VSITE: a scalable and secure architecture for seamless L2 enterprise extension in the cloud

Li Erran Li, Thomas Woo
{erranlli,woo}@research.bell-labs.com
Bell Labs, Alcatel-Lucent

ABSTRACT

This paper presents an end-to-end architecture, called VSITE, for seamless integration of cloud resources into an enterprise's intranet at layer 2. VSITE allows a cloud provider to carve out its resources to serve multiple enterprises simultaneously while maintaining isolation and security. Resources (allocated to an enterprise) in the cloud provider appears "internal" to the enterprise. VSITE achieves this abstraction through the use of VPN technologies, the assignment of different VLANs to different enterprises, and the encoding of enterprise IDs in MAC addresses. Unlike traditional layer 2 VPN technology such as VPLS, VSITE suppresses layer 2 MAC learning related broadcast traffic. VSITE makes use of location IP (represents location area) for scalable migration support. The MAC or IP address of a VM is not visible in data center core. VSITE has built in security mechanisms to prevent enterprises from attacking one another. Thus, VSITE is scalable, secure and efficient, and it facilitates common data center operation such as VM migration. Because VSITE extends enterprise network at layer 2, this offers transparency to most existing applications and presents an easy migration path for an enterprise to leverage cloud computing resources.

1. INTRODUCTION

A particularly important form of cloud computing is called Infrastructure as a Service (IaaS), which refers to the provision of raw resources such as compute servers and storage from a cloud. Amazon's Elastic Cloud Computing (EC2) solution is a well-known example of IaaS cloud computing. In such a solution, a user is allocated compute/storage resources in the cloud where she can use for launching applications of her choice.

For most enterprises, IaaS represents an intriguing option for scaling their computing needs. Specifically, an enterprise can choose to build its internal infrastructure to satisfy only the average demand, and augment that with resources from an IaaS cloud on an on-demand basis to address demand burst above the average. To enable live VM migration and minimize application configuration, it is essential that the network extension can be reached in layer 2.

A public data center needs to accommodate a large number of enterprises (tenants). These tenants need to be logically isolated. Assign each tenant a VLAN throughout the data center is not a scalable solution. There are only 4096 VLANs. It is possible to use VLAN stacking where the outer VLAN ID will be shared by many enterprises, and the inner VLAN ID identifies each enterprise. However, in this case, each enterprise's L2 broadcast traffic will be flooded throughout the outer VLAN and can potentially be visible throughout the data center.

It is tempting to reuse VPLS technology to interconnect an enterprise with its cloud provider. However, VPLS floods L2 broadcasts from the public data center to the enterprise, vice versa. Each VM can potentially appear in the L2 forwarding table of each switch in the enterprise.

In this paper, we provide a systematic treatment of issues related to supporting seamless enterprise cloud computing and present a scalable architecture, called VSITE. VSITE uses a protocol similar to recent CISCO OTV [1] for scalable connection between enterprise networks and cloud resources. The MAC information are exchanged between the cloud edge switch and the enterprise customer edge switch. This is much more scalable than VPLS which uses MAC learning.

VSITE uses VLANs to logically separate enterprises. Unlike Cloudnet [9], VSITE does not assign global VLAN IDs (visible throughout a data center) to enterprises. Instead, VSITE assigns local VLAN IDs that are local to a data center edge location. To ensure packets from one enterprise never reach another enterprise, we encode enterprise IDs in MAC addresses. Specifically, the hypervisor of a VM will ensure that packets can only be sent to or received from MAC addresses of the same enterprise.

VSITE makes minimal assumption of data center internals. VSITE uses Ethernet over IP (e.g. Ethernet over GRE or EnterIP protocol) encapsulation to traverse the **data center core**. VSITE only assumes Ethernet at **data center edge**. L3 switches connect data center edge with data center core[1].

VSITE introduces the concept of location IP for scalable VM migration. Specifically, each VM has a location IP (`locIP`). `locIP` is the IP of any L3 switch that interconnects data center edge and core. The hypervisor of a VM looks up a directory server for the `locIP` of the destination VM if it is not known.

VSITE makes the following contributions:

- VSITE makes public cloud data center resources appear "internal" to an enterprise. VSITE provides a

[1] VSITE design also works well if both core and edge are Ethernet; we will discuss the changes required later

set of networking mechanisms that solves the address space, security and scalability issues.

