Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 2, 2012

M. Hasan
Cisco Systems, Inc.
A. Chari
D. Fahed
France Telecom - Orange Labs
L. Tucker
M. Morrow
M. Malyon
Cisco Systems, Inc.
March 2012

A framework for controlling Multitenant Isolation, Connectivity and Reachability in a Hybrid Cloud Environment
draft-masum-chari-shc-00

Abstract

Multiple enterprises (tenants) consuming resources in a public Cloud shares the physical infrastructure of one or more DCs out of which the Cloud resources are serviced. Hence one of the major features that has to be supported in public Cloud DCs is multitenant isolation, which is realized via various DC isolation technologies, such as VLAN or VxLAN. In a hybrid Cloud environment where a public Cloud (more specifically off-premises public Cloud resources acquired by a tenant) becomes an _extension_ of a tenant intranet or private Cloud, the multitenant isolation capability has to be extended beyond the public Cloud DCs. The multitenant isolation _domain_ has to span end-to-end from the tenant network or on-premises resources via the MAN/WAN and the public Cloud DC networks to tenant off-premises resources. While multitenant isolationI isolates one tenant from another (inter-hybrid Cloud isolation), an enterprise may desire controlled connectivity to a hybrid Cloud from another Cloud or network or tenant or select resources. In addition, there may be need for controlling direct reachability of resources within a hybrid Cloud itself (intra-hybrid Cloud). The tenant network may be connected to the public Cloud (DCs) over the Internet or a private IP/MPLS MAN/WAN owned or operated by a service provider, which also may support PPVPN (Provider Provided VPN) service, such as the L3 MPLS VPN. In this work we consider the latter type of network and describe a framework for supporting inter-hybrid Cloud multitenant isolation, inter-hybrid Cloud connectivity and intra-hybrid Cloud reachability.

Status of this Memo

This Internet-Draft is submitted in full conformance with the

Copyright Notice

Table of Contents

1. Introduction

A Cloud service provider (CSP) offers services to tenants (enterprises, enterprise departments, employees, end consumers) out of one or more Cloud DCs, where a tenant can acquire or release (CRUD: Create, Read, Update, Delete) compute, storage or network resources on-demand and anytime. The CSP exposes Cloud service interfaces to tenants which are used by tenants to CRUD resources.

The NIST definition of Cloud computing [NIST] has defined following Cloud deployment models:

- Private Cloud: A Cloud for use by an enterprise only, where the Cloud infrastructure and services are owned and/or operated by the enterprise IT or a 3rd party.

- Public Cloud: A Cloud that can be used by anyone and owned/operated/offered by a Cloud Service Provider.

- Hybrid Cloud: A Cloud consisting of multiple interoperable Clouds that enables data and application portability.

- Community Cloud: The Cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns.

The NIST definition of the hybrid Cloud refers to multiple interoperable Clouds to which we include a case where an enterprise may not have a private Cloud. Following are the options in which a hybrid Cloud can be architected or deployed:

1. A1: A tenant intranet and a single public Cloud.

2. A2: A tenant Private Cloud and a single public Cloud.

3. A3: A tenant intranet or private Cloud and multiple public Clouds.

4. A4: A Public Cloud and one or more other public Clouds. In this case the source public Cloud may acquire resources from another public Cloud on behalf of its tenant.

5. A5: A3 + A4.

In this document we focus on A1 and A2 (the other options require further considerations and will be addressed in future work). We assume that the same SP owns or operates both the public Cloud (DCs) and the private MAN/WAN. Employing PPVPN (Provider Provided VPN,

such as L3 MPLS VPN) features we describe a framework that facilitates inter-hybrid Cloud multitenant isolation and connectivity and intra-hybrid Cloud direct reachability between resources in the same hybrid Cloud. Note that the realization framework discussed in this document covers only the MAN/WAN segment of the end-to-end network of a hybrid Cloud.

1.1. Acronyms and Definitions

Tenant: An enterprise, enterprise department, enterprise user or end consumer.

CRUD: Create, Read, Update, Delete (of resources, entities).

ONPR: On-premises resources - Tenant intranet or private Cloud resident resources.

