Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 05, 2014

YJ. Stein
Y. Gittik
RAD Data Communications
D. Kofman
K. Katsaros
LINCS
M. Morrow
L. Fang
Cisco Systems
W. Henderickx
Alcatel-Lucent
July 04, 2013

Accessing Cloud Services
draft-stein-cloud-access-03.txt

Abstract

Cloud services are revolutionizing the way computational resources are provided, but at the expense of requiring an even more revolutionary overhaul of the networking infrastructure needed to deliver them. Much recent work has focused on intra- and inter-datacenter connectivity requirements and architectures, while the "access segment" connecting the cloud services user to the datacenter still needs to be addressed. In this draft we consider tighter integration between the network and the datacenter, in order to improve end-to-end Quality of Experience, while minimizing both networking and computational resource costs.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 05, 2014.

Table of Contents

1. Introduction

Cloud services replace computational power and storage resources traditionally located under the user's table or on the user's in-house servers, with resources located in remote datacenters. The cloud resources may be raw computing power and storage (Infrastructure as a Service - IaaS), or computer systems along with supported operating systems and tools (Platform as a Service - PaaS), or even fully developed applications (Software as a Service - SaaS). Processing power required for the operation of network devices can also be provided (e.g., Routing as a Service - RaaS). The inter- and intra-datacenter networking architectures needed to support cloud services are described in [I-D.bitar-datacenter-vpn-applicability].

The advantages of cloud services over conventional IT services include elasticity (the ability to increase or decrease resources on demand rather than having to purchase enough resources for worst case scenarios), scalability (allocating multiple resources and load-balancing them), high-availability (resources may be backed up by similar resources at other datacenters), and offloading of IT tasks (such as applications upgrading, firewalling, load balancing, storage backup, and disaster recovery). These translate to economic

efficiencies if actually delivered. The disadvantages of cloud service are lack of direct control by the customer, insecurity regarding remote storage of sensitive data, and communications costs (both direct monetary and technical such as lack of availability and additional transaction latency).

The cloud service user connects to cloud resources over a networking infrastructure. Today this infrastructure is often the public Internet, but (for reasons to be explained below) is preferably a network maintained by a Network Service Provider (NSP). The datacenter(s) may belong to the NSP (which is the case considered by [I-D.masum-chari-shc]), or may belong to a separate Cloud Service Provider (CSP), and accessible from the NSP's network. In the latter case there may or not be a business relationship between the NSP and CSP, the strongest such relationship being when either the NSP or CSP offers a unified "bundled" service to the customer.

In order to obtain the advantages of cloud service without many of the disadvantages, the cloud services customer enters into a Service Level Agreement (SLA) with the CSP. However, such an SLA by itself will be unable to guarantee end-to-end service goals, since it does not cover degradations introduced by the intervening network. Indeed, if the datacenter is accessed over the public Internet, end-to-end service goals may be unattainable. Thus an additional SLA with the NSP (that may already be in effect for pre-cloud services) is typically required. When the CSP and the NSP are the same entity but not offering a bundled service, these SLAs may still be separate documents.

Cloud services require a fundamental rethinking of the Information Technology (IT) infrastructure, due to the requirement for dynamic changes in IT resource configuration. Physical IT resources are replaced by virtualized ones packaged in Virtual Machines (VMs). VMs can be created, relocated while running (VM migration), and destroyed on-demand. Since VMs need to interconnect, connect to physical resources, and connect to the cloud services user, they need to be allocated appropriate IP and layer 2 addresses. Since these addresses need to be allocated, moved, and destroyed on-the-fly, the cloud IT revolution directly impacts the networking infrastructure. Recent work, such as [I-D.bitar-datacenter-vpn-applicability], has focused on requirements and architectures for connectivity inside and between datacenters. However, the "access segment", that is, the networking infrastructure connecting the cloud services user to the datacenter, has not been fully addressed.

