Scripting And Integration with oVirt

Oved Ourfali, ovedo@redhat.com
Senior Software Engineer, Red-Hat

CloudOpen Europe - October 2013

Agenda

- Part 1 – REST-based APIs
  - Introduction
  - oVirt API
  - oVirt Shell (CLI) (Demo)
  - oVirt SDK
  - Deltacloud APIs
- Part 2 – Extension APIs
  - UI Plugin API (Demo)
  - Scheduling API
  - VDSM hooks

oVirt

Part 1
REST-based APIs

What can I do with it?

oVirt

- Infrastructure configuration
  - Host configuration and management
  - Network and storage configuration
- Virtual machine (VM) configuration and management
  - Networking for the Guest, Virtual disks, VM properties
- User management
- Daily maintenance and VM lifecycle management
- Advanced operations not available in the GUI
- And much more ...

Where would I use it?

oVirt

- Scripting and utilities
- A way to integrate with other software used in the organization
- Automation of tasks – provisioning and management
- Software development – of specific add-ons or applications
- Performance monitoring and reports

API methods

oVirt

REST

https://host:port/api/vms

Returns:

- XML/JSON/...

```
<vm id="aee0dbce-1591-44d4-9052-c2209b3e45b8"
href="/api/vms/aee0dbce-1591-44d4-9052-
c2209b3e45b8">
    <name>Austin</name><actions>
    <link rel="shutdown" href="/api/vms/aee0dbce-
1591-44d4-9052-c2209b3e45b8/shutdown"/>
     <link rel="start" href="/api/vms/aee0dbce-1591-
44d4-9052-c2209b3e45b8/start"/>
.....
```

SDK (Python/Java)

api.vms.list()

Returns:

- list of VM objects

Shell

list vms

Returns:

- Formatted text

```
id                    : 18df94a7-048f-4306-
9cfd-a74e8ea3b907
name                  : Boston
description           : Main service for
Boston
cluster-id            : 99408929-82cf-4dc7-
a532-9d998063fa95
cpu-topology-cores    : 2
cpu-topology-sockets  : 1
```

What method to use?

oVirt

- Depends on
  - Who you are and what is your role?
  - What you are trying to do?
  - What is the level of your expertise?
- Few examples
  - System Administrator who is managing small virtualized servers, using the Shell interface for a quick troubleshoot
  - Datacenter admin, running large virtualized environment, using the SDK to automate provisioning and managing hundreds of virtual machines
  - Software developer who integrates creating his own libraries over REST API

API Concepts

oVirt

- All APIs integrate through the oVirt engine
- All types of APIs are based on the web services interface
  - REST as the core
  - SDK on top of REST
  - Shell implemented on top the Python SDK
- Backward and forward compatibility
- Secure access
- Session-based access

oVirt API

HTTP methods in REST

oVirt

- GET

  Requests a representation of the specified resource. Requests using GET (and a few other HTTP methods) "SHOULD NOT have the significance of taking an action other than retrieval."

- POST

  Submits data to be processed to the identified resource. The data is included in the body of the request.

- PUT

  Uploads a representation of the specified resource

- DELETE

  Deletes the specified resource

Media types

- XML

```
<vms>
  <vm id=”xxx”>
    <name>yyy<name>
  </vm>
</vms>
```

- JavaScript Object Notation (JSON)

```
{
  “vms” : [
    “vm” : {
      “id” : ”xxx”,
      “name” : ”yyy” } ]
}
```

oVirt-API URI structure

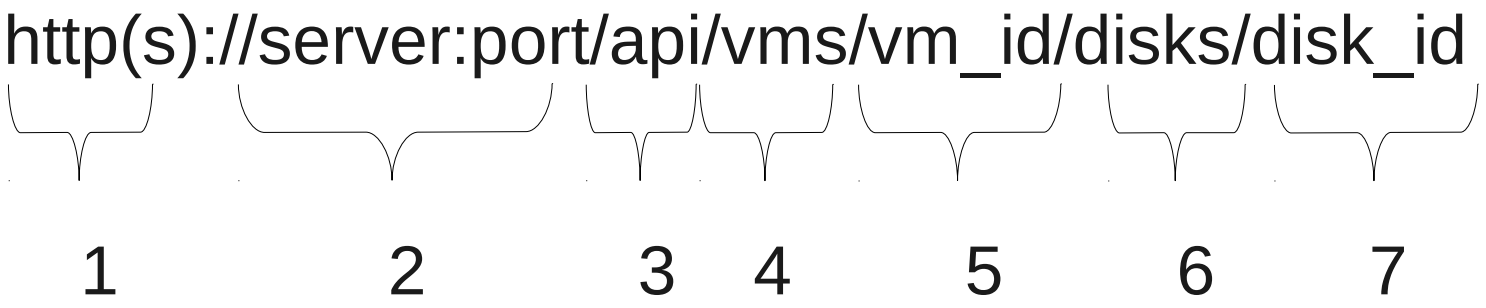

1. protocol
2. server endpoint details
3. entry point (base resource)
4. collection
5. resource
6. sub-collection
7. sub-resource