OFPR: Off-premises resources - Resources acquired by a tenant in a public Cloud on-demand.

MTI: Multitenant Isolation - Isolate traffic and routing/forwarding/switching instances from one tenant or hybrid Cloud from another.

E2E: End-to-end - From ONPR via private MAN/WAN and public Cloud DC networks to OFPR.

PPVPN: Provider Provided VPN - VPN that is configured and managed by a service provider on behalf of a tenant.

L3MV: L3 MPLS VPN.

PE: Provider Edge router.

CE: Customer Edge router.

VRF: Virtual Routing and Forwarding instance.

VRF-Lite: VRF without need for various MPLS and L3MV features typically supported on CE or other DC routers.

CSP: Cloud Service Provider - Owns or operates a public and hybrid Cloud.

SHC: Seamless Hybrid Cloud - An instance of a logical hybrid Cloud.

ECRT: BGP Extended Community Route Target.

SRTR: Source MPLS PE router that originates or exports vpnv4 routes.

DRTR: Destination MPLS PE router that imports vpnv4 routes.

2. Logical Hybrid Cloud

The simplest case of a hybrid Cloud environment is where an enterprise connects to a public Cloud and consumes resources and all-to-all communication, connectivity, reachability and route advertisements are allowed (between the whole enterprise intranet or private Cloud and all the acquired off-premises resources). But enterprises should have the flexibility in defining what constitutes a hybrid Cloud and control communication, connectivity and reachability in and across a hybrid Cloud (for enhanced security, flexibility and enterprise specific needs).

While deploying a hybrid Cloud an enterprise (IaaS admin) should be able to specify following (it is expected that a CSP will support these requirements as part of its hybrid Cloud service and expose relevant tenant facing service interfaces so that an enterprise can choose these features; the full definition of services and tenant facing interfaces is beyond the scope of this document; Readers may check seamless Cloud abstraction, model and interfaces [SHCA]):

1. Logical hybrid Cloud: Creation of instances of hybrid Clouds, for example, each belonging to an organization within the enterprise or for a particular use case or application. We call an instance of a logical hybrid Cloud a _seamless hybrid Cloud (SHC)_. The concept of a tenant is accordingly generalized where each SHC has an associated tenant. Following are the components of an SHC that can be associated to or disassociated from an SHC on-demand (by a tenant IaaS admin) thus allowing flexible hybrid Cloud deployment:

    1. Sites: A set of tenant selected sites (DC, remote/branch). Sites not associated with an SHC will not be able to communicate with the SHC or resources in it. When a whole site is selected all the resources in it become accessible from within the SHC.

    2. ONPR: A set of on-premises (intranet or private Cloud) resources of a particular site. In this case the whole site is not accessible from within the SHC, only the selected ONPR.

    3. OFPR: A set of off-premises resources that are acquired in a public Cloud DC location.

2. Intra-SHC Direct Reachability: ONPR or OFPR that is _directly_ reachable within an SHC. For example, an ONPR or OFPR (such as a web tier) may be directly reachable from within an SHC, but an app tier is reachable indirectly via the web tier.

3. Inter-SHC connectivity/reachability: Which SHC can connect to or reach directly which other SHC.

The multitenant isolation will isolate both the data plane traffic and control plane elements, such as routing/forwarding/switching table instances.