The allocation, management, manipulation, and release of cloud resources is called "orchestration" (see [I-D.dalela-orchestration]). Orchestrators need to respond to user demands and uphold user SLAs

(perhaps exploiting virtualization techniques such as VM migration) while taking into account the location and availability of IT resources, and optimizing the CSP's operational objectives. These objectives include, for example, decreasing costs by consolidating resources, balancing use of resources by reallocating computational and storage resources, and enforcing engineering, business, and security policies. Orchestrators of the present generation do not attempt optimization of CSP's networking resources, but this generalization is being studied [I-D.ietf-nvo3-framework]. Furthermore, these orchestrators are completely oblivious to the NSP's resources and objectives. Hence, there is no mechanism for maintaining end-to-end SLAs, or for optimizing end-to-end networking.

This goal of this Internet Draft is to kick off discussions on requirements and possible mechanisms for improving end-to-end Quality of Experience while minimizing both networking and computational costs.

2. Model of Existing Cloud Services

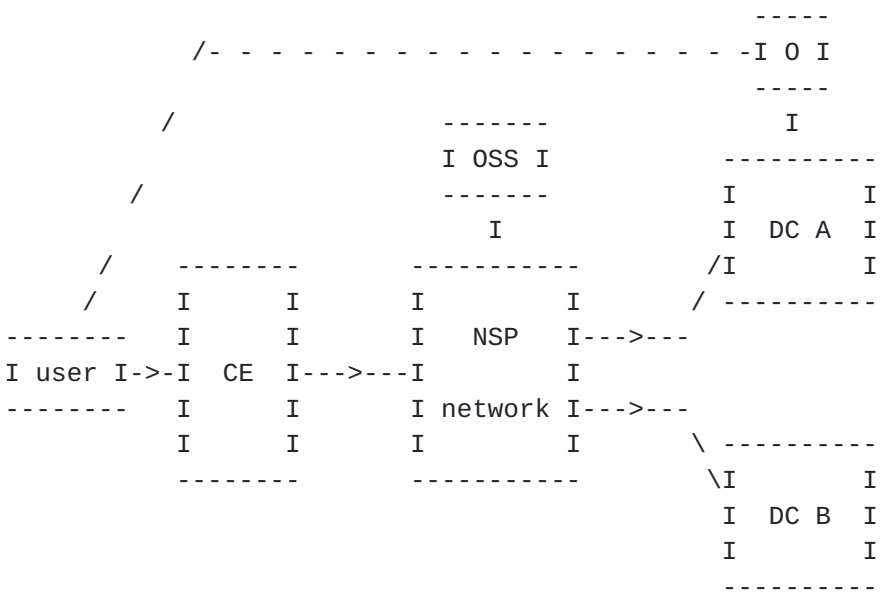

Figure 1: Simplified model of cloud service provided over Service Provider network to an enterprise customer behind a CE device

For concreteness, we will assume the scenario of Figure 1. On the left we see a cloud services user attached to a customer site network. This network connects to the outside world via a Customer Edge (CE), which may be a branch-site router or switch, a special purpose cloud demarcation device, or in degenerate cases the user's computer itself. The NSP network is assumed to be a well-engineered

network providing VPN and other SLA-based services to the customer site. The NSP network is managed from an Operations Support System (OSS), which may include a Business Support System (BSS), the latter being needed for interfacing with the customer for approval of service reconfiguration, billing issues, etc. In some cases, the functionality needed here may be obtained by interfacing with a Looking Glass server or a Policy and Charging Rules Function (PCRF). Connected to this network are datacenters (two are shown - datacenter A and datacenter B), which may belong to the NSP, or to a separate CSP. The orchestrator of datacenter A is depicted as "O". Additionally, Internet access may be available directly from the CE (not shown) or from the NSP network.