- VSITE enables dynamic and scalable service offering. Dynamic creation, removal or migration only involves updating the directory service and the relevant elements such as logical customer premise edge router (CE), and the configuration of the hypervisor of a server. Migration within the data center edge is seamless; MAC learning will enable packet forwarding to the new location.
- VSITE does not require changes to switches and routers. All the changes are relegated to the CEs and the hypervisors. The overhead is the directory query when information are not cached.
- VSITE prevents L2 attacks such as ARP attacks, MAC flooding attacks performed by one enterprise from adversely impacting other enterprises.

The rest of the paper is organized as follows. In section 2, we present the architecture and its design principles. We illustrate the details of VSITE control and data place in Section 3 and 4. We then discuss issues in Section 5, 6, related work in Section 7 and conclude our paper in Section 8.

2. THE VSITE ARCHITECTURE

The problem of creating a cloud extension to an enterprise is a multifaceted one. A complete solution will touch on resource allocation (e.g., where and how the cloud resources are allocated), networking, monitoring, migration, reporting, etc. We focus mostly on networking and security in this paper as it lays the foundation for most other components.

We begin by defining a *virtual stub* network or *vstub* to be the logical network inside the cloud infrastructure formed by the collection of resources allocated to a specific enterprise. A virtual stub network is logically isolated from resources outside the stub in that resources within a stub can communicate freely while no traffic can cross the stub boundary. Essentially, a stub network is like a virtual private subcloud inside the cloud infrastructure.

We break up the problem of seamless enterprise cloud extension into three components: **virtual stub creation**, and **virtual stub attachment** and **virtual stub isolation and scalability**. We present the overview of VSITE architecture, design considerations and architecture entities.

2.1 The End-to-End Picture

Figure 1 shows the different components that comprise the end-to-end problem of extending an enterprise into the cloud.

Starting from the left are the physical enterprise sites. Two different enterprises *A* and *B* are shown here, with enterprise *A* having two existing sites *A*1 and *A*2.

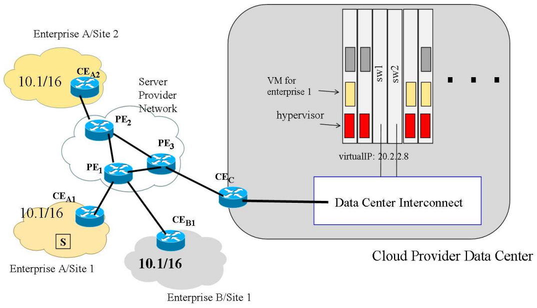

Figure 1: End-to-end architecture

Virtual Stub Attachment: Connectivity between different enterprise sites including vstub are provided over the service provider network. This is typically provided using some form of VPN technologies. Using standard terminology, CE (customer edge) routers in enterprise sites interface to PE (provide edge) routers in the service provider network to create VPN connectivity. These VPN solutions can include IPsec, Layer 3 MPLS VPN, and Layer 2 VPN such as VPLS, OTV [1]. For example, in a site extension using L3 MPLS VPN case, the stub site will be given an IP subnet prefix that is within the space of the enterprise, all resources within the stub site will be allocated an IP address within that subnet. The subnet prefix is advertised via BGP within the MPLS VPN to all other sites within the enterprise, which then establishes reachability to the stub site. Communication between internal and stub site resources proceeds via standard MPLS encapsulation as defined for L3 MPLS VPN. Some technologies such as IPsec VPN and Cisco OTV [1] do not require special support from Internet service providers. That is, PE routers do not have any specific VPN state.

A virtual stub attachment connects a virtual stub to the enterprise VPN which extends over the service provider network and reach all enterprise sites.

Virtual Stub Creation: A cloud provider is shown on the right side of the picture. In fact, it shows a data center of the cloud provider. The cloud provider can be the service provider or a third party provider, and there can be multiple physical data centers. The cloud provider data center interfaces to the service provider network just like an enterprise site. The data center hosts the physical resources such as the compute and storage servers. Typically, these resources are hosted in chassis-based servers (i.e., blade servers) and are virtualized. A blade server has multiple slots each of which can hold a physical server blade. These blades are redundantly connected to a pair of embedded switch within the chassis. Each server blade runs a *hypervisor* which in turn hosts *virtual machines* (VMs) that can be allocated to an enterprise. The different resources inside a data center is connected by a data center network.

