LOMO: Location Resolution of Mobile Hosts Using Internet Routing Structures

Ted Taekyoung Kwon
Computer Science and Eng.
Seoul National University
tkkwon@snu.ac.kr

Jaeyoung Choi
Computer Science and Eng.
Seoul National University
jychoi@mmlab.snu.ac.kr

Eunsang Cho
Computer Science and Eng.
Seoul National University
escho@mmlab.snu.ac.kr

Yanghee Choi
Computer Science and Eng.
Seoul National University
yhchoi@snu.ac.kr

ABSTRACT

The DNS has performed the resolution for quasi-static hosts so far. However, mapping from the names of mobile hosts to their locators is increasingly important. We propose a new resolution architecture, LOMO, that keeps track of the locators of mobile hosts as they are changing the points of attachment. LOMO takes a host-based approach and leverages current IP routing structures to route the resolution traffic in a scalable manner and to balance the workload of resolution.

Categories and Subject Descriptors

C.2 [**Computer-Communication Networks**]: Network Architecture and Design

General Terms

Design

Keywords

Resolution, Mobility, Locator, Location

1. INTRODUCTION

The domain name system (DNS) has performed the resolution between a hostname and its IP address successfully. The DNS assumes that hosts with domain names are quasi-static, and hence it works well only if caching the IP addresses is valid for a sufficiently long time. However, we believe the DNS should be fundamentally changed to provide *host mobility*. We argue that mobile hosts be allowed to have domain names (along with [2,10]) since the high computing power and broadband mobile access can enable mobile devices to provide user generated content, e.g. mobile web server. If mobile hosts have domain names, the current DNS structure may not be able to handle the dynamic update/query traffic to locate mobile hosts.

Even though increasingly more Internet hosts become mobile, *host mobility* is not supported well in the current Internet. The related studies to solve mobility can be classified into two categories: (i) network layer proposals [5, 8], and (ii) host-based proposals [2, 10].

Among the network layer proposals, the Mobile IPv6 solution enables us to support mobility at the network layer with route optimization (i.e. without indirection) [8]. However, it focuses on unpopular *macro-mobility*, and hence deploying the Mobile IP solution is still questionable. In addition, the Proxy Mobile IP solution [5] may not be scalable as the traffic of mobile hosts keeps increasing [6] [1]because it needs some anchor points relay the traffic to mobile hosts as they are changing their points of attachment.

The host-based approaches provide host mobility at the transport layer or underneath. However, they assume that the DNS performs the resolution and focus on how to modify the TCP/IP kernel stack and prevent security vulnerabilities. To the best of our knowledge, there is no solid study on how to modify the DNS or how to deploy a new infrastructure to handle the resolution traffic in a scalable fashion.

We propose a new resolution architecture, LOMO, by which the Internet can support host mobility. LOMO takes the lookup-by-name approach; that is, it performs the mapping between the names of mobile hosts and

[1]According to [4], mobile data traffic has been almost tripling each year lately.

AINTEC'11, November 9–11, 2011, Bangkok, Thailand.

their locators. As it would be too costly or even infeasible to deploy a new infrastructure like the DNS from scratch, LOMO leverages the Internet routing structures to minimize the upfront cost. As proposed in [2, 10], we assume that end hosts deal with the mobility signaling as follows. First, a mobile host informs the infrastructure of its current locator whenever its point of attachment is changed. Likewise, if the locator of a mobile host is changed in the middle of a flow, it will notify the other counterpart of the change.

2. LOMO: LOCATION RESOLUTION OF MOBILE HOSTS

LOMO is designed based upon the current IP networking; thus, a mobile host configures its new IP address (i.e. the current locator) whenever it moves to a new IP subnet. However, LOMO needs some extensions for the current IP routing protocols.

2.1 Routing Protocol Extensions

We need 3 extensions: (i) when a BGP speaker in an origin AS constructs a BGP UPDATE message, its IP address is recorded and retained while the BGP routing information is propagated throughput the Internet; (ii) the IP address of a BGP speaker of an origin AS, as well as the origin AS number (ASN) in the BGP UPDATE message is propagated to all the other ASs. Even though BGP UPDATE messages cannot sometimes cross the boundary of two ASs depending on their relation, we assume that the above information (the ASN and the BGP speaker's IP address) is propagated between the two ASs regardless of their relation; (iii) a BGP speaker of an AS also participates in intra-domain routing. When a router constructs an intra-domain routing message, its IP address is recorded, so that the BGP speaker learns the IP addresses of the routers in its AS.

