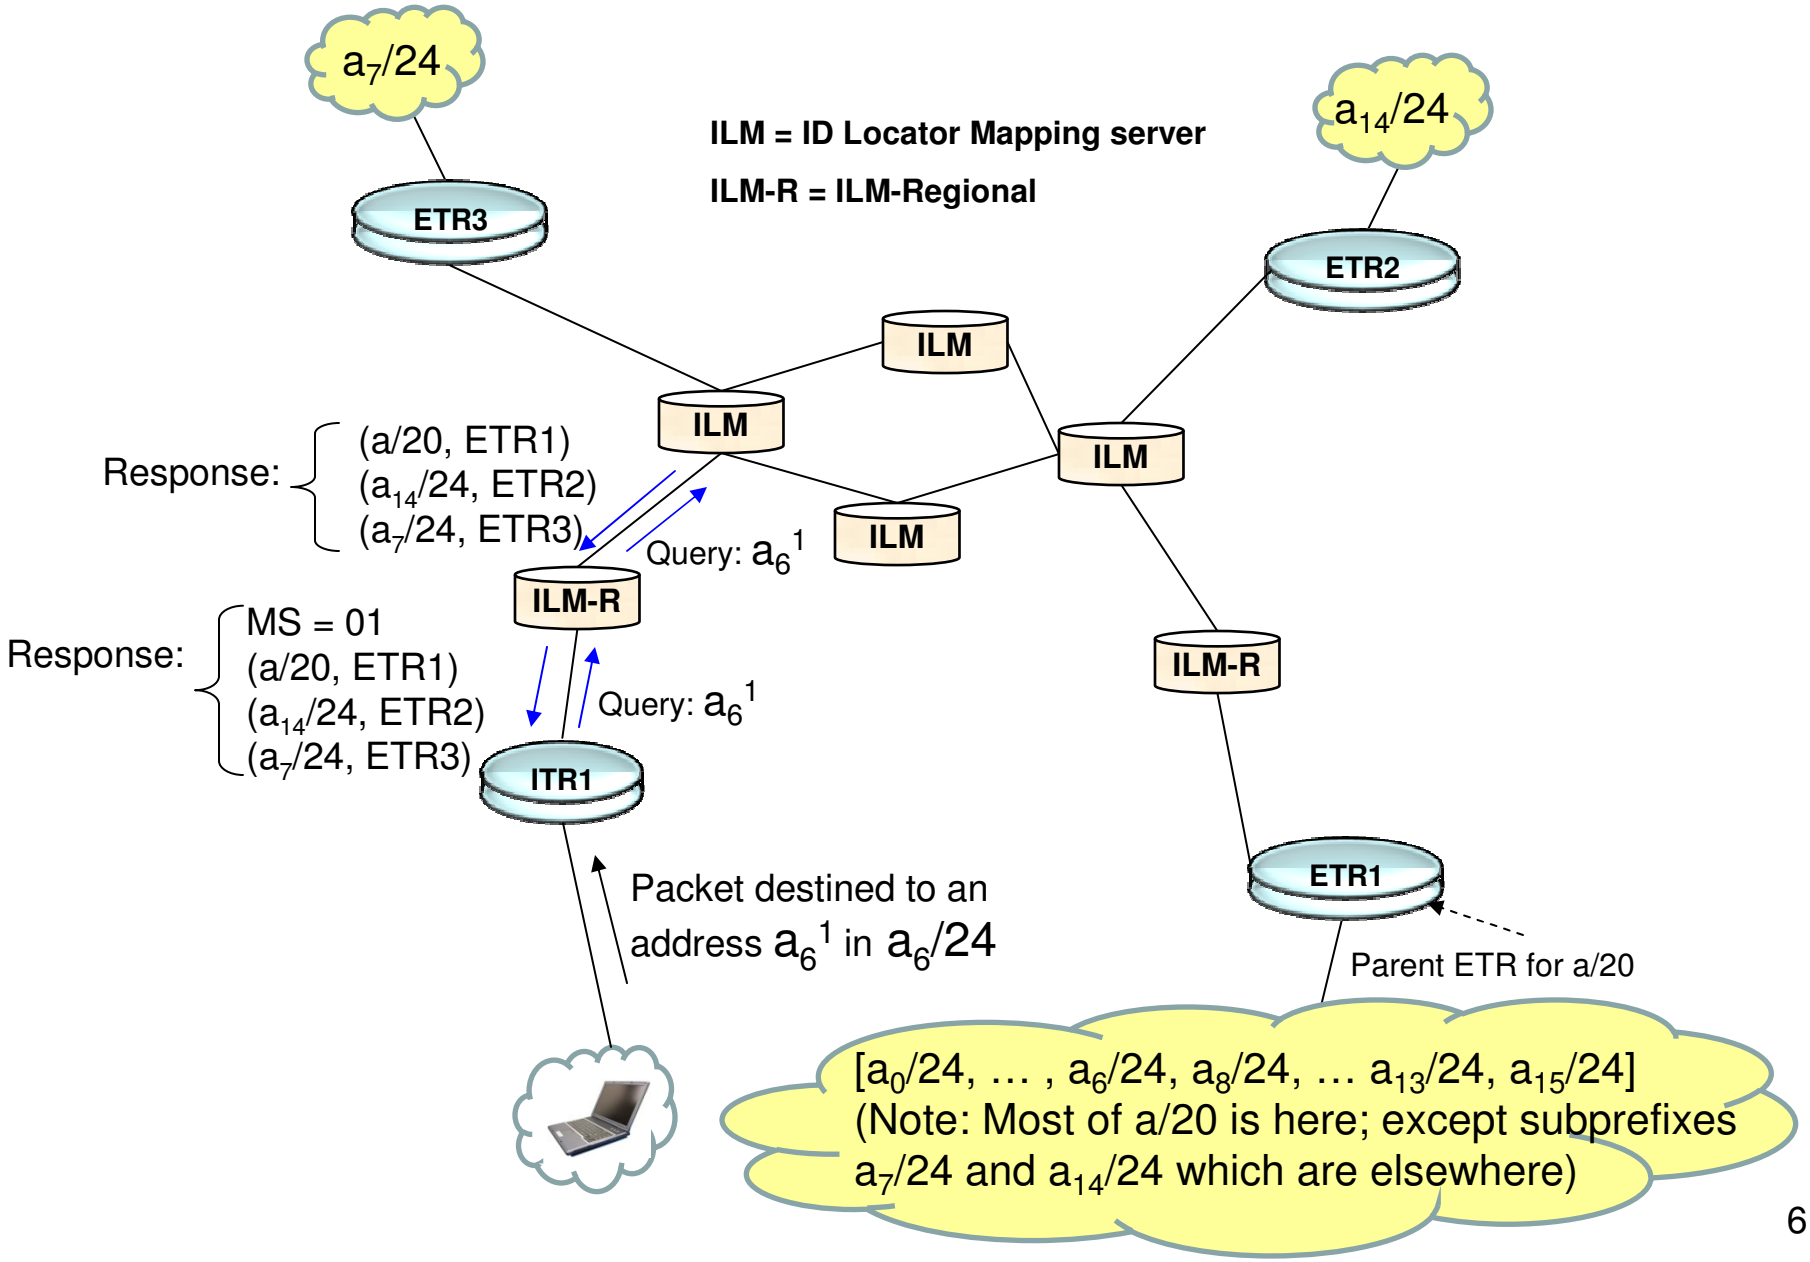

Case 3: Exception More-Specific Communicated without ETR Info (Lots of Mobile Nodes)