oVirt-API Resource Structure

```
<vm id="xxx" href="/api/vms/xxx">  identification details
    <name>vm1_iscsi</name>
    <status>DOWN</status>
    <memory>10737418240</memory>
    <cpu>
        <topology cores="1" sockets="1"/>
    </cpu>
    <start_time>2011-07-13T12:05:34.931Z</start_time>
    <creation_time>2011-05-31T16:47:51+03:00</creation_time>
    <actions>
        <link rel="start" href="/api/vms/xxx/start"/>
        <link rel="stop" href="/api/vms/xxx/stop"/>
    </actions>
    <link rel="disks" href="/api/vms/xxx/disks"/>
    <link rel="nics" href="/api/vms/xxx/nics"/>
    <cluster id="zzz" href="/api/clusters/zzz"/>
    <template id="yyy" href="/api/templates/yyy"/>
</vm>
```

Metadata

Resource details

Actions

Links to related resources

oVirt-API How-to (the methods)

- To list all collection resources, use GET

  GET http(s):/server:port/api/vms
- To retrieve specific resource, use GET

  GET http(s)://server:port/api/vms/{ID}
- To create a resource, use POST

  POST http(s)://server:port/api/vms

  <vm>...</vm>
- To update the resource, use PUT

  PUT http(s)://server:port/api/vms/{ID}

  <vm><name>new_name</name></vm>
- To remove the resource, use DELETE

  DELETE http(s)://server:port/api/vms/{ID}

Clients / Tools

- Any HTTP library/client can be used as a

  a client for oVirt-API
- Common used clients are
  - Firefox REST Client
  - Chrome REST-Client
  - Linux: curl / wget

  ...

Example - GET

GET http(s)://server:port/api/vms/{ID}

```
curl -v -u "user@domain:password" -H "Content-
type: application/xml" -X GET
http(s)://server:port/api/vms/{ID}
```

Example - CREATE

- Create VM

  POST http(s)://server:port/api/vms

```
<vm>
  <name>my_new_vm</name>
  <cluster id="xxx" />
  <template id="yyy" />
</vm>
```

- ```
  curl -v -u "user@domain:password"   -H "Content-type: application/xml"
     -d '<vm> <name>my_new_vm</name>
     <cluster><name>cluster_name</name></cluster>
     <template><name>template_name</name></template>
      </vm>' 'http(s)://server:port/api/vms'
  ```
  ```

Example - UPDATE

- Update

  PUT http(s)://server:port/api/vms/xxx
  <vm>
  <name>new_name</name>
  </vm>

- `echo "<vm><name>new_name</name></vm>" > /tmp/upload.xml`

- ```
  curl -v -u "user@domain:password"
    -H "Content-type: application/xml"
    -T /tmp/upload.xml
    'http(s)://server:port/api/vms/xxx'
  ```
  ```

Example - DELETE

- Delete

  DELETE http(s)://server:port/api/vms/xxx

- `curl -v -u "user@domain:password" -X DELETE http(s)://server:port/api/vms/xxx`

RSDL - RESTful Services Description Language

- Describes parameter constraints
- Easy way to understand
  - How to create the resource
  - What actions are available on a collection
  - What parameters to pass
    - Mandatory/optional/read-only
    - Type
    - Overloads

RSDL link description (http(s)://engine:port/api?rsdl)

```
<link href="/api/vms" rel="add">
    <request>
        <http_method>POST</http_method>
        <headers>
            <header>
                <name>Content-Type</name>
                <value>application/xml</value>
            </header>
            <header>
                <name>Expect</name>
                <value>201-created</value>
            </header>
        </headers>
        <body>
            <type>VM</type>
            <parameters_set>
                <parameter required="true">
                    <name>vm.name</name>
                    <type>xs:string</type>
                </parameter>
                <parameter required="true">
                    <name>vm.template.id|name</name>
                    <type>xs:string</type>
                </parameter>
                <parameter required="true">
                    <name>vm.cluster.id|name</name>
                    <type>xs:string</type>
                </parameter>
                ...
            </parameters_set>
        </body>
    </request>
    <response>
        <type>VM</type>
    </response>
</link>
```