In the usual cloud services orchestration model the user requests a well-defined resource, for example over the telephone, via a web-based portal, or via a function call. The orchestrator, after checking correctness, availability, and updating the billing system, allocates the resource, e.g., a VM on a particular CPU located in a particular rack in datacenter A. In addition, the required networking resources are allocated to the VM, e.g., an IP address, an Ethernet MAC address, and a VLAN tag. The VM is now started and consumes CPU power, memory, and disk space, as well as communications bandwidth between itself and other VMs on the same CPU, within the same rack, on other racks in the same datacenter, between datacenters, and between itself and the user. If it becomes necessary to move the VM from its allocated position to somewhere else (VM migration), the orchestrator needs to reallocate the required computational and communications resources. An example case is "cloudbursting" where a customer who finds himself temporarily with insufficient local resources reaches out to the cloud for supplementary ones [I-D.mcdysan-sdnp-cloudbursting-usecase]. A priori this requires allocating new addresses and rerouting all of the aforementioned traffic types, while maintaining continuous operation of the VM. When the user informs the CSP that it no longer requires the VM, the orchestrator needs to clear the routing entries, withdraw the communications resources, release storage and computational resources, and update the billing system.

The operations of the previous paragraph are all performed by the orchestrator, with possible cooperation with orchestrators from other datacenters. The needed routing information is advertised to the NSP via standard routing protocols, without taking into account possible effects on the NSP network. If, for example, the path in the NSP network to datacenter A degrades, while the path to datacenter B is performing well, this information is neither known by the orchestrator, nor is there a method for the orchestrator to take it into account. Instead, the NSP must find a way to reach datacenter A, even if this path is expensive, or of high latency, or problematic in some other way.

This predicament arises due to the orchestrator communicating (indirectly) with the user, but not with the NSP's OSS. In addition, although the CE may be capable of OAM functionality, fault and performance monitoring of the communications path through the NSP network are not employed. Finally, while the user can (indirectly) communicate with the orchestrator, there is no coordinated path to the NSP's OSS/BSS.

3. Optimized Cloud Access

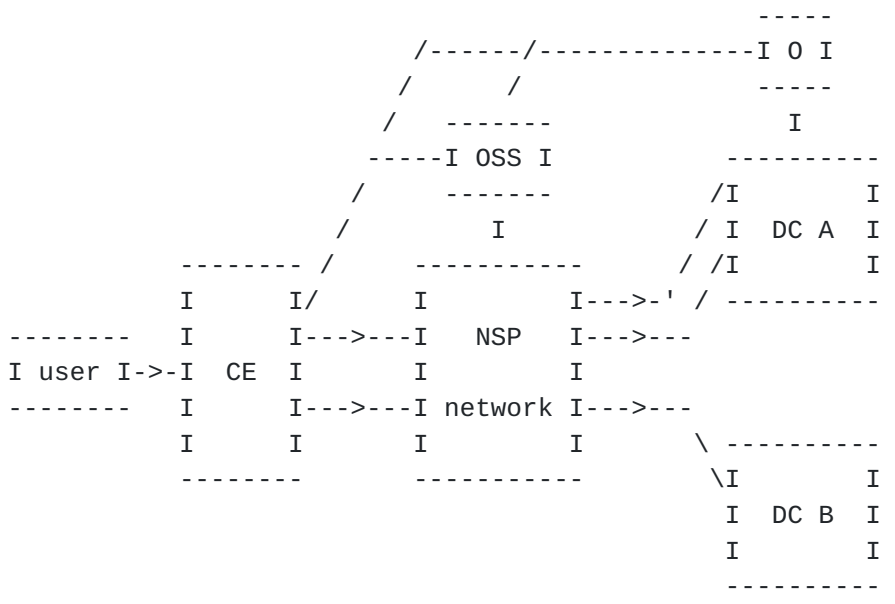

Figure 2: Cloud service with dual homing between a cloud-aware CE and NSP network, and coordination between CE, NSP OSS/BSS, and orchestrator

Figure 2. depicts two enhancements to the previous scenario. The trivial enhancement is the providing of dual-homing between the CE and the NSP network. This is a well-known and widely deployed

feature, which may be implemented regardless of the cloud services. We shall see that it acquires additional meaning in the context of the solution described below.