In simple terms, a virtual stub network begins at a data center CE, and extends over the data center network to reach all the VMs allocated to an enterprise.

Virtual Stub Isolation and Scalability: Virtual stubs between different enterprise sites must be isolated be default. A cloud provider must be able to support a large number of vstubs. One can make use of VLANs to isolate different enterprises. However, there are only 4096 IDs. VLAN stacking can increase the number of VLANs to 4096*4096. However, broadcast traffic of VLANs can seriously stress the data center networks. All broadcast with the same outer VLAN ID will reach all ports configured with that VLAN ID.

As VMs are created or terminated dynamically in a much shorter time scale, traditional MAC learning can seriously impact the underlying network: (1) broadcast traffic will reach all switches with the same VLAN configured; (2) end host MAC entries will populate switch MAC forwarding tables.

2.2 Design Considerations

In the following, we look at the key technical requirements and motivate the specific VSITE design feature that address it.

VM configuration transparency and mobility: Many applications may require layer 2 connectivity to different application components, e.g. L2 reachability among front end or backend severs in three-tiered applications. L2 connectivity also makes it easy to support VM mobility. Thus, we choose L2 vstub attachment. Among L2 technologies, there are VPLS, pseudo-wire and OTV. Pseudo-wire is a point-to-point solution. We would like to interconnect multiple sites of a given enterprise. Thus, we choose mutipoint-to-multipoint solutions. We make use of OTV-like protocols (OTV is proprietary). The key advantage of OTV-like protocols is that it exchanges MAC addresses among sites using a control plane protocol. This eliminates the cross-site MAC learning related flooding which increases the scalability of vstub attachment.

To further support VM configuration transparency, a VM's hypervisor will intercept the VM's DHCP and ARP messages. a VM will be configured the same way as if it is in the enterprise.

Data center scalability: Each VM is configured with its enterprise IP `entIP`. Rapid shrinking or growing of virtual machines (VM) as well as rapid VM migration calls for the separation of names from locations. For each VM, we associate it with a location IP (`locIP`). One option for `locIP` is to use the switch IP to which the VM logically connects. If the server hosting a VM is multi-homed to more than one rack switch, then `locIP` is the virtual IP of the switch master (switches are configured in master or slave by virtual router redundancy protocol (VRRP)). Note that, we could have used server virtual NIC IP as `locIP` of an VM. Virtual NIC using NIC teaming technique can provide fault tolerance again one NIC failure. However, VM mobility is handed better when the `locIP` is the switch IP. If a VM migrates to a different server behind the same switch, then there is no need for routing updates as the `locIP` does not change.

A specific VM's MAC address is only visible within its location area (the data center edge). We only assume the data center edge is Ethernet. The limited VM MAC visibility increases the scalability of data center core. At the same time, it supports VM mobility much better as the `locIP` stays the same within the location area.

Data center core transparency: There are two main reasons that we want VSITE to be data center core transparent (i.e. data center core is agnostic of vstub). The first is scalability. Limiting VM's MAC within the edge increases data center core scalability. The second is to accommodate different data center designs. Data center core is still evolving, and there are many competing technologies, some are composed of purely L2 nodes, some are L3 nodes, and some even proposes the use of optical switches.

Traffic separation and security: A vstub can be part of an enterprise VLAN logically. To support this, we organize a vstub in a VLAN in the data center. This VLAN does not reach into data center core. In order to support a large number of vstubs, VLAN ID has only local significance. Thus, the VLAN ID of a vstub can be different from the VLAN ID of its enterprise portion. The data center CE does the translation. We will discuss the details later.

Since VLAN IDs are not assigned to enterprises statically, there may be transient situations where traffic of VMs of one enterprise gets sent to VMs of a different enterprise (due to VLAN ID reuse). To prevent this, we choose to encode enterprise ID in the MAC address. There are many encoding schemes. The simplest one is to partition the 48 bit MAC address into a x bit portion which uniquely identifies the enterprise, and a $48-x$ bit portion that identifies a VM that is local to ToR switch. Other encoding schemes are possible. All we need is that we should be able to easily figure out the enterprise a given MAC address belongs to. The hypervisor of a VM will check the validity of MAC addresses. That is, if a VM sends a packet to a VM of a different enterprise, the hypervisor will drop the packet. Similarly, if the hypervisor receives a packet where the source and destination enterprise ID are different, the packet will be dropped.