Figure 1 shows an example of full view of a network and Cloud. An example of an SHC (as viewed by a tenant) is shown in Figure 2.

```
   -------------------------
   | O ONPR-DB  O Other    | T1 Site 1       FIGURE 1
   | |          | Resources|
   | O FW -------          |
   | |                     |   HSW: Hypervisor Switch
   | O ONPR-App1-0         |   VLB/FW: Virtual Load-balancer/Firewall
   | |                     |   ER: Edge Router, CE: Customer Edge
   | O FW       O ONPR-    |   VM: Virtual Machine, PE: Provider Edge
   | |          | App-D-0  |
   | O ONPR-DMZ |          |
   | |          |          |   T2 Site 1       T1 Site 5      T1 Site 7
   |---------------------- | -------------  -------------   ------------
   | Tenant DC Network   | | |  |        |  |  |        |   |  |       |
   |---------------------- | |  |        |  |  |        |   |  |       |
   |    |                  | |  |        |  |  |        |   |  |       |
   |----O CE11-------------| |--O CE21---|  |--O CE15---|   |--O CE17--|
        |                       |              |               |
        -------------------------              -----------------
                 |                                     |
                 O PE-TS1------------------------------O PE-TS2
                 |   SP Private (IP/MPLS) MAN/WAN      |
   CSP           O PE-CL1------------------------------O PE-CL2 CSP
   DC-Loc1       |                                     |        DC-Loc2
   |-------------O ER-CL1-------------------| |--------O ER-CL2--------|
   |             |                          | |        |               |
   |            -------------------------   | | ---------------------- |
   |            | CSP DC 1 Core/Aggr    |   | | |CSP DC 2 Core/Aggr  | |
   |            ---------O Access SW1----   | | ------O Access SW2---| |
   |                     |                  | |       | \              |
   |         ------------|-----------       | |       |   \            |
   |         |                      |       | |       |     \          |
   |      ---O HSW 1 ----------     O HSW 2 | | HSW 3 O       O HSW 4  |
   |      |           |       |     |       | |       |       |        |
   |      |           |       |     |       | |       |       |        |
   |      |           O VFW   |     |       | |       O VFW   |        |
   |      |           | T11   |     |       | |       | T12   |        |
   | ---------   ----------  ---- ------    | | -----------  --------- |
   |  |     |     |     |      |     |      | |  |      |       |      |
   |  O ... O     O ... O      O ... O ...  | |  O  ... O       O ...  |
   |  T1    T1    T1    T1     T2    T1     | |  T1     T1      T1     |
   |  VM11  VM12  VM13  VM14   VM21  VM15   | |  VM16   VM17    VM18   |
   |     OFPR-      OFPR-           OFPR-  | |     OFPR-       OFPR-  |
   |     DMZ1       App1-1          App-D-1| |     App1-2      App-D-2|
   ------------------------------------------ --------------------------
```

```
   -------------------------
   | O ONPR-DB             | T1 Site 1       FIGURE 2
   | |                     |
   | O FW                  |
   | |                     |   HSW: Hypervisor Switch
   | O ONPR-App1-0         |   VLB/FW: Virtual Load-balancer/Firewall
   | |                     |   SHC Cloud Abstractions:
   | O FW       O ONPR-    |     ST: Site  VST: Virtual Site
   | |          | App-D-0  |     CDCL: Cloud DC Location
   | O ONPR-DMZ |          |
   | |          |          |                                 T1 Site 7
   |---------------------- |                                ------------
   |                     | |                                |  |       |
   |---------------------- |                                |  |       |
   |    |                  |                                |  |       |
   |----O ST11-------------|                                |--O ST17--|
        |                                                      |
        |---------                                     ---------
                 |                                     |
                 O-------------------------------------O
                 |              T1-SHC-1               |
   CSP           O-------------------------------------O        CSP
   DC-Loc1       |                                     |        DC-Loc2
   |-------------O CDCL1---------------------| |--------O CDCL2--------|
   |           /                             | |        |              |
   |         /                               | |        |              |
   |         |                               | |        |              |
   |         |                               | |        |              |
   |         |                               | |        |              |
   |         |                               | |        |              |
   |         |                               | |        |              |
   |   VST-  O ------------                  | |   VST- O --------     |
   |   T11 / |            |                  | |   T12  |        |     |
   |     /   |            |                  | |        |        |     |
   |    |    O VFW        |                  | |        O VFW    |     |
   |    |    | T11        |                  | |        | T12    |     |
   | ------ --------     ----                | | -----------  ---------|
   |  |        |          |                  | |  |      |      |      |
   |  O ...    O ...      O ...              | |  O  ... O      O ...  |
   |  T1       T1         T1                 | |  T1     T1     T1     |
   |  VM11     VM13       VM15               | |  VM16   VM17   VM18   |
   |  OFPR-    OFPR-      OFPR-              | |     OFPR-      OFPR-  |
   |  DMZ1     App1-1     App-D-1            | |     App1-2     App-D-2|
   ------------------------------------------ --------------------------
```