More significantly, Figure 2 depicts three new control communications channels. The CE device is now assumed to be cloud-aware, and may communicate directly with the NSP OSS/BSS, and with the CSP orchestrator. In addition, the latter two may communicate with each other. These control channels facilitate new capabilities, that may improve end-to-end QoE while optimizing operational cost. An alternative to a combined cloud/network CE is a separate "cloud demarcation device" placed behind the network CE.

Consider the provisioning of a new cloud service. With this new architecture the user's request is proxied by the cloud-aware CE to both the OSS/BSS and to the orchestrator. Before commissioning the service, the orchestrator initiates network testing between the datacenter and the CE, and with the NSP's assistance QoS parameters are determined for alternative paths to various relevant datacenters. The NSP and CSP (whether a single SP or two) can now jointly decide on placement of the VM in order to optimize the user's end-to-end Quality of Experience (QoE) while minimizing costs to both SPs. The best placement will necessitate the solution of a joint CSP + NSP optimization problem, while the latter minimization may only be reliable when a single SP provides networking and cloud resources. The joint optimization calculation will input the status of computational and storage resources at all relevant datacenters; as well as network delay, throughput, and packet loss to each datacenter. In some cases re-allocation of existing computational and networking resources may needed.

Similarly, the NSP OSS may trigger VM migration if network conditions degrade to the point where user QoE is no longer at the desired level, or may veto a CSP initiated VM migration when its effect would be too onerous on the NSP network.

The cloud-aware CE may be configured to periodically test path continuity and measure QoS parameters. The CE can then report that the estimated QoE drops under that specified in the SLA (or dangerously approaches it), in order to promote SLA assurance even when neither OSS nor orchestrator would otherwise know of the problem. Additionally, the cloud-aware CE may report workload changes detected by monitoring the number of active sessions (e.g., the number of "flows" or n-tuple pairs). The OSS and orchestrator can jointly perform root cause analysis and decide to trigger VM migration or network allocation changes or both. Finally, over-extended network segments may be identified, and pro-active VM migration and/or rerouting performed to better distribute the load.

When the CE is dual homed to the NSP network, the secondary link may be utilized in the conventional manner when the primary link fails, or may be selected as part of the overall optimization of QoE vs. cost. Load balancing over both links may also employed. The datacenters may also be connected to the network with multiple links (as depicted for DC A in Figure 2), enabling further connectivity optimization.

In addition, popular yet stationary content may be cached in the NSP network, and optimization may lead to the NSP network providing this content without the need to access the datacenter at all. In certain cases (e.g., catastrophic failure in the NSP network or of the connectivity between that network and the datacenter), the cloud-aware CE may choose to bypass the NSP network altogether and reach the datacenter over the public Internet (with consequent QoE reduction). In other cases, it may make sense to locally provide standalone resources at the cloud demarcation device itself.

4. Security Considerations

Perceived insecurity of the customer's data sent to the cloud or stored in a datacenter is perhaps the single most important factor impeding the wide adoption of cloud services. At present, the only solutions have been end-to-end authentication and confidentiality, with the high cost these place on user equipment. The cloud-aware CE may assume the responsibility for securing the cloud services from the edge of the customer's walled garden, all the way to the datacenter.

Isolation of CSP customers is addressed in [I-D.masum-chari-shc]. Security measures such as hiding of network topology, as well as on-the-fly inspection and modification of transactions are listed as requirements in [I-D.dalela-orchestration], while [I-D.dalela-sop] specifies encryption and authentication of orchestration protocol messages.

