

IP Mobility Overview

The increasing use of virtualization in the data center has enabled an unprecedented degree of flexibility in managing servers and workloads. One important aspect of this newfound flexibility is mobility. As workloads are hosted on virtual servers, they are decoupled from the physical infrastructure and become mobile by definition.

As end-points become detached from the physical infrastructure and are mobile, the routing infrastructure is challenged to evolve from a topology centric addressing model to a more flexible architecture. This new architecture is capable of allowing IP addresses to freely and efficiently move across the infrastructure.

There are several ways of adding mobility to the IP infrastructure, and each of them addresses the problem with different degrees of effectiveness. LISP Host Mobility is poised to provide a solution for workload mobility with optimal effectiveness. This document describes the LISP Host Mobility solution, contrasts it with other IP mobility options, and provides specific guidance for deploying and configuring the LISP Host mobility solution.

IP Mobility Requirements

The requirements for an IP mobility solution can be generalized to a few key aspects. To make a fair comparison of existing solutions and clearly understand the added benefit of the LISP Host Mobility solution, we will quickly touch on the different functional aspects that must be addressed in an IP mobility solution.

- **Redirection** The ultimate goal of IP mobility is to steer traffic to the valid location of the end-point. This aspect is generally addressed by providing some sort of re-direction mechanism to enhance the traffic steering already provided by basic routing. Redirection can be achieved by replacing the destination address with a surrogate address that is representative of the new location of the end-point. Different techniques will allow the redirection of traffic either by replacing the destination's address altogether or by leveraging a level of indirection in the addressing such as that achieved with tunnels and encapsulations. The different approaches impact applications to different degrees. The ultimate goal of IP mobility is to provide a solution that is totally transparent to the applications and allows for the preservation of established sessions, as end-points move around the IP infrastructure.
- **Scalability** Most techniques create a significant amount of granular state to re-direct traffic effectively. The state is necessary to correlate destination IP addresses to specific locations, either by means of mapping or translation. This additional state must be handled in a very efficient manner to attain a solution that can support a deployable scale at a reasonable cost in terms of memory and processing.

- **Optimized Routing** As end-points move around, it is key that traffic is routed to these end-points following the best possible path. Since mobility is based largely on re-direction of traffic, the ability to provide an optimal path is largely a function of the location of the re-directing element. Depending on the architecture, the solution may generate sub-optimal traffic patterns often referred to as traffic triangulation or hair-pinning in an attempt to describe the unnecessary detour traffic needs to take when the destination is mobile. A good mobility solution is one that can provide optimized paths regardless of the location of the end-point.
- **Client Independent Solution** It is important that the mobility solution does not depend on agents installed on the mobile end-points or on the clients communicating with these end-points. A network based solution is highly desirable and is key to the effective deployment of a mobility solution given the precedent of the large installed base of end-points that cannot be changed or managed at will to install client software.
- **Address Family Agnostic Solution** The solution provided must work independently of IPv4 or IPv6 end-points and networks. Since mobility relies on the manipulation of the mapping of identity to location, address families with lengthier addresses tend to provide alternatives not available with smaller address spaces. These address dependent solutions have limited application as they usually call for an end to end deployment of IPv6. To cover the broad installed base of IPv4 networking and end-points, the ideal solution should work for IPv4 or IPv6 independently.

Existing IP Mobility Solutions

The following IP Mobility technology solutions are available and described below:

- [Route Health Injection \(RHI\) and Host Routing](#), page 1-2
- [Mobile IPv4](#), page 1-2
- [Mobile IPv6](#), page 1-3
- [DNS Based Redirection: Global Site Selector \(GSS\)](#), page 1-3

Route Health Injection (RHI) and Host Routing

One simple way to redirect traffic to a new location when a server (or group of servers) moves is to inject a more specific route to the moved end-point(s) into the routing protocol when the moves are detected. In the extreme case, this means injecting a host route from the "landing" location every time a host moves. Load balancers with the Route Health Injection (RHI) functionality implemented can provide an automated mechanism to detect server moves and inject the necessary host routes when the servers move.