3. Realization Framework

The realization mechanism described in this document is based on the assumption that the CSP not only owns or operates the public Cloud DCs, but also owns or operates the private IP/MPLS MAN/WAN connecting tenant sites and the CSP DCs (note that with proper framework or interfaces in place this restriction can be lifted). We outline a framework for supporting inter-SHC isolation (multitenant isolation), intra SHC isolation (reachability rules between SHC components) and inter-SHC connectivity (akin to extranet), which is as follows (note that the realization framework covers only the MAN/WAN segment of the E2E network of an SHC):

o The inter-SHC isolation is realized by mapping an SHC to an instance of an L3 MPLS VPN [RFC4364] (L3MV) identified with a VPN ID [RFC2685] to uniquely identify the SHC (SHC-L3MV). The L3MV technology provides a framework for multitenant isolation (of traffic and routes) in the IP/MPLS MAN/WAN. For SHC, the multitenant isolation has to be extended into the public Cloud DCs to incorporate the OFPR in the SHC-L3MV. As shown in Figure 2, the set of OFPR can be considered as virtual L3MV sites, which in the CSP DC network spans from the DC edge (router) to the OFPR resources. On the DC edge the MTI can be realized via the L3MV multitenat feature on customer edge (CE) router, where multiple customers of an SP (such as in a multitenant building) are supported (via VRF-lite together with subinterface or other mechanism to separate traffic and routes of each tenant in the multitenant CE). Following are various options of mapping an SHC to L3MV:

  * Assuming that the tenant is already on an L3MV with a SP (where all the tenant sites are connected via the private IP/MPLS MAN/WAN), which is also a CSP:

    1. V1: Extend the existing L3MV (L3MV-Original) to include OFPR DC locations (to include OFPR resources) as _extended multitenant L3MV sites_ (set of OFPR of the SHC in a CSP DC becomes virtual L3MV site) of the L3MV-Original. No new L3MV created (see below). This option requires updates of existing configurations every time SHC or its elements (described above) are CRUD on-demand.

    2. V2: Create a separate L3MV for an SHC (L3MV-SHC-Extr) identified by its unique VPN ID (different from L3MV-Original). The L3MV-SHC-Extr is then _connected_ as an _extranet_ to the L3MV-Original.

3. V3: Create a separate L3MV for an SHC (L3MV-SHC) identified by its unique VPN ID, but L3MV-SHC stands on its own without being connected to the L3MV-Original.

* In the case where a tenant is not already on an L3MV with a CSP, the V3 option above will cover it.

o Map SHC components (Site, ONPR and OFPR) that are specified as directly reachable to BGP Extended Community Route Targets [RFC4360] (ECRT) in L3 MPLS VPN [RFC4364]. While a single resource can be mapped to an ECRT, typically an ONPR or OFPR will be an IP address prefix, subnet or DC tier (such as web, app and DB tiers). With this mapping only routes for selected SHC components will be exchanged between _relevant_ MPLS PEs.

o Export the ECRTs from the _source MPLS PE router_ (export of ECRT results in MP-BGP update message to relevant PEs with MP_REACH_NLRI attribute containing VPNv4 address, ECRT and other parameters). Referring to Figure 1, examples of source MPLS PEs are PE-TS1, which is the source for ONPR-DMZ and PE-CL1 for OFPR-App-D2 .

o Import the ECRTs in all _relevant_ _destination MPLS PE routers_ of the SHC. Referring to Figure 1, examples of _relevant_ destination PEs are PE-TS1, PE-TS2 and PE-CL1 for OFPR-App-D2.

o On public Cloud DC networks the L3MV corresponding an SHC can be mapped to any of existing (such as VRF-lite, VLAN) or new generation DC isolation technologies, such as VxLAN [VXLN].

o When a component is detached from an SHC, withdraw the ECRT (resulting in MP-BGP update message with MP_UNREACH_NLRI).

4. Use Cases

We provide two use cases to explain the framework.

4.1. Use Case 1 - One SHC

This use case shows creation of a flexible logical hybrid Cloud or SHC by selectively associating ONPR, OFPR and enterprise sites with the SHC. It also shows intra-SHC direct reachability.