The first two extensions are needed to ensure the following. Every BGP speaker can learn (i) which ASs are active, and (ii) at least one BGP speaker's IP address for each active AS. The active ASs are sorted in the same way (say, in increasing ASN order) across all the BGP speakers in the Internet. Let A be the number of active ASs. Suppose $A = 2^{12}$ ASs are currently up and running in the Internet. Then, each AS takes responsibility for the resolution of $1/2^{12}$ of the ID space. Assume that the three smallest active ASNs are 1, 3, and 6. In this case, the resolution of IDs that starts with 0x000 will be assigned to ASN 1. With $A = 2^{12}$, the first 12 bits are called the first prefix (FP), whose length is determined by $\lfloor \log_2 A \rfloor$.

With the third extension, a BGP speaker learns how many and which routers are running in its AS. Let R_i be the number of routers in AS i. Then, each of R_i routers will be responsible for $1/(A * R_i)$ of the ID space. If AS i is in charge of the IDs starting with 0x002 and $R_i = 16$, the 16 routers are responsible for IDs starting with 0x0020, 0x0021, ... , 0x002F, respectively. The $\lfloor \log_2 R_i \rfloor$ bits after the FP are called the second prefix (SP) of AS i.

2.2 LOMO Operations

Let us now describe how LOMO performs the location resolution of a mobile host. Note that a mobile host informs the resolution infrastructure of its new locator in the same way.

Suppose a solicitor wishes to find out the locator of a particular mobile host. First, it uses a hash function (e.g. MD-5 and SHA-1) to map the hostname[2] into a fixed-length identifier (ID). Then, it constructs a resolution request (ResReq) packet that contains the host's ID and the solicitor's address, and sends the ResReq message to its access router.

When the access router of the solicitor receives the ResReq message, it will look at the FP of the ID. We assume all the access routers know which FP is assigned to their AS. (For instance, the FP is distributed across the AS by extending its intra-domain routing protocol.) Suppose the FP does not belong to the AS of the access router. The ResReq packet will then be sent to one of the BGP speakers of the AS. On receipt of the ResReq packet, the BGP speaker looks at the FP and figures out that AS i is responsible of the ID in the packet. It will be sent to one of BGP speakers of AS i.

The ResReq packet now arrives at the BGP speaker of AS i. The BGP speaker maintains the list of the IP addresses of the routers in AS i. It looks at the SP (of AS i) and figures out which router is responsible of the ID in the packet. Finally, the ResReq packet is sent to the router that has the locator of the ID. The router will reply with a Resolution Response (ResRsp) packet containing the locator to the solicitor.

For sake of simplicity, we assume a router performs the resolution by itself. More practically, a resolution server may co-locate with the router, or it can be located in the same subnet with the router.

Figure 1 illustrates how LOMO performs the resolution service. Suppose there is a solicitor, which wishes to find out a host H_1, connected to an access router in AS 20. The solicitor hashes the name of H_1 into a fixed-length identifier (ID). Suppose the ID is 0b0101000011. Then, the solicitor sends a resolution request (ResReq) packet to its access router. In the ResReq packet, the Request (Req) flag is set, and the source and destination addresses are the solicitor's and the access router's addresses, respectively. The ID and the solicitor's address are also included. Recall that the access router

[2]It can be a domain name or a session initiation protocol (SIP) uniform resource identifier (URI). Or even a home address is fine.

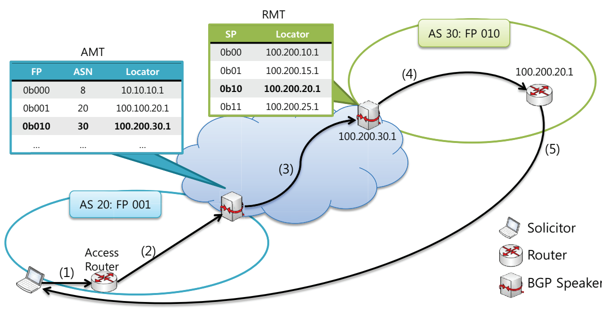

Figure 1: Resolution in LOMO

belongs to AS 20. For sake of exposition, suppose that there are total 8 active ASs, two of which are shown in Figure 1 (ASs 20 and 30). Thus, each AS is in charge of 1/8 of the ID space, and hence the FP length is three bits. Assume that ASs 20 and 30 are responsible of IDs whose FPs are 0b001 and 0b010, respectively.