2.3 VSITE Architecture Entities

A number of entities work together to create the VSITE architecture. We describe these entities and their roles here.

Cloud Manager: The cloud manager takes enterprise customer requests, and creates vstubs. This includes instantiating the VMs, configuring appropriate network elements such as cloud data center CE.

Directory Server: On an enterprise basis, there is a mapping from entIP to locIP and MAC, possible other informations such as VLAN ID. Such mappings are collected in a directory server maintained by the cloud provider; and like a DNS server, queries can be made to retrieve the current mapping.

Cloud Data Center CE (CE_c): Typically, a site CE serves as a gateway for a single enterprise site. In the cloud data center case, its CE must support virtualization as it provides gateway functions for multiple tenants. Cloud data center CE communicates with the directory server to retrieve the up-to-date mapping. Cloud data center CE implements OTV-like protocol which exchanges MAC reachability information with enterprise CE and directory server.

Hypervisor: We assume a hypervisor understands (i.e., is configured with) the enterprise information for each VM it host. Specifically, it knows the enterprise membership (i.e., enterprise ID) for each VM, and their corresponding security parameters. A hypervisor is also responsible for intercepting certain traffic (e.g., L2 broadcast), ARP request and DHCP request. The hypervisor communicates with VSITE data plane to assign an appropriate MAC address for a VM. The hypervisor also relays DHCP messages to DHCP servers.

3. VSITE CONTROL PLANE

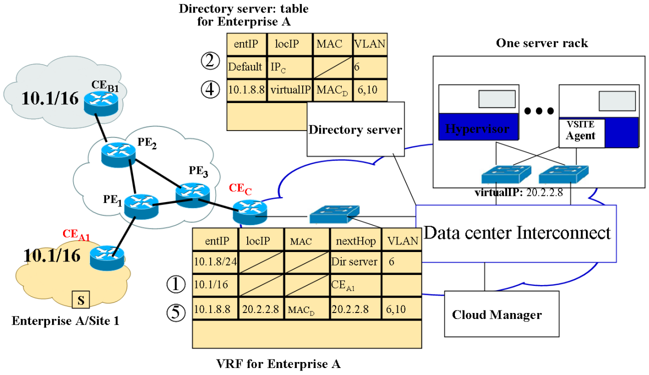

	entIP	locIP	MAC	VLAN
②	Default	IP_C		6
④	10.1.8.8	virtualIP	MAC_D	6,10

	entIP	locIP	MAC	nextHop	VLAN
	10.1.8/24			Dir server	6
①	10.1/16			CE_{A1}	
⑤	10.1.8.8	20.2.2.8	MAC_D	20.2.2.8	6,10

Figure 2: Control Plane of VSITE

VSITE control plane is shown in Figure 2. An enterprise can request cloud resources via a web portal or programmatic API. These interface to a cloud manager that kicks off actions within the cloud to instantiate/remove a vstub and/or add/subtract resources from a vstub. Such a manager should have interface for an enterprise to configure address plan for its vstub (including internal DHCP server addresses), VM images, security keys, etc.

The intelligence of deciding which data center and servers should specific resources be allocated from is beyond the scope of this paper. We focus on the networking actions that tie those resources into the enterprise once they are allocated.

We describe the typical control plane scenarios below.

vstub Instantiation: As shown in Figure 2, for enterprise A to dynamically instantiate a subnet (say 10.1.8/24) in the cloud provider, a logical CE (virtual router) will be created in CE_c if it does not exist. Once the logical CE is created, a routing entry for the subnet will be configured in the enterprise's VRF table of CE_c. The provider-facing interface of enterprise A's CEs (such as CE_{A1}) will be properly configured. A default entry mapping all enterprise address to the logical CE's IP_C will be installed at the directory server. Packets originating from a vstub VM to an internal enterprise address now will be forwarded to the logical CE, which will then forward it to the enterprise.

VM Instantiation: A VM is identified by its entIP, MAC, and locIP. All these information are obtained during the boot process. Upon booting, the VSITE hypervisor intercept the DHCP request and relay it to a previously configured per-enterprise DHCP server. If the server sits inside the enterprise, the packet will be properly tunneled to CE_c first using the regular tunneling procedure (will become clear in the next section). The locIP is the address of the switch to which the host hypervisor homes. The entIP, locIP and MAC mapping is then installed in the directory server (step 4). The directory server may optionally install a routing entry in routing table of the logical CE (step 5).