Referring to Figure 1, a tenant T1 creates an SHC T1-SHC-1 and associates following components to the SHC:

1. ONPR-DB, which is not directly reachable from within T1-SHC-1 (only via ONPR-App1-0), but accessible in the SHC.

2. ONPR-App1-0, which is not directly reachable (only via ONPR-DMZ), but accessible in the SHC.

3. ONPR-App-D-0, which is directly reachable.

4. ONPR-DMZ, which is directly reachable.

5. T1 Site 1 as a whole is not associated with the SHC. Hence other resources of Site 1 will not be accessible from within T1-SHC-1.

6. T1 Site 5 is not associated with T1-SHC-1.

7. T1 Site 7 is associated with T1-SHC-1. Hence all the resources in it will be accessible from within T1-SHC-1.

8. T1 acquires OFPR-DMZ1 resources in public Cloud DC location DC-Loc1. These resources are directly reachable.

9. T1 acquires OFPR-App1 resources (OFPR-App1-1) in public Cloud DC location DC-Loc1. These resources are not directly reachable, rather via ONPR-DMZ or OFPR-DMZ1. It is assumed that the DMZ web-servers are globally load-balanced to serve requests to instances of App1.

10. T1 acquires OFPR-App-D resources (OFPR-App-D-1) in public Cloud DC location DC-Loc1. These resources are directly reachable.

11. T1 acquires OFPR-App1 resources (OFPR-App1-2) in public Cloud DC location DC-Loc2. These resources are not directly reachable.

12. T1 acquires OFPR-App-D resources (OFPR-App-D-2) in public Cloud DC location DC-Loc2. These resources are directly reachable.

The realization of T1-SHC-1 is as follows:

- The T1-SHC-1 is mapped to an L3MV with a unique VPN ID (L3MV-1) that facilitates inter-SHC MTI. Each SHC component described above is mapped as follows:

-

  1. ONPR-DB is not mapped to an ECRT, that is, its routes will not be advertised in the L3MV-1.

  2. ONPR-App1-0 is not mapped to an ECRT, that is, its routes will not be advertised in the L3MV-1.

  3. ONPR-App-D-0 is mapped to an ECRT, that is, its routes will be advertised in the L3MV-1. The ECRT is exported by the PE-TS1 and imported by PE-TS2 (to be directly reachable from the Site 7), PE-CL1 (to be directly reachable from the OFPR resources in DC-Loc1) and PE-CL2 (to be directly reachable from the OFPR resources in DC-Loc2).

  4. ONPR-DMZ is mapped to an ECRT and exported by PE-TS1 and imported by PE-TS2, PE-CL1 and PE-CL2.

  5. T1 Site 1 is not mapped to an ECRT.

  6. T1 Site 5 is not mapped to an ECRT.

  7. T1 Site 7 is mapped to an ECRT, exported by PE-TS2 and imported by PE-TS1, PE-CL1 and PE-CL2.

  8. OFPR-DMZ1 is mapped to an ECRT, exported by PE-CL1 and imported by PE-TS1, PE-TS2 and PE-CL2.

  9. OFPR-App1-1 is not mapped to an ECRT.

  10. OFPR-App-D-1 is mapped to an ECRT, exported by PE-CL1 and imported by PE-TS1, PE-TS2 and PE-CL2.

  11. OFPR-App1-2 is not mapped to an ECRT.

  12. OFPR-App-D-2 is mapped to an ECRT, exported by PE-CL2 and imported by PE-TS1, PE-TS2 and PE-CL1.

4.2. Use Case 2 - Multiple SHC

This use case shows the case of inter-SHC connectivity, where two SHCs are connected in a controlled way. In this use case tenant at the Site 7 accesses App1 instances located on-premises in T1 Site 1 via load balancing through ONPR-DMZ or OFPR-DMZ1. Site 7 is associated with a (dedicated) SHC T1-SHC-7 and App1, DB and DMZ resources are grouped into a second SHC T1-SHC-2.