Additional Functionality

oVirt

- User-level API
  - oVirt version 3.2 and above support user-level API
  - In order to use it, you need to pass the “filter” HTTP header, with a value of “true”
- Session support
  - We use cookies to allow using the REST API without having to re-authenticate on every request
  - See more details in http://wiki.ovirt.org/Features/RESTSessionManagement

oVirt Shell

oVirt Shell

- Concepts
  - Interactive shell
  - Use tab for auto-completion
  - Use arrow keys to browse previous commands
  - Ability to execute scripts
  - Pipe output into shell commands

oVirt Shell - Smart help

- help is dynamically created for each command
  - help <command> [arguments] [options]
- Help command
  - help add
- Help command argument
  - help add storagedomain
- Or if it's for operation on a subcollection context
  - help add storagedomain --datacenter-identifier yyy

oVirt Shell - Connect

- When running the ovirt-shell, you'll get a disconnected prompt
- To connect as an administrator, run
  - [oVirt shell (disconnected)]# connect --url "http://server:port/api" --user "user@domain" --password "password"

  and as a user run
  - [oVirt shell (disconnected)]# connect --url "http://server:port/api" **--filter** --user "user@domain" --password "password"

  then you should get the “connected” prompt
  - [oVirt shell (connected)]#

oVirt Shell - Smart auto-completion

available commands

```
[oVirt shell (connected)]# <TAB><TAB>
EOF         clear       console     remove      echo        help        ping        show        update
action      connect     add      disconnect  exit        list        shell       status
```

available options for specific command

```
[oVirt shell (connected)]# add <TAB><TAB>
cdrom          datacenter     group          network        permission     role           storagedomain  template       vm
cluster        disk           host           nic            permit         snapshot       tag            user           vmpool
```

available options for command on specific resource

```
[oVirt shell (connected)]# add vm <TAB><TAB>
cluster-id                          display-type                        os-boot-dev                 template-id
cluster-name                        domain-name                         os-cmdline                    template-name
cpu-topology-cores                  high_availability-enabled           os-initRd                     timezone
cpu-topology-sockets                high_availability-priority          os-kernel                     type
custom_properties-custom_property   memory                              os-type                       usb-enabled
description                         name                                placement_policy-affinity
display-monitors                    origin                              stateless
```

available options for command on specific sub-resource

```
[oVirt shell (connected)]# add nic --vm-identifier xxx <TAB><TAB>
interface     mac-address   name          network-id    network-name
```

oVirt Shell – Querying for Resources

- list <resource>
  - Provides a list of these resources, minimal details
- list <resource> --show-all
  - Provides a list of these resources, full details
- list <resource> --kwargs "param=val"
  - list resources using client side filtering
  - list vms --kwargs name=p*
- list <resource> --query "param=val"
  - list resources using oVirt query engine filtering
- Etc...

oVirt Shell – Resource Manipulation

- add
- update
- remove

oVirt Shell – Actions

- Action on resource
  - action <resource> <name> <action>
- Action on sub-resource
  - action <sub-resource> <name> <action> --<resource>-identifier <name>

oVirt Shell – Help on Actions

```
[oVirt shell (connected)]# help action vm demo start

OBJECT IDENTIFIERS

  Some objects can only exist inside other objects. For example, a disk
  can only exist in the content of a virtual machine. In this case, one
  or more object identifiers needs to be provided to identify the
  containing object.

  An object identifier is an option of the form '--<type>id <id>'. This
  would identify an object with type <type> and id <id>. See the
  examples section below for a few examples.

ATTRIBUTE OPTIONS

  The following attribute options are understood. Note: this lists all
  available attribute options for action 'start'.

    [--vm-os-initRd: string]
    [--vm-domain-name: string]
    [--host-id|name: string]
    [--async: boolean]
    [--vm-display-type: string]
    [--vm-os-kernel: string]
    [--grace_period-expiry: long]
    [--display-type: string]
    [--vm-stateless: boolean]
    [--vm-os-cmdline: string]
    [--vm-domain-user-username: string]
    [--pause: boolean]
    [--vm-os-boot-dev: string]
    [--vm-domain-user-password: string]

EXAMPLES

  - This example starts a vm named "myvm" with few optional parameters:

    $ action vm iscsi_desktop start --vm-display-type vnc --vm-cpu-topology-sockets 2 --vm-cpu-topology-cores 2
```

