

Virtual Multi-homing: On the Feasibility of Combining Overlay Routing with BGP Routing

Zhi Li, Prasant Mohapatra, and Chen-Nee Chuah

University of California, Davis, CA 95616, USA

{lizhi, prasant}@cs.ucdavis.edu, chuah@ece.ucdavis.edu

Abstract. This paper proposes a framework called *Virtual Multi-Homing (VMH)* to achieve loose source-based path selection and improve inter-domain path diversity. VMH is based on the concept of *Virtual Peering* and *Multi-Homing Overlay* to set up flexible inter-domain relationships. By interacting with BGP whenever possible, VMH can achieve scalable inter-domain route discovery without introducing duplicate work. In addition, VMH is a complementary approach to the existing Internet routing infrastructure and can be incrementally deployed.

1 Introduction

Even though there exists high degree of redundancies in the current Internet topology, Border Gateway Protocol (BGP-4) cannot fully utilize it to provide desirable inter-domain routing services [2]. In some cases, such as content distribution networks (CDNs) and service composition networks, the source Autonomous System (AS) may want to traverse or avoid specific ASes due to performance or security concerns. Unfortunately, the end users and low-tier ASes usually cannot select their preferred inter-domain paths and control how the packets are routed.

To improve path diversity and fault tolerance, some ASes use the concept of multi-homing by subscribing to more than one upstream service providers. However, multi-homing can only increase source ASes' flexibility in selecting the immediate next-hop ASes. In addition, the number of providers an AS can subscribe is usually limited due to economic reasons. Alternative routing architectures have been proposed to improve inter-domain routing performance and provide additional flexibility that BGP lacks, e.g., NIRA [6], RON [2], and RPC [3]. However, these approaches either require the change of the whole Internet routing architecture [6] or are not intended for large-scale deployment [2]. Overlay networks have been proposed as an effective way to improve inter-domain routing performance. However, existing overlays can only reach the destinations within the same overlay networks.

In this paper, we propose a new framework called Virtual Multi-Homing (VMH) to investigate the feasibility of combining the strength of overlay and BGP routing to improve inter-domain routing performance. VMH can be used by overlay service providers, such as Akamai [1] or PlanetLab, to provide wide area

inter-domain routing service to customers. Compared to existing approaches, VMH has the following main advantages. First, VMH enables flexible source-based routing by allowing a source AS to send packets through a series of selected overlay nodes within a multi-homing overlay network (MON). This loose source-based forwarding feature can achieve flexible inter-domain path selection and enhanced QoS support. Second, VMH uses overlay routing technique to enhance inter-domain path diversity for fault tolerance or load balancing. It can also be used to provide proactive re-routing during transient inter-AS path failures without relying on the global path re-convergence. Third, unlike most overlay networks that use a totally separate control plane to discover and exchange path reachability information, VMH interacts with BGP to learn about inter-AS connectivity and routing policies. By associating a subset of the BGP prefixes with a chosen overlay node of MON (e.g., based on proximity to an AS), VMH extends its overlay routing services to a larger range of destinations.

We will introduce the framework of VMH and the preliminary simulation results in Section 2 and discuss the related issues and future work in Section 3.

2 Virtual Multi-Homing (VMH)

2.1 Proposed VMH Framework

In VMH, each AS can have one or more *Multi-Homing Servers (MHS)*. An MHS can either co-exist with a border router or locate at a separate host. An MHS sets up one or more I-BGP sessions with local BGP routers to receive BGP updates and determine its inter-domain paths to other destinations. However, an MHS is a passive peer, i.e., it will not generate route updates to its BGP peers.

The MHSes from different ASes cooperate with each other to form a *Multi-Homing Overlay Network (MON)*. The connections between MHSes are based on the *Overlay Transit* relationships among VMHes. Similar to current inter-AS relationship, *Overlay Transit* determines whether a MHS can forward packets to its neighboring MHSes within the MON. That is, if there is an *Overlay Transit* relationship between two ASes, there is a corresponding overlay link in the MON overlay topology.

A link state based overlay routing protocol runs among the MHSes. This provides resilient overlay routings path connecting each pair of source and destination MHSes. Similar to RON[2], each MHS continuously probes its neighboring MHSes and sends the latest overlay link performance to every other MHS. If there is an IP-layer path failure or service degradation, MON can quickly provide an alternative overlay path.

Virtual Multi-Homing (VMH) runs on top of MON. A new inter-domain transit service called *Virtual Peering* can be set up between two remote ASes. There is a virtual BGP peering session (using an overlay path via MON) between these two ASes' MHSes. One end (an MHS) is a *Virtual Peering Provider (VPP)* while the other end (an MHS) is a *Virtual Peering Customer (VPC)*. Different from *Overlay Transit* relationship, a VPP can send its customer's traffic to any destinations. A VPC can receive BGP updates from its VPPs in addition to the

messages for its local BGP router via I-BGP session. However, a VPC will not send any BGP update messages to its VPPs. This restriction ensures that VMH will not introduce any extra routing states into the BGP routing system and affect its performance. Based on MON, an AS can subscribe several VPPs in addition to its directly-connected provider(s). If necessary, it can also subscribe new VPPs on-demand. To send a packet to its destination, a source AS can either (a) send a packet via its direct physical providers, or, (b) send the packet using an overlay path to one of its VPPs, which will then forward the packet to its destination. We can see that the traditional multi-homing service can be deemed as a special case of VMH. It can be observed that VMH can help a source AS to explore inter-domain path diversity, which will potentially improve the inter-domain paths service quality.