Referring to Figure 3, a tenant T1 creates an SHC T1-SHC-7 and associates following components to the SHC:

1. T1 Site 7.

The realization of T1-SHC-7 is as follows:

- The T1-SHC-7 is mapped to an L3MV with a unique VPN ID (L3MV-7) that facilitates inter-SHC MTI. Each SHC component described above is mapped as follows:

-

    1. T1 Site 7 is mapped to an ECRT ECRT7, that is, its routes will be advertised into the L3MV-7. The ECRT7 is exported by PE-TS2 (L3MV-7 VRF) and imported into L3MV-2 at PE-TS1 and PE-CL1 (since connectivity to T1-SHC-2 is allowed; see below).

Referring to Figure 3, a tenant T1 creates an SHC T1-SHC-2, allows connectivity to T1-SHC-7, and associates following components to the SHC:

1. ONPR-DB, which is not directly reachable (only via ONPR-App1).

2. ONPR-App1, which is not directly reachable (only via ONPR-DMZ).

3. ONPR-DMZ, which is directly reachable from within T1-SHC-2 and T1-SHC-7.

4. T1 acquires OFPR-DMZ1 resources in public Cloud DC location DC-Loc1. These resources are directly reachable from within T1-SHC-2 and T1-SHC-7.

The realization of T1-SHC-2 is as follows:

- The T1-SHC-2 is mapped to an L3MV with a unique VPN ID (L3MV-2) that facilitates inter-SHC MTI. Each SHC component described above is mapped as follows:

- o

1. ONPR-DB is not mapped to an ECRT.

2. ONPR-App1 is not mapped to an ECRT.

3. ONPR-DMZ is mapped to an ECRT, exported by PE-TS1 and imported into L3MV-2 at PE-CL1 and into L3MV-7 at PE-TS2.

4. OFPR-DMZ1 is mapped to an ECRT, exported by PE-CL1 and imported into L3MV-2 at PE-TS1 and L3MV-7 at PE-TS2.

```
   -------------------------
   | O ONPR-DB             | T1 Site 1     FIGURE 3
   | |                     |
   | O FW                  |
   | |                     |   HSW: Hypervisor Switch
   | O ONPR-App1-0         |   VLB/FW: Virtual Load-balancer/Firewall
   | |                     |
   | O FW                  |
   | |                     |
   | O ONPR-DMZ            |
   | |                     |                                 T1 Site 7
   |---------------------- |                                ------------
   ||                    | |                                |  |       |
   |---------------------- |                                |  |       |
   |    |                  |                                |  |       |
   |----O CE11-------------|                                |--O CE17--|
        |                                                      |
        -------------------------              -----------------
                 |                                     |
                 O PE-TS1------------------------------O PE-TS2
                 |   SP Private (IP/MPLS) MAN/WAN      |
   CSP           O PE-CL1------------------------------O PE-CL2
   DC-Loc1       |
   |-------------O ER-CL1-------------------|
   |             |                          |
   |            -------------------------   |
   |            | CSP DC 1 Core/Aggr    |   |
   |            ---------O Access SW1----   |
   |                     |                  |
   |         ------------|                  |
   |         |                              |
   |      ---O HSW 1                        |
   |      |                                 |
   |      |                                 |
   |      |                                 |
   |      |                                 |
   | ---------                              |
   |  |                                     |
   |  O                                     |
   |  T1                                    |
   |  VM11                                  |
   |     OFPR-                              |
   |     DMZ1                               |
   ------------------------------------------
```

5. Discussion

In this section we discuss various issues that needs future considerations.

5.1. Network Management

The framework described in this document is a network management (NM) framework. Hence the framework can be used on any _existing network without any changes_. But NM frameworks typically are not built for on-demand operations, as required in a Cloud environment. As described above, creation or deletion of an SHC should result in creation or deletion of L3MV on-demand so that pay-per-use accounting are turned on or off on-demand. Similarly, association or disassociation of SHC components should result in creation or deletion and export or withdrawal of ECRT on-demand. Currently, these actions can be performed via on-demand application of configuration. But current network (CLI or even NetConf based) configuration/provisioning frameworks are cumbersome to use in an on-demand environment. Hence proper Cloud-ready NM framework and interfaces are needed.