oVirt Shell - Actions

```
(oVirt cli) > show vm new_vm

id          : 62004129-a806-4e48-9b39-f6a54c97cba6
name        : new_vm
status      : down
memory      : 1024
os          : unassigned
display     : spice
monitors    : 1
stateless   : False
template    : 94f5ad88-a12a-4f48-af9f-f2ba28b7285b
cluster     : 99408929-82cf-4dc7-a532-9d998063fa95

(oVirt cli) > action vm new_vm start

status: complete

(oVirt cli) > show vm new_vm

id          : 62004129-a806-4e48-9b39-f6a54c97cba6
name        : new_vm
status      : powering_up
memory      : 1024
os          : unassigned
display     : spice
monitors    : 1
stateless   : False
template    : 94f5ad88-a12a-4f48-af9f-f2ba28b7285b
cluster     : 99408929-82cf-4dc7-a532-9d998063fa95
```

oVirt Shell - Examples

- Action with parameters
  - action host Atlantic install --root_password 123456
- Create and update
  - add vm --name BaseRHEL6 --memory 1073741824 --template-name Blank --cluster-name Default
  - update vm BaseRHEL6 --os-type rhel_6x64
  - add nic --vm-identifier BaseRHEL6 --name nic1 –network-name ovirtmgmt

oVirt Shell - Examples

- Create and attach Disks
  - create disk --size 10737418240 --wipe_after_delete False --format raw --sparse True --interface virtio --storage_domains-storage_domain "storage_domain.name=myImageStore" --bootable True
  - create disk --id bc2454d0-2539-41c3-85c2-91cfbdfbede3 --vm-identifier BaseRHEL6
- Or
  - create disk --size 10737418240 --wipe_after_delete False --format raw --sparse True --interface virtio --storage_domains-storage_domain "storage_domain.name=myImageStore" --vm-identifier BaseRHEL6

oVirt Shell - Examples

- Create Template
  - create template --vm-name BaseRHEL6 --name BaseRHEL6
- Create VM from Template
  - create vm --template-name BaseRHEL6 --name myRHEL --cluster-name Default
- Remove VM
  - remove vm my_new_vm

oVirt

oVirt SDK

oVirt SDK

- Mainly used for integration or advanced automation
- Object oriented
- Current bindings
  - Java
  - Python
  - Libgovirt (GObject wrapper for the oVirt REST API)
  - Rbovirt – ruby binding for the oVirt REST API
- Slides will demonstrate the python SDK
- Java - http://www.ovirt.org/Java-sdk
- Libgovirt - https://github.com/GNOME/libgovirt
- Rbovirt - https://github.com/abenari/rbovirt

oVirt SDK

- Concepts
  - Complete protocol abstraction
  - Full compliance with the oVirt API architecture
  - Auto-completion
  - Self descriptive, intuitive and easy to use
  - Auto-generated

oVirt SDK - Examples

- Creating the proxy

```
#create proxy
api = API(url='http://localhost:8080', username='user@domain', password='password')
```

Additional options:

filter=True for user-level API (default is False – admin API)
persistent_auth=True for using API REST sessions (default is False)

- Listing all collections

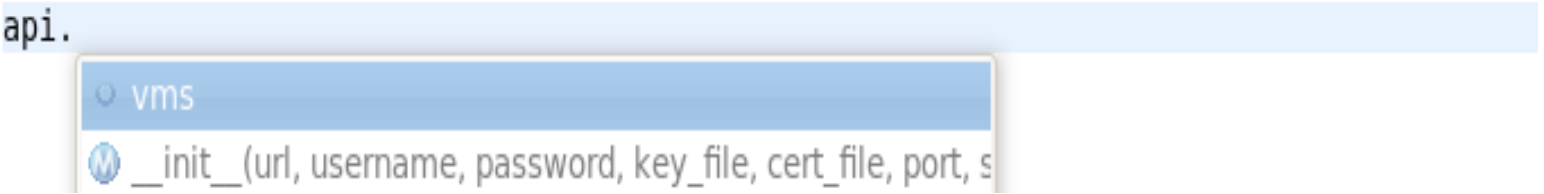

- Listing collection's methods.

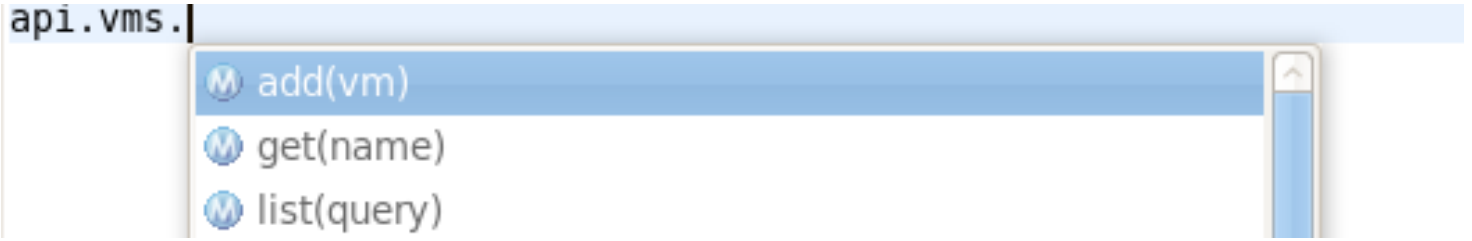

oVirt SDK - Examples

- Querying collection with the search engine.