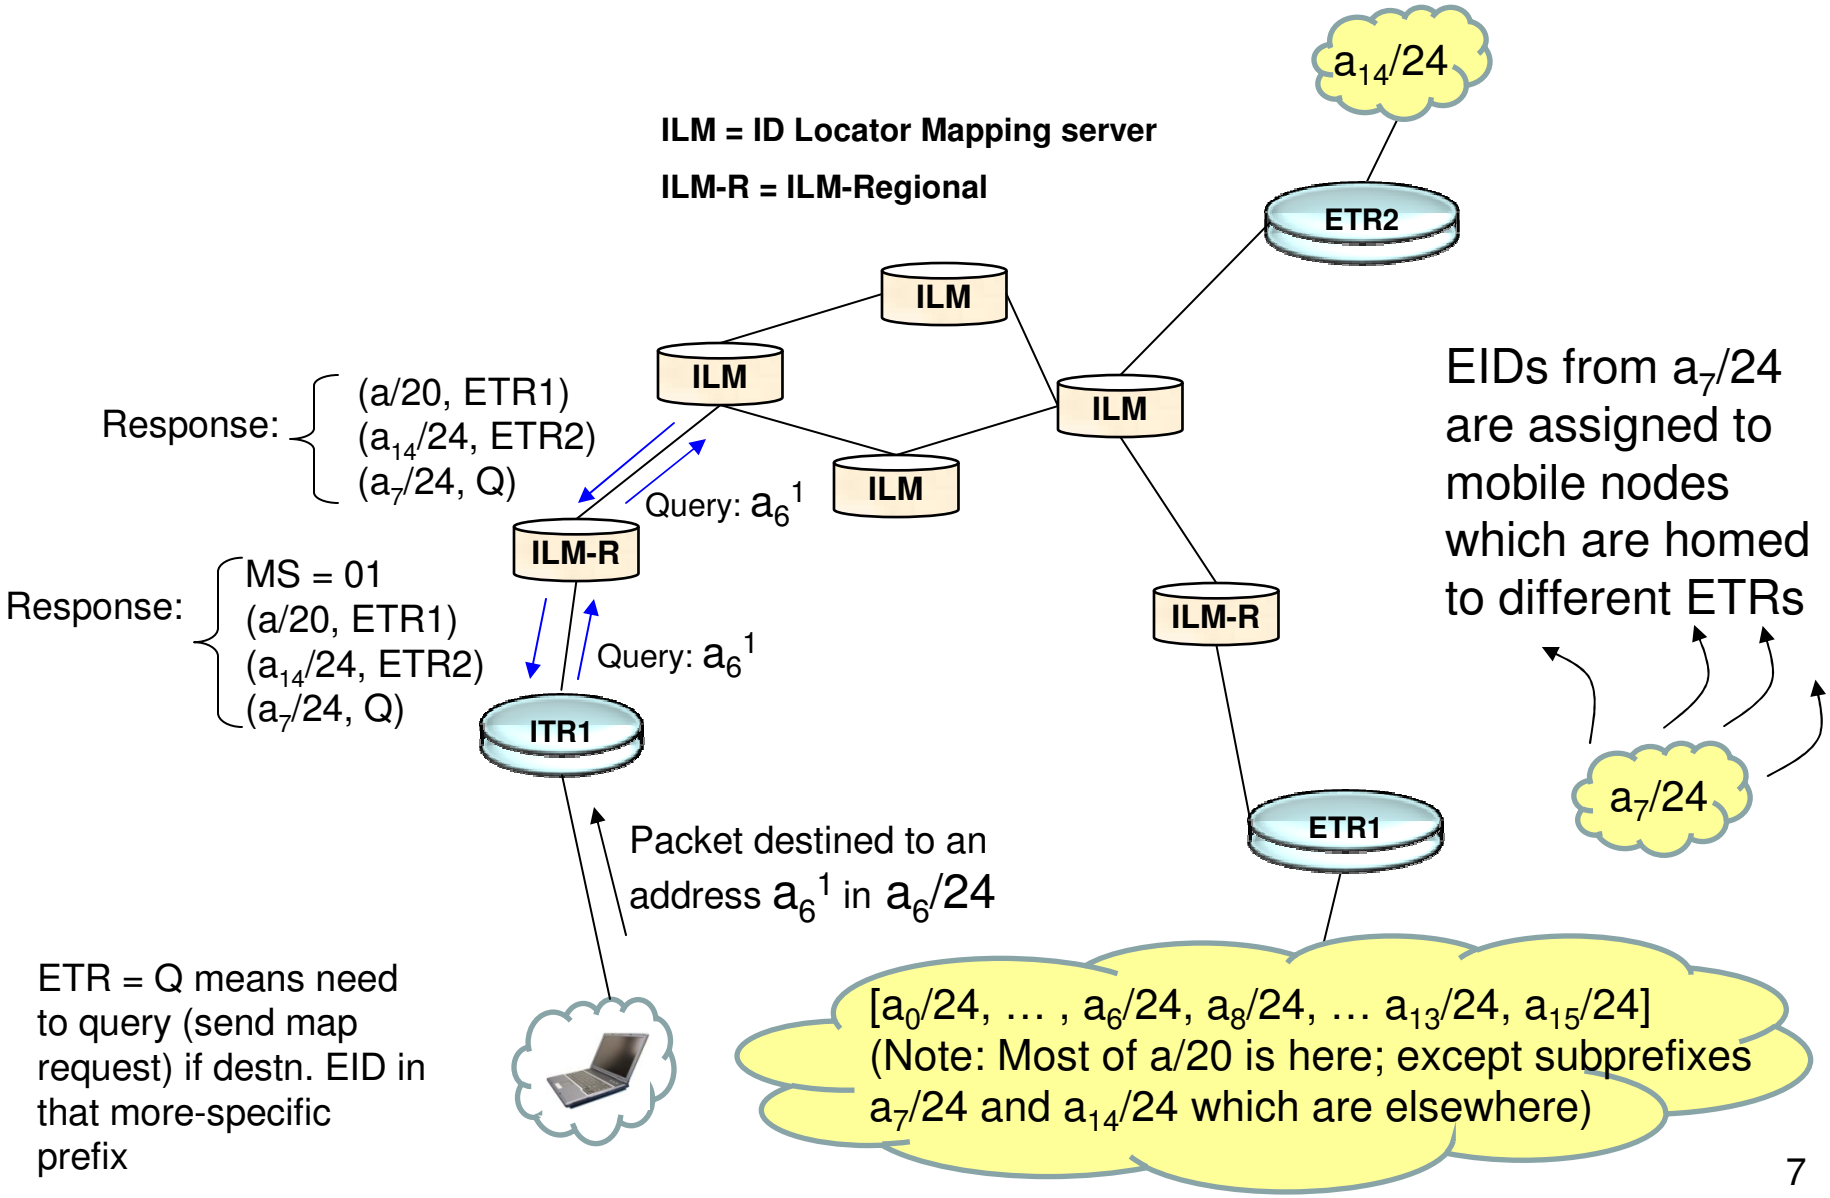

Case 4: NML Communicated; More-Specifics Not Communicated

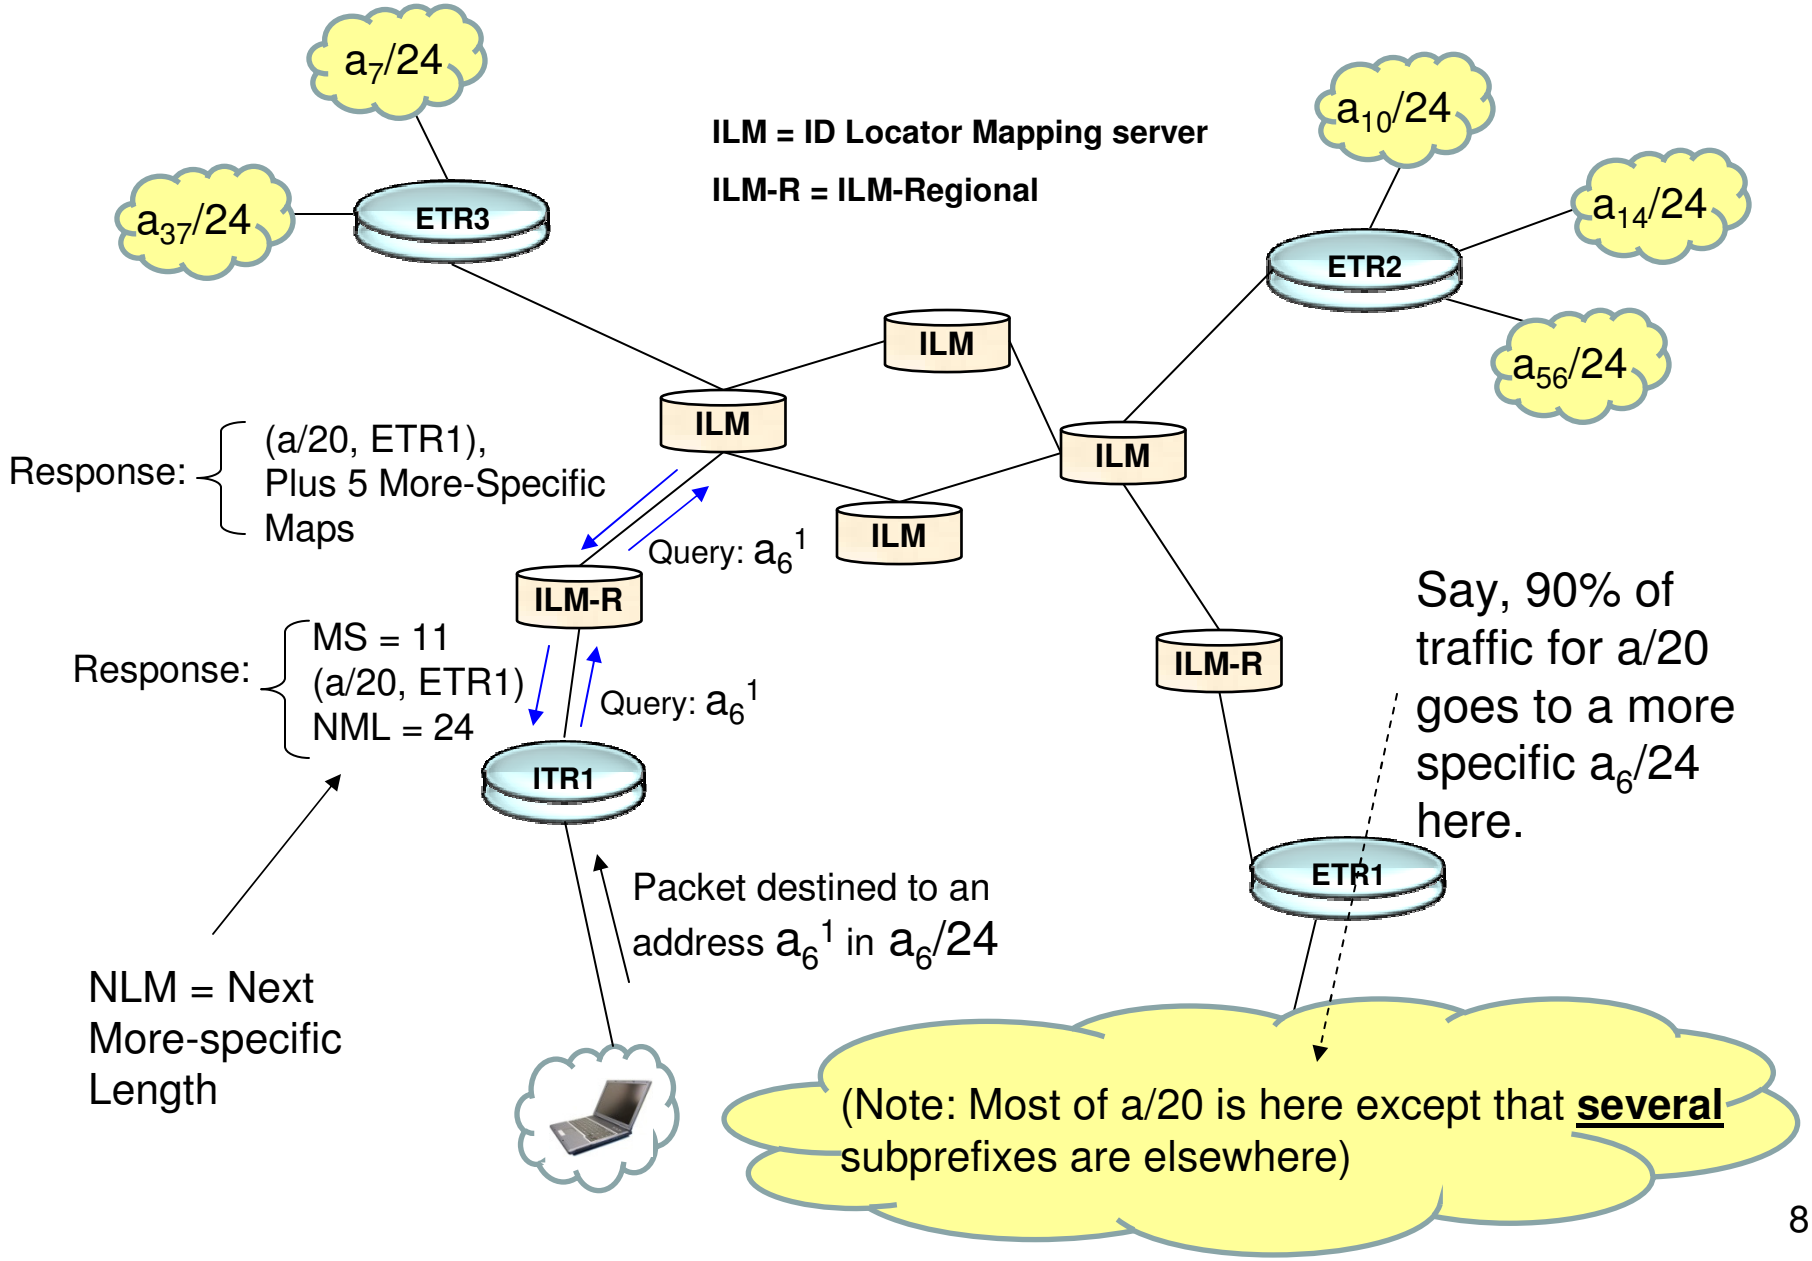

Conceptual Format for the Enhanced Map Response

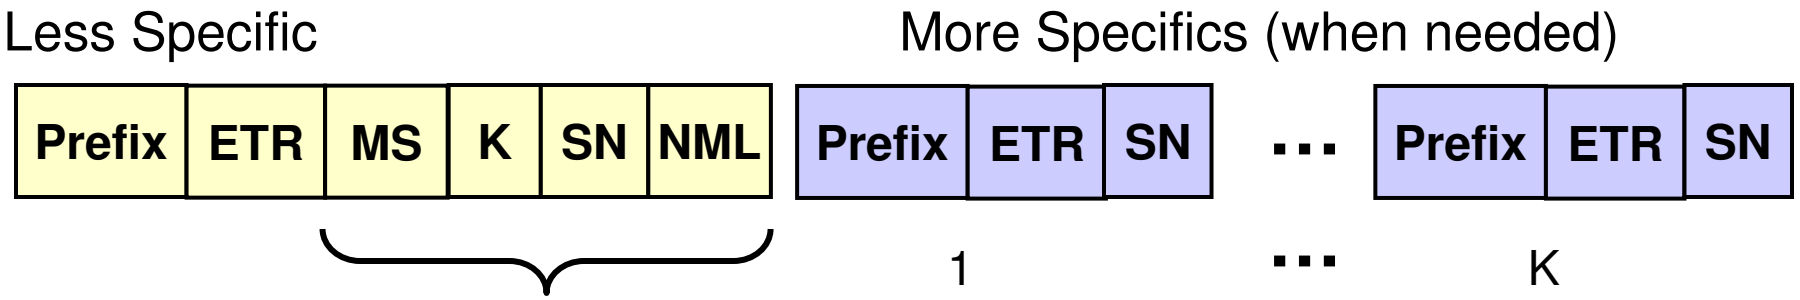

MS = More Specific indicator

K = # Maps to follow

SN = Sequence Number ($\leq$ K +1)

NML = Next longer Mask-Length

If ETR = Q, it means ITR needs to query (send map request) for destn. EID in that more-specific prefix

Algorithm Description

