DEA libvirt deployment prototype

This is an example of how to deploy a libvirt KVM setup with a DEA YAML file.

The file is created from an already deployed Fuel installation using the create_dea script and helper files which are to be present on the Fuel master and run from there.

The install is kicked off from the host by running deploy.sh and providing the ISO file to deploy and (optionally) an DEA file name as an argument. If the DEA file is omitted the example one will be used instead.

Pre-condition 1: The host needs to be Ubuntu 14.x

Pre-condition 2: Necessary packages installed by running sudo genesis/fuel/prototypes/libvirt/setup_vms/setup-vm-host.sh

Pre-condition 3: Example VM configuration deployed by running genesis/fuel/prototypes/libvirt/setup_vms/apply_setup.sh The VMs and networks to be setup are in genesis/fuel/prototypes/libvirt/examples: "vms" and "networks" sudo mkdir /mnt/images cd setup-vms sudo ./apply_setup.sh /mnt/images 50

In order to run the automated install, it's just a matter of running genesis/fuel/prototypes/libvirt/deploy.sh <isofile> [<deafile>] The deafile will be optional, if not specified the example one in genesis/fuel/prototypes/libvirt/examples/libvirt_dea.yaml will be used. sudo ./deploy.sh ~/ISO/opnfv-P0000.iso ~/DEPLOY/deploy/dea.yaml

Now either this will succeed (return code 0) or fail. I'll have a three hours safety catch to kill off things if something is hanging, may need to be adjusted for slow environments (see deploy.sh).

All the steps above should be run with sudo.

In principle the deploy.sh is assuming the example vm setup (one fuel, three controllers, two computes) and will always deploy with full HA and Ceilometer.

TODO: Copy also the deployment mode in my dea.yaml creation script genesis/fuel/prototypes/libvirt/create_dea/create_dea.sh so it's a real xerox of the running deploy.