AS 20 is in charge of resolution service for IDs starting with 0b001, so the ResReq packet should be sent to one of BGP speakers (step (2) in Figure 1). Here, the source and destination addresses are the access router's and the BGP speaker's ones. Upon receipt of the ResReq message, the BGP speaker of AS 20 figures out that AS 30 is responsible for FP 0b010 by looking up its AS mapping table (AMT), then sends the message to the BGP speaker of AS 30 (step (3) in Figure 1). Here, the source and destination addresses are the addresses of the BGP speakers of ASs 20 and 30, respectively.

A router mapping table (RMT) is used to keep track of the running routers in the AS. Suppose there are four running routers in AS 30 now, then the length of the SP is 2 (i.e. the 4th and 5th bits of an ID constitute the SP of AS 30). Based on the RMT, ResReq packet is sent to the router 100.200.20.1 (step (4) in Figure 1). Finally, router 100.200.20.1 receives the ResReq packet, and replies with the resolution response (ResRsp) packet containing the locator of host H_1 (step (5) in Figure 1). Here, the response (Rsp) flag is set to 1, and the locator in the ResRsp packet is the current IP address of host H_1.

3. RESULTS

From a worst case perspective, we compare LOMO, Chord [11], and the DNS in terms of: (i) number of resolution entries, (ii) lookup cost (or delay), and (iii) resolution workload on a server. Recall that a resolution function is assumed to be handled by a router in LOMO. Similarly, Chord nodes and DNS servers are co-located with routers for fair comparison. Thus, we use routers and servers interchangeably in this section.

The number of resolution entries represents how many entries need to be maintained in a router, which indicates the routing scalability. In Chord, each node in the resolution structure maintains the finger table of $logS$ resolution entries where S is the number of routers performing resolution functions. It has $O(logS)$ lookup delay in terms of overlay hop count. A node in the DNS only needs to know its child nodes because of its hierarchical tree structure. However, the number of child nodes in the DNS is highly variable. According to [12], `.com` top level domain servers have around 90 million child nodes at the end of year 2010. Therefore, the number of (worst case) resolution entries for the DNS is set to 90 million here. In case of LOMO, a BGP router should have resolution entries of the number of active ASs and the number of running routers in the AS to which it belongs. From [7], the number of ASNs advertised to be active in BGP in the end of year 2010 is around 35,000. As for the number of routers in an AS, the largest AS in terms of AS connectivity is Level 3 communications, which is connected to 2643 ASs [3]. Thus, we conservatively assume almost 10 times AS connectivity, 25,000, for the worst case (or largest) number of routers in as AS. Thus, the number of resolution entries in a BGP speaker is 60k in LOMO. Figure 2a shows the number of resolution entries in log scale as the number of routers S increases.

Figure 2b shows the lookup cost in terms of the hop count asymptotically as the number of routers S increases. It is known that the average hop count between an arbitrary pair of nodes in a network of N nodes is approximated by $logN$ [1]. We assume all the routers also serve as resolution servers and hence N is equal to S. LOMO routes the packets over the path decided by the current IP routing, and hence its lookup delay is $2{*}logS$. In the DNS, a solicitor sends/receives a query/response packets with randomly distributed servers for a few rounds (we use 3 round trip times for plotting the DNS lookup delays). Chord performs worst since mapping query packets go back and forth for $logS$ overlay hops, and each overlay hop takes $logS$ (physical) hop count as Chord nodes are randomly located. Hence, the lookup delay of Chord is $(logS)^2$.