- Querying collection by custom constraint.

```
#list by query
vms = api.vms.list(query = 'name=python_vm

#search vms by property constraint
vms = api.vms.list(memory=1073741824)
```

- Querying collection for specific resource.

```
#get by constraints
vm = api.vms.get(id = '02f0f4a4-9738-4731-83c4-293f3f734782')
```

- Accessing resource methods and properties.

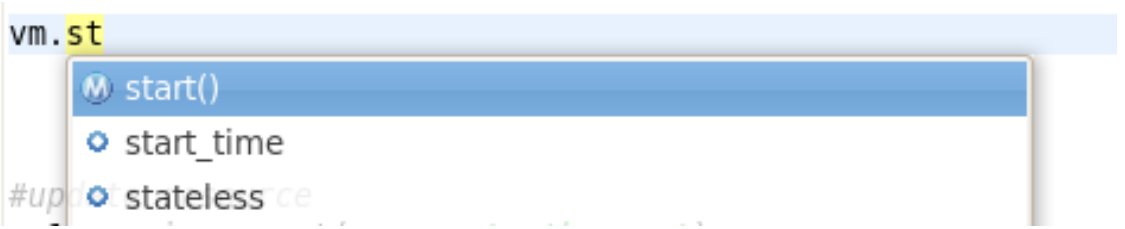

oVirt SDK - Examples

- Accessing resource properties and sub-collections

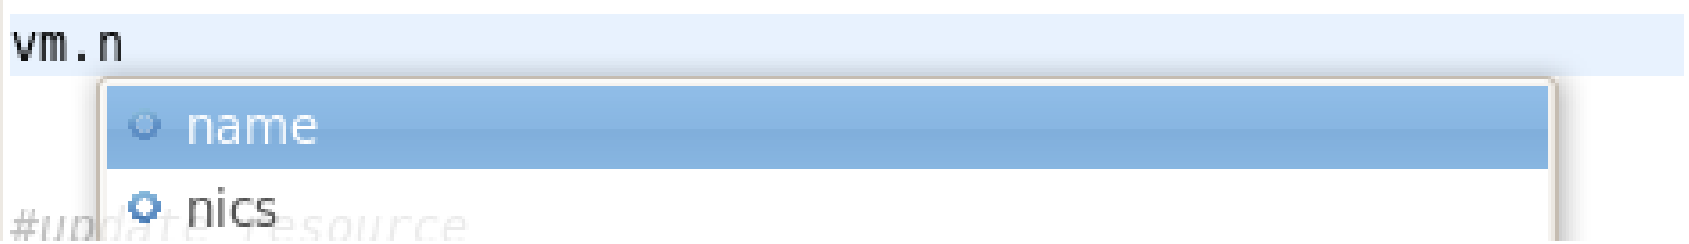

- Accessing sub-collection methods

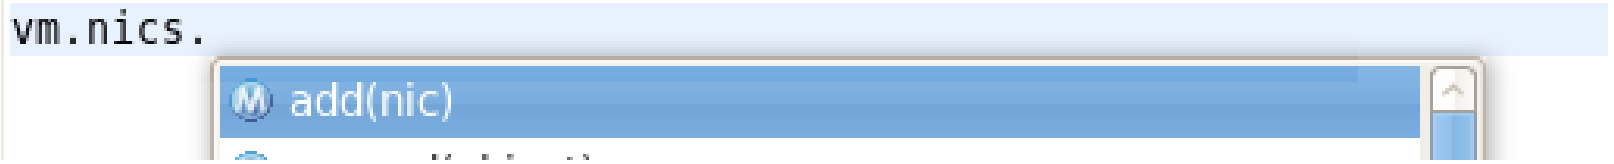

oVirt SDK - Examples

- Creating a VM resource

```
cluster = api.clusters.get(name='Default_nfs')
template = api.templates.get(name='nfs_desktop_tmpl')
param = params.VM(name='pythond_sdk_poc2', cluster=cluster, template=template, memory=1073741824)
vm6 = api.vms.add(param)
```

- Creating a Disk resource

```
param = params.StorageDomains(storage_domain=[api.storagedomains.get('iscsi_data')])
api.disks.add(params.Disk(name='new_disk', provisioned_size=10,
            size=1, status=None, interface='virtio',
            format='cow', sparse=True, bootable=False,
            storage_domains=param))
```