MS =	K =	SN =	NML =	Interpretation	Application Notes
00	1	1	Don't Care	Map response has no exceptions.	Normal case; ITR caches one map.
01	$k + 1$	i ($1 \leq i \leq k + 1$)	Don't Care	Map response has exceptions; Additional k maps for the exception subnets are also included.	ITR caches the main map and also those for all exception subnets.
10	$k + 1$	i ($1 \leq i \leq k + 1$)	Don't Care	Map response has exceptions; Additional k maps for the exception subnets follow automatically but the ETR information for one or more specific subnets is "Q (Query)" (because the specific subnet is further split, i.e., multiple sub-subnets exist in said subnet with different ETRs).	This is very useful, for example, when a /24 subprefix is split from a corporate /16 prefix and company's mobile devices are allocated IP addresses from that one /24.
11	1	1	m	Map response has exceptions; Additional maps are not provided because # subnets (at next longer mask-length value) exceeds threshold (H). The next longer mask-length value NML = m is provided here. No need to request new map if the first m-bits of destination EID match the same for the first packet.	This is a very useful when, for example, an Org-A has a /16 prefix; bulk of it resides at headquarters (ETR1); several /24s subprefixes (in the /16 prefix) are homed to other ETRs elsewhere; majority of the traffic goes to a specific /24 located at the headquarters (ETR1).

Each of the (k+1) maps may be assigned a serial number (SN)