Any ARP request will also be intercepted by the hypervisor and answered with the destination MAC address retrieved from the directory server.

4. VSITE DATA PLACE

VSITE data plane is shown in Figure 3. VSITE data plane makes use of the following techniques:

- Ethernet frame are carried using Ethernet over IP protocols such as Ethernet over GRE or EtherIP protocol (RFC3378) between CE_{A1} and CE_C, between CE_C and top-of-rack (ToR) switch (assume ToR is its `locIP`). For example, CE_C encapsulates the Ethernet frame received from CE_{A1} with an IP header destined to ToR switch of destination (the protocol field of IP header is 97 for EtherIP).

- Since ToR switch will remove the IP header and process the Ethernet frame. To prevent overlapping VLAN IDs from multiple enterprises, The VLAN ID has to be translated into a locally unique one. This is done at data center CE for traffic from enterprise site. For a Ethernet frame from cloud VMs, the translation is done at the hypervisor, and the Ethernet frame is then encapsulated with an outer IP header.

- For the MAC address of a VM, we can identify which enterprise it belongs to.

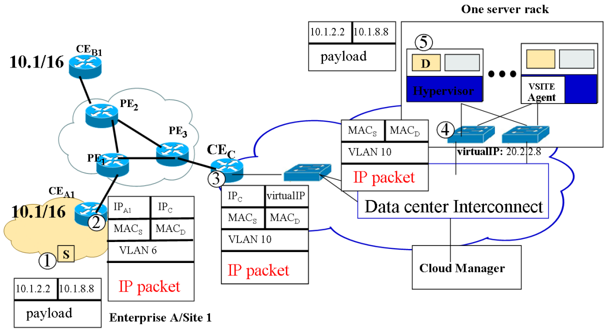

Figure 3: VSITE data plane

Inbound traffic: Figure 3 illustrates our data plane when S sends a packet to D in the vstub. To avoid excessive ARP flooding, CE_{A1} is configured to be the ARP proxy for IP prefixes in the cloud. CE_{A1} and CE_C exchanges MAC reachability information using OTV control plane protocol. ARP request from S will be answered by CE_{A1} with D's MAC address MAC_D. After obtaining destination MAC, S sends its packet (Step 1). When CE_{A1} receives the packet from S, it consults its MAC routing table for destination D. The outgoing interface is the IP address of cloud CE CE_c. Thus, the Ethernet frame is encapsulated with an IP header with destination CE_C (Step 2). The packet is then routed to CE_C. Packets received at CE_C are first decapsulated. They then get encapsulated with another IP header. For each packet, the source is the CE itself, the destination is the `locIP` of the destination. The MAC is the MAC of the destination. We need to translate the VLAN to a locally unique one. VLAN 6 is translated into VLAN 10. The VRF entry for MAC_D is kept up-to-date by the directory server. These are done in Step 3. The packet is then sent to the ToR switch. When the ToR switch receives the packet, it will remove the IP header and forwarding in layer 2 (Step 4). Since MAC_D is unique, the packet will be sent directly to the hypervisor of D.

Outbound traffic: When the hypervisor receives a packet from a VM with source, 10.1.8.8 and destination 10.1.2.2. It checks whether it has an entry in its forwarding table. If not, it queries the directory server for the destination `locIP`. The `locIP` is the IP of CE_C. The packet is then encapsulated with an IP header (destination is `locIP`) and routed. Because VM D's default gateway is its ToR switch, the MAC layer of the hypervisor will append a MAC header with VLAN 10 and its ToR's MAC as L2 destination. Upon receipt at the logical CE, the packet will first be decapsulated, then encapsulated with another IP header (destination CE_{A1}) before routed to the enterprise site with prefix 10.1/16.

Cloud internal traffic: If the source and destination have the same `locIP`, then the packet will be sent in layer 2 with the appropriate source MAC and destination MAC as well as VLAN ID in the header. If they are different, then the process is the same as outbound traffic except that the `locIP` and destination MAC will correspond to the destination VM's `locIP` and destination VM's MAC.