5.2. Protocol, Control Plane Features

In current L3MV the import of ECRT in respective L3MV VRF can only be done via (static) pre-configuration, which is not suitable in an on-demand Cloud environment. This process can be automated, if an L3MV is identified by a unique VPN ID and the information is carried in MP-BGP ECRT update messages.

As described above, the multitenant isolation (MTI) has to reach all the way to the OFPR. But currently E2E MTI (as required in a hybrid Cloud environment) can only be achieved by stitching together multiple isolation technologies (such as L3MV + VRF-Lite + VLAN or L3MV + VxLAN) with their own limitations. Homogenous E2E MTI technology is desirable. For example, L3MV all the way to a TOR switch (TS), where the TS can functions as a PE and the virtual access switches on hypervisors as L3MV multitenant CE. The L3MV technology handles the case of overlapping private IP addresses of multiple tenants. Hence bringing the technology all the way to the hypervisor will be useful.

6. Security Considerations

Typical security considerations of L3MV configuration will apply. In an on-demand dynamic Cloud environment certain security issues may be amplified. For example, uncontrolled BGP updates as resources are CRUD very dynamically (by a rogue entity). The dynamic application of configuration may amplify the situation of mistaken route leaking from one SHC to another. Hence proper steps should be taken.

7. IANA Considerations

There is no IANA consideration.

8. Conclusion

We have discussed a framework for realizing inter-hybrid Cloud end-to-end multitenant isolation, intra-hybrid Cloud reachability and inter-Hybrid Cloud controlled connectivity. The framework allows creation of a logical hybrid Cloud, we call a seamless hybrid Cloud (SHC), consisting of tenant selected subset of enterprise sites, on-premises resources (resident in tenant network) and off-premises resources (acquired by the tenant on-demand in public Cloud DCs). An SHC is mapped to an L3 MPLS VPN (L3MV) to support inter-SHC multitenant isolation and SHC components are mapped to BGP extended community route targets (ECRT) so that route advertisements of SHC associated components can be controlled and isolated from any other Cloud or non-Cloud L3MV. The intra-SHC reachability and inter-SHC connectivity are also controlled via SHC-component to ECRT mapping and via proper import/export policies involving ECRTs. The framework is network management based allowing support of it on existing infrastructure. We have discussed few issues above (in the Discussion section), which should be addressed. We have not covered the cases A3-A5 discussed in the Introduction section, which also should be addressed.

9. References

9.1. References

9.2. Normative References

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.

[RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, February 2006.

9.3. Informative References

[NIST] Mell, P. and T. Grance, "The NIST Definition of Cloud Computing", 800-145 NIST, September 2011, <http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>.

[SHCA] Hasan, M., Morrow, M., Tucker, L., Gudreddi, S., and S. Figueira, "SEAMLESS CLOUD ABSTRACTION, MODEL AND INTERFACES", In Proc. ITU/IEEE Kaleidoscope Conference, December 2011, Cape Town, South Africa, December 2011, <http://www.itu.int/dms_pub/itu-t/oth/29/05/T29050000160001PDFE.pdf>.

[VXLN] M.Mahalingam, M. and et.al, "VXLAN: A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", draft-mahalingam-dutt-dcops-vxlan-00.txt, August 2011, <http://tools.ietf.org/html/draft-mahalingam-dutt-dcops-vxlan-00>.

[RFC2685] Fox, B. and B. Gleeson, "Virtual Private Networks Identifier", RFC 2685, September 1999, <http://tools.ietf.org/html/rfc2685>.

Authors' Addresses

Masum Z. Hasan
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
USA

Email: masum@cisco.com

Abdelhadi Chari
France Telecom - Orange Labs
2, avenue Pierre Marzin
Lannion ,   22307
France

Email: abdelhadi.chari@orange.com

David Fahed
France Telecom - Orange Labs
2, avenue Pierre Marzin
Lannion,   22307
France

Email: david.fahed@orange.com

Lew Tucker
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
USA

Email: letucker@cisco.com

Monique Morrow
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
USA

Email: mmorrow@cisco.com

Mark Malyon
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
USA

Email: mmalyon@cisco.com