- Attach a disk to a VM, and activate it

```
vm.disks.add(params.Disk(id = disk_id, active = True))
```

oVirt SDK - Examples

- Querying sub-collection by custom constraint.

```
nics = vm.nics.list(interface='e1000')
```

- Retrieving sub-collection resource.

```
nic = vm.nics.get(name='eth0')
```

- Accessing sub-collection resource properties and methods.

```
#get sub resource
nic = vm6.nics.get(name='eth0')

nic.u
```

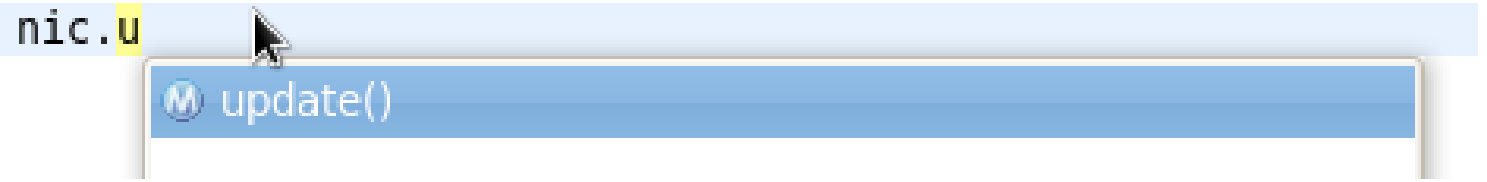

oVirt

Deltacloud

Deltacloud

- Open source Apache project
- Abstracts the differences between cloud providers
- Supports Deltacloud, EC2 and CIMI APIs
- Supports many cloud providers
  - EC2
  - oVirt
  - Eucalyptus
  - OpenNebula
  - OpenStack
  - .....

Motivation

oVirt

- Heterogeneous cloud and virtualization environment
- Existing software working with common cloud APIs like EC2
  - Synaps – CloudWatch implementation over EC2 API
  - Heat (Openstack project)
    - CloudFormation and CloudWatch support
  - Automated scripts
  - ....

Cloud APIs

oVirt

- DMTF CIMI API
  - Cloud Infrastructure Management Interface
  - http://dmtf.org/standards/cloud
  - Version 1.0.1 was published in October 2012
  - Still new API, but aims to be the cloud standard API
- EC2 API – Amazon Elastic Cloud API
- Deltacloud API