4.1 VLAN issues

VLAN processing: VM D needs to join in its appropriate VLAN through the hypervisor. Each local switch will only need to be configured with the VLANs of its attached VMs.

Reuse of VLAN ID: VLAN ID may be reused. If a VLAN ID gets reassigned from enterprise A to B due to VM shutdown or migration, VMs of enterprise A may have stale entries and still send packets to VMs of the same VLAN ID in a particular location. However, since the MAC address of different enterprises are partitioned. The packets from enterprise A will be dropped at the destination hypervisor. If MAC address filtering is enabled on VLANs, then enterprise A's packet will be dropped at ToR because the MAC does not match enterprise B (the VLAN is currently allocated to B).

5. DISCUSSION

5.1 Security

vstub attachment security: By default, a OTV like protocol only transport Ethernet frames across geographically separated data centers. It does not specify virtual private connection. However, this can be easily achieved by using IPsec encapsulation between CE_{A1} and CE_C.

VM attack from other tenants: The hypervisor implements a virtual switch with VLAN support. Broadcast and multicast traffic is only sent to the VMs with the same VLAN ID of the same enterprise. Recall we encode enterprise ID in the MAC address. Thus, it is not possible to carry out MAC address spoofing attack without compromising the hypervisor. We assume the hypervisor is secure.

It is not possible for a VM to impersonate as a VM of another enterprise. The hypervisor knows which VM belongs to which enterprise. It is also not possible to carry out a DDoS attack. The attacker's hypervisor will drop those packets.

VM attacks through cover channels have been uncovered recently [7]. Our design mitigates these attacks as VMs do not have cloud provider's publicly reachable IP addresses configured. Since these VMs are configured with enterprise IP addresses, in most cases, these VMs should not be reachable from Internet hosts. Thus, inducing workload by VMs of other tenants is not possible. Cover channels attacks are hard to rule out as pointed out in [7]. Thus, to be free from these attacks, VMs of different tenants should not co-locate at the same physical machine.

VLAN networking attack from other tenants: In traditional Ethernet, it is possible to mount MAC flooding attack of one VLAN from hosts belong to another VLAN. MAC flooding exhaust MAC forwarding table entries. Thus, MAC frames whose destinations are kicked out the table will have to be flooded throughout the VLAN. The problem is that forwarding table of different VLANs are shared. In VSITE, all Ethernet frames have to go through the hypervisor. Hypervisor will drop frames with invalid MAC addresses. In VSITE, data center edge switches are designed to hold maximum MAC entries if all physical machines in the edge are instantiated with the maximum number of VMs.

5.2 Ethernet data center core and edge

In the case that both core and edge are Ethernet, all we need is to replace the current encapsulation protocol with MACinMAC protocol. Instead of location IP, we will have location MAC.

5.3 Bandwidth isolation

Bandwidth isolation will depend on the bandwidth allocation scheme. If each enterprise is allocated a pre-specified amount of total bandwidth, then we can use cloud-wide rate limiting [6] to achieve bandwidth isolation.

Cloud provider may not fix the bandwidth allocation for each enterprise. Rather it may just enforce fairness. Fairness can be enforced in VSITE by fair bandwidth allocation among VLANs. In particular, each enterprise gets the same weight (if all enterprises have the same priority). This weight is split among all VLANs of the enterprise. A control protocol can adjust the weight in switches (e.g. ToR switch).

5.4 L2 broadcast

L2 broadcast of a VLAN is nicely support in VSITE. If S broadcasts, CE_{A1} will encapsulates the Ethernet frame with the IP address of CE_C (if there are multiple MAC forwarding entries for the broadcast address, it will be done for each entry). The packet will be routed first to CE_C, then ToR switch where it will be released into L2.

When VM D broadcasts, each VM in the local VLN 10 will receive the packet. When ToR receives the packet, it will be transported to CE_C using Ethernet over IP protocol. This in turn gets sent to CE_{A1}.