Once a VMH user's packet arrives at a MHS, the MHS first picks a destination MHS from its set of VPPs based on its BGP path information. The selection of a destination MHS is based on the distances between VPPs and the packet destination as well as the path disjointness between the default BGP path and the paths via VPPs. Next, the source MHS determines the Loose Forwarding Path (LFP) from itself to the destination MHS using one of the two mechanisms: 1) By default, the source MHS can find an overlay path (list of MHSes) with best service quality (least loss rate or shortest delay); Or 2) the source MHS can find an alternate overlay path through MON based on the constraints specified by packet sender or local routing policy.

After finding the LFP to a destination MHS, the source MHS can encapsulate the original data packets with an additional header which includes the list of MHSes (LFP). When a MHS receives a loose forwarding packet from one of its neighboring MHSes, if the current MHS is the destination MHS on the LFP and the packet source is one of its VPCs, it just removes the LFP header and sends the packet to its destination via the normal IP path. Otherwise, if the packet comes from one of its *Overlay Transit* customers, it locates the LFP from the packet header part and sends out the packet to the next hop MHS. Step-by-step, the packet will be transmitted to the destination MHS following the LFP. It should be noted that, when the data pass through the overlay links, they will be transmitted via IP-Tunneling through corresponding IP-layer paths. However, the destination MHS will remove the LFP packet headers and directly send the data to its destination.

In summary, under VMH framework, the end-to-end routing via VMH is a combination of IP and overlay routing. This will result in less redundant work and is inherently "topology-aware". As shown in [2], the inter-domain forwarding performance can be greatly improved via one or two intermediate forwarding points, we expect that most of LFPs will pass through no more than 3 MHSes. As a result, the extra LFP header will not incur too much overhead.

2.2 Preliminary Performance Evaluation

Our simulation is based on a real Internet inter-domain topology provided by [5] with Tier 5 ASes pruned out, which left a topology with 2473 ASes. For each

simulation, we randomly construct a MON whose size varies from 25 to 200 and choose a VPC from those Tier 4 ASes with only one physical provider. We also vary the number of the customer's VPPs from one to eight. For each setup, we use the heuristic method proposed in [4] to choose virtual peering providers. We evaluate VMH's performance in terms of exploring Inter-domain path diversity by the following two metrics [4]:

- End-to-End Path Availability (*EEPA*): it is defined as the ratio of existing a valid inter-domain path from a VPC to a destination given a link/AS availability ratio.
- Resilient Serviced Destination Ratio (*RSDR*): it is the ratio of destinations to which a VPC's paths can benefit from VMH.




Fig. 1. Path Availability (MON size 100)




Fig. 2. Average Serviced Dest Paths

Figure 1 shows the average EEPA for a customer AS passing through a *MON* of 100 nodes. For a customer AS, without any VPP, the customer only has one BGP path provided by its service provider. We show that VMH can increase EEPA by providing the customer more flexibility to select loose forwarding path via *MON*. The EEPA improves with the increasing number of VPPs but saturates when the number of VPPs is more than 4. This is because the VMH's performance is also restricted by the underlying physical topology and the composition of *MON*, which together determine the maximum benefit VMH can provide to the customer ASes.

Similarly, VMH cannot provide a customer with resilient source-based routing paths (alternate AS-level paths) to all the destinations restricted by the underlying physical topology. Figure 2 shows the average RSDR for a VMH customer. From the figure, we can observe that both *MON* size and the number of VPPs determine the number of serviced destination paths. This means that careful selection of VPPs is important for a VMH customer.

3 Discussion and Future Work

VMH can effectively utilize the inter-domain path diversity and provide flexible routing service to a large user population. The proposed method utilizes BGP

routing states whenever possible, hence reducing the possibility of introducing duplicate work at overlay layer. However, one may raise the following concerns of VMH:

Security Issue from Source-based Routing: In VMH, a *Virtual Peering* session or *Overlay Transit* is set up between two trust-worthy MHSes. An MHS will only deliver packets for its *Virtual Peering* or *Overlay Transit* customers. This can facilitate source tracing and prevent DoS attack without introducing extra security concerns.

Multi-Hop E-BGP Session Stability Analysis over MON: In our technical report [4], we have analyzed the real multi-hop EBGP session stability and shown that overlay can effectively reduce the possibility of session resets. It is shown that the virtual peering BGP session is feasible via MON.

Impact on BGP Routing Performance: We can see that VMH does not introduce any BGP routing messages to the BGP routing system. However, similar to other overlay networks, VMH will certainly affect the intra-domain and inter-domain traffic distribution. How ISPs or BGP can effectively deal with the dynamic traffic shifting is still an open issue.

In future, we plan to deploy a test-bed on top of Planet-Lab to study VMH's performance and related management issues.

References

1. Akamai Corporation. <http://www.akamai.com>.
2. D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris. Resilient overlay network. In *18th ACM SOSP*, Oct. 2001.
3. N. Feamster, H. Balakrishnan, J. Rexford, A. Shaikh, and J. Merwe. The case for separating routing from routers. In *ACM SIGCOMM FDNA workshop '04*, Sep. 2004.
4. Z. Li, P. Mohapatra, and C.-N. Chuah. Virtual Multi-Homing: on the Feasibility of Combining Overlay Routing with BGP Routing, UC Davis ECS Technical Report, March 2005.
5. L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the Internet hierarchy from multiple vantage points. *Proc. IEEE INFOCOM*, 2002.
6. X. Yang. NIRA: A New Internet Routing Architecture. In *ACM Sigcomm FDNA workshop*, 2003.