Deltacloud

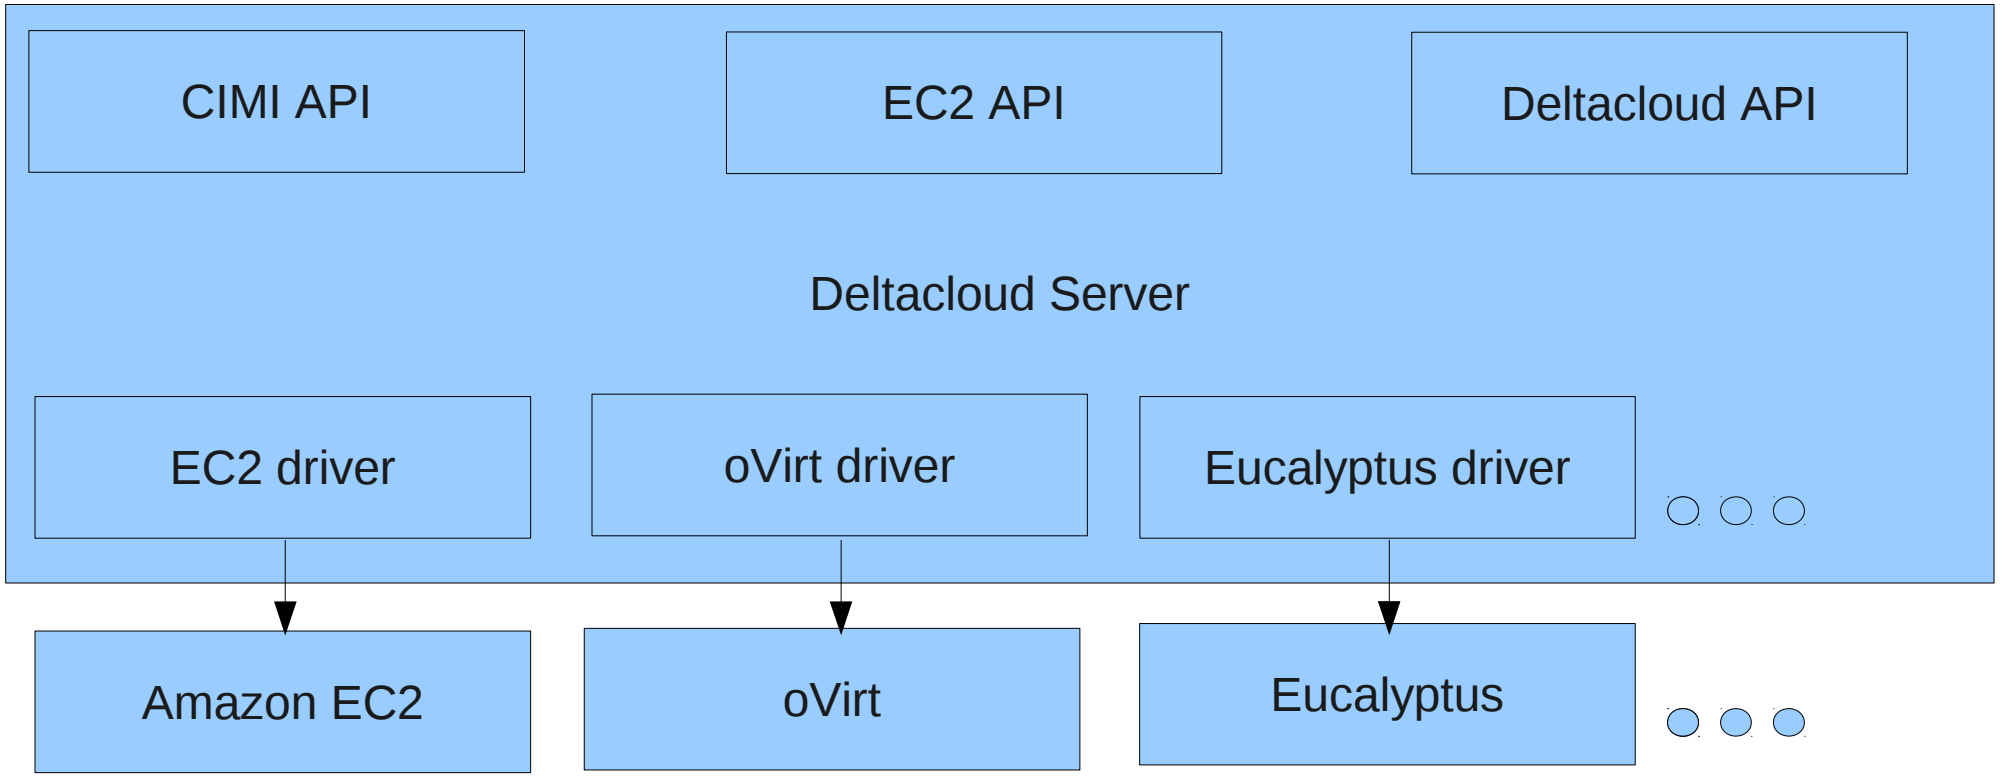
CIMI API
EC2 API
Deltacloud API
Deltacloud Server
EC2 driver
oVirt driver
Eucalyptus driver
Amazon EC2
oVirt
Eucalyptus

Deltacloud Links and Further Reading

- Deltacloud
  - Developer mailing list dev@deltacloud.apache.org
  - http://deltacloud.apache.org
  - #deltacloud on Freenode
- My blog with some useful examples on top of oVirt
  - http://ovedou.blogspot.com

oVirt

Part 2
Extension APIs

Agenda

- Part 1 – REST-based APIs
  - Introduction
  - oVirt API
  - oVirt Shell (CLI) (Demo)
  - oVirt SDK
  - Deltacloud APIs
- **Part 2 – Extension APIs**
  - **UI Plugin API (Demo)**
  - **Scheduling API**
  - **VDSM hooks**

