

Physical Networking Migration

Or: migrate networks to new physical network

Feature Reference

[CLOUDSTACK-10024](#) - Jira project doesn't exist or you don't have permission to view it.

Feature Specifications

As a Cloud Owner

I want to migrate my existing guest networks to another physical network infrastructure

In order to leverage new network technologies, e.g. leverage SDN.

Hypervisors:

1. KVM
2. Vmware

Network Types

1. Isolated
2. VPC

Shared networks are not supported at this stage.

Architecture and Design

CloudStack supports multiple physical networks with guest traffic, but only when the guest traffic is given a tag. The network offering then also needs to specify the tag, in order to correlate it to a matching physical network.

CloudStack Core will be extended to support moving networks to another Physical Network, by adding API commands.

- Migrating a network will keep the VM running on the same host, so there are no issues for multinic VM's, when you would migrate all networks one by one.
- VR will be recreated. The new VR will get a new guest IP if it is no longer the gateway (when the target Network offering doesn't specify VR as provider of source nat). In that case, the new VR will get the first free IP i.s.o. Gateway IP.
- We will update the DB object with the details of the new implementation.
To be able to correctly cleanup the old implementation, we make a temporary copy of the network.

API Changes:

New API's:

- migrateNetwork
 - Parameter: networkId (UUID)
 - Parameter: networkOfferingId (UUID)
 - Parameter: resume (boolean)
- migrateVpc
 - Parameter: vpcId (UUID)
 - Parameter: vpcOfferingId (UUID)
 - Parameter: tierNetworkOfferings[] (to defines the new network offering for each tier)
 - networkId
 - networkOfferingId
 - Parameter: resume (boolean)

The networkOfferingId/vpcOfferingId specifies the target network/vpc offering, which as mentioned, uniquely determines the target physical network using its associated tag.

The API implementation will be made as re-entrant as possible.

E.g. if 2 network objects are found in the db for the specified networkId (UUID), it will

- By default indicate the error and fail
- Unless "resume" is explicitly set to True, then try to resume as much as possible from the corrupt entry state. Read more on "resume" in below section.

Manager Implementation

Migrate Network

During migration of a network the UUID of objects are swapped so that external system (like NuageVsp) will perceive the migration as a delete of the network on the old physical network followed by a reimplement on the new physical network. Because of this, orchestration systems can still use the same UUID's as before (the resulted objects in CloudStack will have the same UUID's except for the VR.).

- Make a temporary copy of the network before migration to keep the implementation details (moving along VR and nics).
- Implement original network on new Physical network using guest network guru, found by looking up the tag of the offering
- Implement services (Dhcp, static nat, source nat, ...)
 - Deploy VR if needed
- Migrate vnics to new physical network (move them back from the copy to the original network that is upgraded now)
 - Prepare Nic on Guru (add vport, vm interface, vm)
 - Prepare Nic on Element (apply acl rules, ...)
 - Send ReplugNicCommand to host
 - Remove nic on native element and guru
- Shutdown network services on old physical network (for this the copy object is used)
 - Undeploy VR if needed
- Shutdown network infrastructure on old physical network
- Delete original network

Migrate VPC

During migration of a network the UUID of objects are swapped so that external system (like NuageVsp) will perceive the migration as a delete of the network on the old physical network followed by a reimplement on the new physical network. Because of this, orchestration systems can still use the same UUID's as before (the resulted objects in CloudStack will have the same UUID's except for the VR.).

- Make copy of original VPC and implement VPC on new Physical network using guest network guru, found by looking up the tag of the network offerings
- Migrate Tiers to newly created VPC in the new Physical Network (same pattern as migrate network happens for every tier.)
 - Make copy of original network (moving along VR and nics)
 - Re-implement original network on new Physical network using guest network guru
 - Implement services (Dhcp, static nat, source nat, ...)
 1. Deploy VR if needed
 - Migrate vnics to new physical network
 1. Prepare Nic on Guru (add vport, vm interface, vm)
 2. Prepare Nic on Element (apply acl rules, ...)
 3. Send ReplugNicCommand to host
 4. Release nic on native element and guru
 - Shutdown network services on old physical network
 1. Undeploy VR if needed
 - Shutdown network infrastructure on old physical network
 - Delete original network
- Shutdown (original) vpc in old physical network
- Delete original vpc

Hypervisor Changes

Provide Agent command to change the bridge of a nic in one call.