A further extension to the model is to explicitly include security levels as parameters of the QoE optimization process. This parameter may be relatively coarse-grained (for example, 1 for services which must be provided only over secure links, 0.5 for those for which access paths under direct control of the NSP is sufficient, 0 for general services that may run over out-of-footprint connections). Security may also take regulatory restrictions into account, such as limitations on database migration across national boundaries. Thus, the placement and movement of a VM will be accomplished based on full optimization of computational and storage resources; network delay, throughput, and packet loss; and security levels. For example, for an application for which the user can not afford denial of service

the joint optimizaton would need to find the needed resources as close as possible to the end user.

5. IANA Considerations

This document requires no IANA actions.

6. Acknowledgements

The work of Y(J)S, YG, DK, and KK was conducted under the aegis of ETICS (Economics and Technologies for Inter-Carrier Services), a European collaborative research project within the ICT theme of the 7th Framework Programme of the European Union that contributes to the objective "Network of the Future".

7. References

[I-D.bitar-datacenter-vpn-applicability]
Bitar, N., Balus, F., Lasserre, M., Henderickx, W., Sajassi, A., Fang, L., Ikejiri, Y., and M. Pisica, "Cloud Networking: Framework and VPN Applicability", draft-bitar-datacenter-vpn-applicability-02 (work in progress), May 2012.

[I-D.bitar-datacenter-vpn-applicability]
Bitar, N., Balus, F., Lasserre, M., Henderickx, W., Sajassi, A., Fang, L., Ikejiri, Y., and M. Pisica, "Cloud Networking: Framework and VPN Applicability", draft-bitar-datacenter-vpn-applicability-02 (work in progress), May 2012.

[I-D.dalela-orchestration]
Dalela, A. and M. Hammer, "Service Orchestration Protocol (SOP) Requirements", draft-dalela-orchestration-00 (work in progress), January 2012.

[I-D.dalela-sop]
Dalela, A. and M. Hammer, "Service Orchestration Protocol", draft-dalela-sop-00 (work in progress), January 2012.

[I-D.masum-chari-shc]
Hasan, M., Chari, A., Fahed, D., Tucker, L., Morrow, M., and M. Malyon, "A framework for controlling Multitenant Isolation, Connectivity and Reachability in a Hybrid Cloud Environment", draft-masum-chari-shc-00 (work in progress), February 2012.

[I-D.mcdysan-sdnp-cloudbursting-usecase]
McDysan, D., "Cloud Bursting Use Case", draft-mcdysan-sdnp-cloudbursting-usecase-00 (work in progress), October 2011.

[I-D.ietf-nvo3-framework]
Lasserre, M., Balus, F., Morin, T., Bitar, N., and Y. Rekhter, "Framework for DC Network Virtualization", draft-ietf-nvo3-framework-02 (work in progress), February 2013.

Authors' Addresses

Yaakov (Jonathan) Stein
RAD Data Communications
24 Raoul Wallenberg St., Bldg C
Tel Aviv 69719
Israel

Email: yaakov_s@rad.com

Yuri Gittik
RAD Data Communications
24 Raoul Wallenberg St., Bldg C
Tel Aviv 69719
Israel

Email: yuri_g@rad.com

Daniel Kofman
LINCS
23 Avenue d'Italie
Paris 75013
France

Email: daniel.kofman@telecom-paristech.fr

Konstantinos Katsaros
LINCS
23 Avenue d'Italie
Paris 75013
France

Email: katsaros@telecom-paristech.fr

Monique Morrow
Cisco Systems
Richtistrase 7
CH-8304 Wallisellen
Switzerland

Email: mmorrow@cisco.com

Luyuan Fang
Cisco Systems
300 Beaver Brook Road
Boxborough, MA 01719
US

Email: lufang@cisco.com

Wim Henderickx
Alcatel-Lucent
Copernicuslaan 50
2018 Antwerp
Belgium

Email: wim.henderickx@alcatel-lucent.com