According to a measurement study [9], the popularity of domain names follows Zipf-like distribution. In the Zipf distribution, the probability of the x-th popular object is modeled by a function c/x^{α}, where c and α are the normalizing constant and the popularity index, respectively. As α grows, the disparity of popularity of objects (or domain names) becomes severer. Figure 2c shows the worst-case workload on a server with D domain names being queried as α of the Zipf distribution varies. [9] studied the DNS traffic and found that as the level in the DNS hierarchy goes up, α increases. That is, the popularity distribution of second-level domain names is skewer than that of fully qualified domain names. Considering that the measured DNS traffic in [9] does not count the cache hits in the local DNS server, we conjecture that α (that reflects the orig-

(a) Number of resolution Entries (b) Lookup delay (c) Workload

Figure 2: Result Graphs

inal popularity without the caching effect) may increase more sharply since the locators of mobile hosts will be valid much shorter in the DNS cache. On the other hand, the popularity distribution of domain names in LOMO and Chord will be closer to the uniform distribution (i.e. smaller α) since they have flat structures.

4. DISCUSSIONS

The LOMO scheme in the previous section partitions the ID space equally among ASs regardless of the AS sizes (i.e. number of routers). We need to address the disparity in AS sizes. We suggest that the ID space assigned to each AS should be proportional to the IP address space assigned to the AS. The rationale behind this assignment is that the more IP address space an AS has, the more routers is the AS likely to operate.

The LOMO can be deployed incrementally by adding a flag to indicate whether an AS supports LOMO or not in a BGP UPDATE message. When, initially, only a handful of ASs may start the resolution service with the described BGP extensions for LOMO, the AS mapping table will include them only. Likewise, an intra-domain routing message generated by each router has a flag to indicate whether the router participates in LOMO or not.

The LOMO can also support multi-homing easily by letting each resolution entry of a mobile host maintain multiple locators. In this case, some extra attributes (e.g. time of registration, stability or available bandwidth of each interface) will be useful for a solicitor to decide which locator should be contacted preferably.

5. CONCLUSIONS

Mapping from names of mobile hosts to their locators will be crucial for future Internet. Due to problems of the prior mobility management solutions at the network layer, we follow the approach of host-based mobility management proposals. For the name-to-locator resolution, the prior host-based mobility solutions simply rely on the DNS, which may not scale. We propose LOMO, which is a new resolution architecture that uses the IP routing structures. By extending IP routing protocols slightly, LOMO can handle the resolution traffic in a scalable and load-balanced fashion. To facilitate any indirection in the future Internet, we will study how LOMO can serve as a general rendezvous system that provides mapping between a name/identifier and its corresponding information.

6. ACKNOWLEDGMENTS

This work was supported by NAP of Korea Research Council of Fundamental Science & Technology. This research was supported by the MKE (The Ministry of Knowledge Economy), Korea, under the ITRC (Information Technology Research Center) support program supervised by the NIPA (National IT Industry Promotion Agency) (NIPA-2011-(C1090-1111-0004)). The ICT at Seoul National University provides research facilities for this study.

7. REFERENCES

[1] Albert, R., and Barabási, A.-L. Statistical mechanics of complex networks. *Rev. Mod. Phys. 74*, 1 (Jan 2002), 47–97.
[2] Atkinson, R., Bhatti, S., and Hailes, S. Mobility through naming: Impact on dns. In *Mobiarch '08* (2008), ACM.
[3] CAIDA. http://as-rank.caida.org/, 2011.
[4] Cisco. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf, 2010.
[5] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and Patil, B. Proxy mobile ipv6. RFC 5213, 2008.
[6] Henderson, T. Host mobility for ip networks: A comparison. *IEEE Network 17*, 6 (2003), 18–26.
[7] Huston, G. http://www.potaroo.net/tools/asn16/, 2010.
[8] Johnson, D., Perkins, C., and Arkko, J. Mobility support in ipv6. RFC 3775, 2004.
[9] Jung, J., Sit, E., Balakrishnan, H., and Morris, R. Dns performance and the effectiveness of caching. *ACM/IEEE Trans. on Networking. 10*, 5 (Oct 2002), 589–603.
[10] Moskowitz, R., and Nikander, P. Host Identity Protocol (HIP) Architecture. RFC 4423, 2006.
[11] Stoica, I., Morris, R., Karger, D., Kaashoek, M. F., and Balakrishnan, H. Chord: A scalable peer-to-peer lookup service for internet applications. In *ACM SIGCOMM* (2001), ACM, pp. 149–160.
[12] VeriSign. http://www.verisigninc.com/assets/domain-name-report-feb-2011.pdf, 2011.