This approach, although simple, pollutes the routing tables considerably and causes large amount of churn in the routing protocol. Forcing churning of the routing protocol is a risky proposition as it could lead to instabilities and overall loss of connectivity, together with adding delays to roaming handoffs.

Mobile IPv4

Mobile IP is defined for IPv4 in IETF RFC 3344. Basically mobile IPv4 provides a mechanism to redirect traffic to a mobile node whenever this node moves from its "Home Network" to a "Foreign Network." Every host will have a "Home Address" within a "Home Network" which is front-ended by a router that acts as a "Home Agent" and that advertises the "Home Network" into the routing protocol. Traffic destined to the "Home Address" will always be routed to the "Home Agent." If the mobile node

is in its "Home Network" traffic will be forwarded directly in the data plane to the host as per regular routing. If the host has moved to a "Foreign Network", traffic will be IP tunneled by the "Home Agent" to a "Care-of- Address" which is the address of the gateway router for the "Foreign Network."

With Mobile IPv4 there is always a triangular traffic pattern. Also, Mobile IPv4 does not offer a solution for multicast. Since the mobile node is usually sourcing traffic, if the Foreign Agent is not directly connected, there is the need for host route injection at the foreign site to get RPF to work. In addition, multicast traffic from the mobile node has to always hairpin through the home agent since the distribution tree is built and rooted at the "Home Agent."

Mobile IPv6

IETF RFC 3775 defines mobility support in IPv6. IPv6 takes a step beyond IPv4 mobility and provides optimal data paths between server and client. The process in IPv6 is similar to that of IPv4 with a few additions.

Rather than having the Home Agent always redirect the traffic to the Care-of-Address (CoA) for the server that has moved, the Home Agent is taken out of the data path by distributing the CoA to Home Address Binding information to the client itself. Once the client has the CoA information for a particular server, it can send traffic directly to the CoA rather than triangulating it through the Home Address. This provides a direct path from client to server.

Although Mobile IPv6 provides direct path routing for mobile nodes, it is limited to IPv6 enabled end-points, it requires that the entire data path be IPv6 enabled, and it also requires that the end-points have IPv6 mobility agents installed on them.

DNS Based Redirection: Global Site Selector (GSS)

It may be possible to direct traffic to a moving server by updating the DNS entries for the moving server as the server moves locations. This scheme assumes that every time a server moves it is assigned a new IP address within the server's "landing" subnet. When the server moves, its DNS entry is updated to reflect the new IP address. Any new connections to the server will use the new IP address that is learnt via DNS resolution. Thus traffic is redirected by updating the mapping of the DNS name (identity) to the new IP address (location).

The new IP address assigned after the move may be assigned directly to the server or may be a new Virtual IP (VIP) on a load balancer front-ending the server at the new location. When using load balancers at each location, the load balancers can be leveraged to determine the location of a host by checking the servers' health with probes. When a change of location is detected, the integration of workflow in vCenter (VMware) updates the Global Site Selector (GSS) of the new VIP for the server and the GSS will in turn proceed to update the DNS system with the new VIP to server-name mapping. Established connections will continue to try to reach the original VIP, it is up to the load balancers to be able to re-direct those connections to the new host location and create a hair-pinned traffic pattern for the previously established connections. New connections will be directed to the new VIP (provided the DNS cache has been updated on the client) and will follow a direct path to this new VIP. Eventually all old connections are completed and there are no hair-pinned flows.

The main caveats with this approach include:

- Rate of refresh for the DNS cache may impact either the convergence time for the move or the scalability of the DNS system if the rate is too high.
- Works only for name-based connections while many applications are moving to an IP based connection model.

- Previously established connections are hair-pinned. This implies that there is a period of time where there are active connections to the old address and some new connections to the new address in the second data center. During this state the network administrator may not be able to ascertain that these two addresses are the same system (from the point of view of the application).

**Note**

For more information on the DNS based redirection functionality leveraging Cisco Load Balancers and Global Site Selector please refer to the Cisco Virtualized Workload Mobility design guide:
http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DCI/4.0/EMC/EMC.pdf