6. IMPLEMENTATION AND SCALABILITY

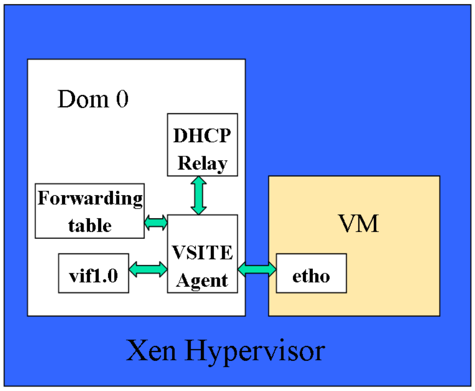

Figure 4: VSITE hypervisor implementation in Xen

An example Xen-based implementation of our hypervisor is shown in Figure 4. The key operation in control plane is a directory query when a forwarding entry is not found. This has been shown to be very scalable [3]. The operations in the data plane are IP encapsulation. This overhead is small.

We briefly discuss the scalability of our solution here. Suppose a large data center design [8] is built with 100 containers, each container having 32 racks and each rack containing 64 servers, which yields a total number of 204,800 physical servers. Assuming each server can host 8 VMs; each rack can hold 2 /24 subnets if every address is allocated, and 4 /24 subnets if only half of each subnet is allocated. Assume we need ToR switches as the `locIP`s. With all ToR switches in a rack sharing one virtual IP (using VRRP), this results in 100*32 `locIP` and at most 3200*4 subnets. Each subnet can be used to extend a VLAN from an enterprise. L2 traffic for a VLAN is confined to reach only the 64 servers in the rack.

7. RELATED WORK

The work most closely related to ours is CloudNet [9], Amazon VPC and VPN-Cubed [2]. CloudNet relies on the VLAN feature for tenant seperation and thus is much less scalable in terms of the number of virtual private clouds (VPCs) or enterprises it can support. In addition, VPLS floods L2 broadcast to the enterprise. Amazon VPC offers seamless extension in L3. Not much technical details are known. VPN-Cubed is a commercial product; it is basically an overlay solution.

Other related work are data center networking [3, 5, 4]. The idea of separating location address from actual address has been applied in these data center networking papers. However, they do not discuss any support for private address space and security isolation among enterprise users.

8. CONCLUSION AND FUTURE WORK

Today's enterprises can benefit significantly if resources can be tapped seamlessly and dynamically from a public cloud. We study the problem of providing such cloud extension in this paper. We divide the problem into two related parts: vstub creation and vstub attachment, which adds significant clarity to the whole design space. We then present a scalable and secure architecture called VSITE that offers a layer 2 solution for enterprise vstub construction. VSITE uses a combination of location IP and enterprise ID (encapsulated in MAC) to tunnel enterprise packets with potentially overlapping addresses. A key design goal of VSITE is to minimize the impact of VM dynamics on the underlying network configuration, thus making it highly scalable. We are implementing VSITE in our test network.

9. REFERENCES

[1] Cisco. Overlay transport virtualization (otv). `http://www.cisco.com/en/US/prod/switches/ps9441/nexus7000_promo.html`.

[2] C. F. T. Corp. Vpn-cubed: customer controlled security for the cloud. `http://www.cohesiveft.com/Cube/VPN/VPN-Cubed_Custom_Enterprise_Configur%ations/`.

[3] A. Greenberg, N. Jain, S. Kandula, C. Kim, P. Lahiri, D. Maltz, P. Patel, and S. Sengupta. Vl2: A scalable and flexible data center network. In *ACM SIGCOMM*, 2009.

[4] C. Kim, M. Caesar, and J. Rexford. Floodless in seattle: a scalable ethernet architecture for large enterprises. In *SIGCOMM '08: Proceedings of the ACM SIGCOMM 2008 conference on Data communication*, pages 3–14, New York, NY, USA, 2008. ACM.

[5] A. Pamboris, N. Farrington, N. Huang, P. Miri, S. Radhakrishnan, V. Subramanya, and A. Vahdat. Portland: A scalable fault-tolerant layer 2 data center network fabric. In *ACM SIGCOMM*, 2009.

[6] B. Raghavan, K. Vishwanath, S. Ramabhadran, K. Yocum, and A. C. Snoeren. Cloud control with distributed rate limiting. In *SIGCOMM '07: Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 337–348, New York, NY, USA, 2007. ACM.

[7] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In *CCS '09: Proceedings of the 16th ACM conference on Computer and communications security*, pages 199–212, New York, NY, USA, 2009. ACM.

[8] R. systems. Ice cube modular data center.

[9] T. Wood, A. Gerber, K. K. Ramakrishnan, and J. V. D. Merwe. The case for enterprise-ready virtual private clouds. In *USENIX HotCloud*, 2009.