Agent API

ReplugNicCommand

KVM

We need to allow both OVS and LinuxBridge on one host at the same time. This is not possible in the current implementation (only one VifProvider can be specified).

We changed the code to look for the correct VifDriver, by asking each one to see, if a bridge with the given name exists, and let the default create the bridge if it isn't found.

You can load both VifDrivers by having the following lines in /etc/cloudstack/agent/agent.properties

```
network.bridge.type=native
libvirt.vif.driver.Vpn=com.cloud.hypervisor.kvm.resource.OvsVifDriver
```

Resume support

Resume support is implemented for operators to have the ability to finish off a previously submitted migration command which halfway failed during its new network implementation. After correcting the cause of that failure, the operator can re-issue the same migration command, now with a resume parameter set to True.

The resume logic follows below table :

resume parameter value \ Network migration "state"	Good	Bad (but error resolved)
False (default)	Success	Fail (again)
True	Success	>> Success <<

Based on the resume parameter and the related field in the networks table, we "force" migration of a network even if the network offering of the network is already the same as the current network offering. In that case, we will use the network that is stored in the related field (DB table) as the already generated copy. Otherwise we follow the normal migration procedure.

The creation and migration of the network copy happens in a transaction block. The same happens for each nic that needs to be reassigned. The use of those transactions allow us to mark critical sections and never leave the DB in an unrecoverable state.

(TL;DR) Use cases

1: Migration of Isolated Network on Vmware 6.0, without VM's deployed

Given a VMware host setup, prepared for migration

And a physical network with VLAN encapsulation providing guest, management and storage networking

And a new physical network providing guest traffic tagged as "target"

And an isolated guest network in the VLAN physical network

And a network offering equivalent to the offering of the guest network tagged with "target",

When I migrate the network using the target network offering

Then the network is configured on the new physical network

This can be refined to:

- US-1-1 : Migrating from non-persistent (no VR deployed) to non-persistent offering
- US-1-2 : Migrating from persistent to non-persistent offering (with removal of VR)
- US-1-3 : Migrating from non-persistent (no VR deployed) to persistent offering
 - Without VR contained in destination offering
 - With VR contained in destination offering (we would need to deploy the VR)
- US-1-4 : Migrating from persistent to persistent offering
 - Without VR contained in destination offering (undeploy VR)
 - With VR contained in destination offering (redeploy VR)

Note: US-1-3 and US-1-4 need persistent isolated network offering support with Nuage (upstream PR is available)

2: Migration of Isolated Network on Vmware 6.0, with only stopped VM('s)

Given a VMware host setup

And a physical network with VLAN encapsulation providing guest, management and storage traffic

And a new physical network providing guest traffic tagged as "target"

And an isolated guest network in the VLAN physical network

And a network offering equivalent to the offering of the guest network tagged with "target",

And 2 stopped vms in the guest network

When I migrate the network using the target network offering

Then the network is configured on the new physical network

And the VMs and their interfaces are configured on the new physical network

This can be refined to :

- US-2-1 : Migrating from non-persistent Isolated Network offering to Nuage based offering without VR in offering or non-persistent Nuage based offeri
- US-2-2 : Migrating from non-persistent Isolated Network offering to persistent Nuage based offering with VR in offering (deploy VR)
- US-2-3 : Migrating from persistent Isolated Network offering to Nuage based offering without VR in offering or to non-persistent Nuage based offering (undeploy VR)
- US-2-4 : Migrating from persistent Isolated Network offering to persistent Nuage based offering with VR in offering (redeploy VR)

3: Migration of Isolated Network on Vmware 6.0, with running VM('s)

Given a VMware host setup

And a physical network with VLAN encapsulation providing guest, management and storage networking

And a new physical network with Guest Traffic tagged as "target"

And an isolated guest network in the VLAN physical network

And a network offering equivalent to the offering of the guest network with tag "target"

And 2 vms in the guest network

When I migrate the network using the target network offering

Then the network is configured on the new physical network

And the VMs and their interfaces are configured on the new physical network

And the VMs can reach each other

This can be refined to :

- US-3-1 : Migrating to an offering without VR providing any service (VR to be undeployed)
- US-3-2 : Migrating to an offering with VR providing services (VR to be redeployed)