oVirt

UI Plugins

Web Admin user interface

- Extend oVirt Web Admin user interface

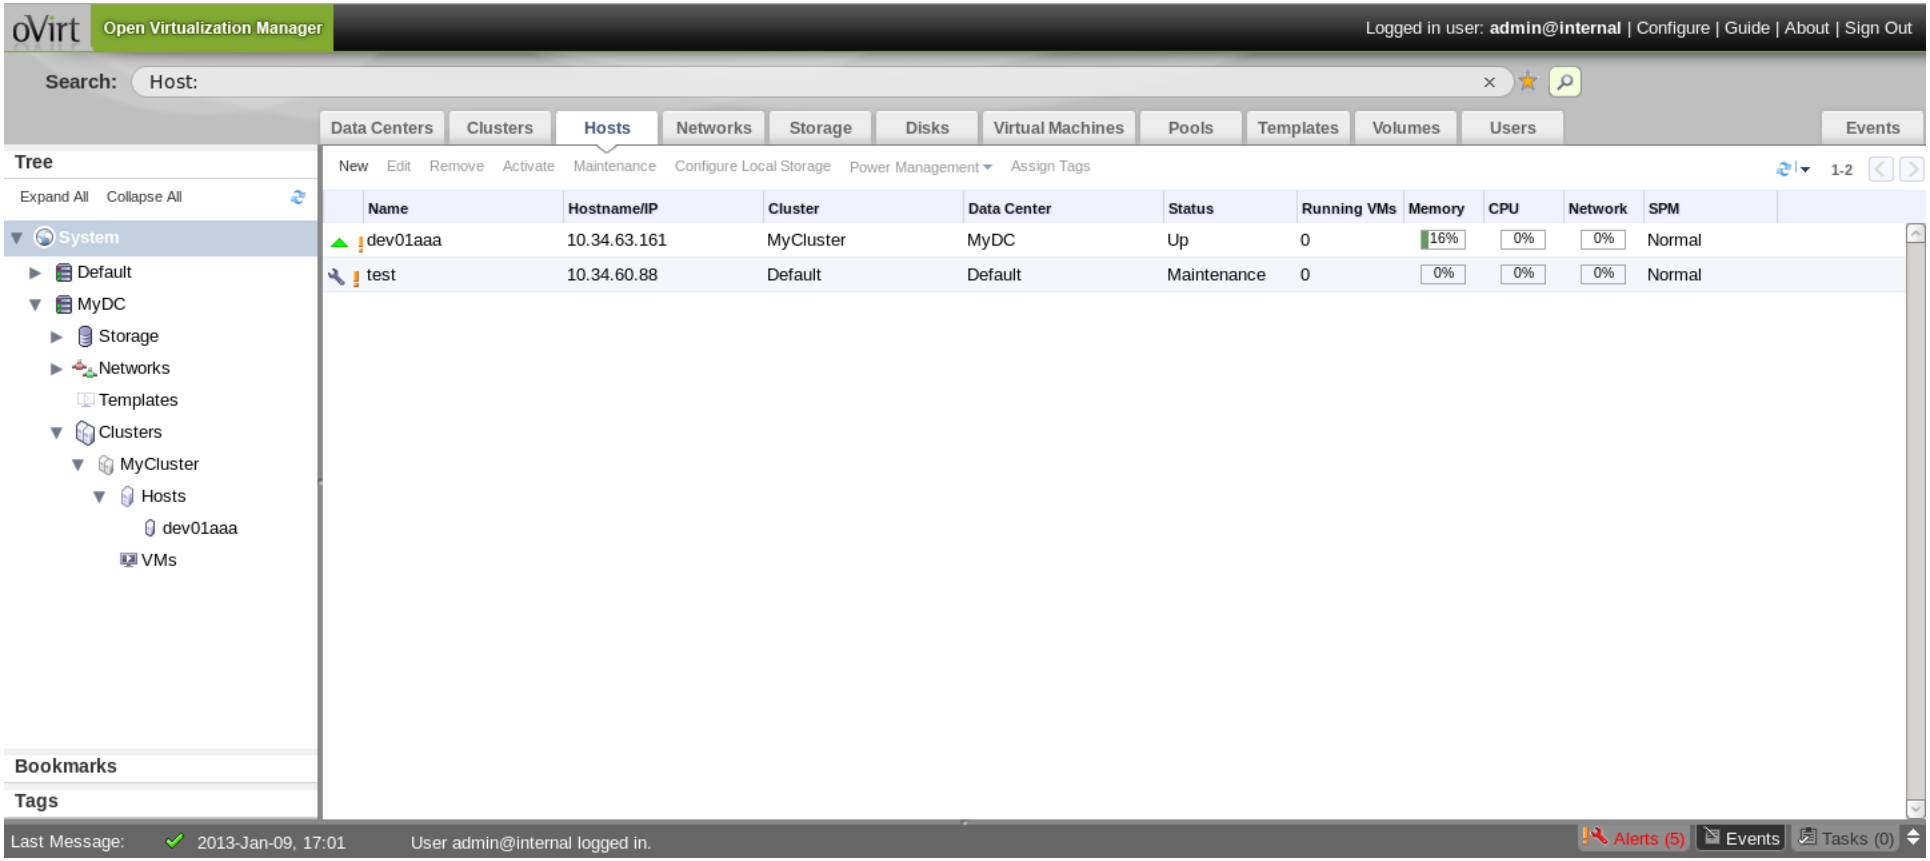

Web Admin user interface

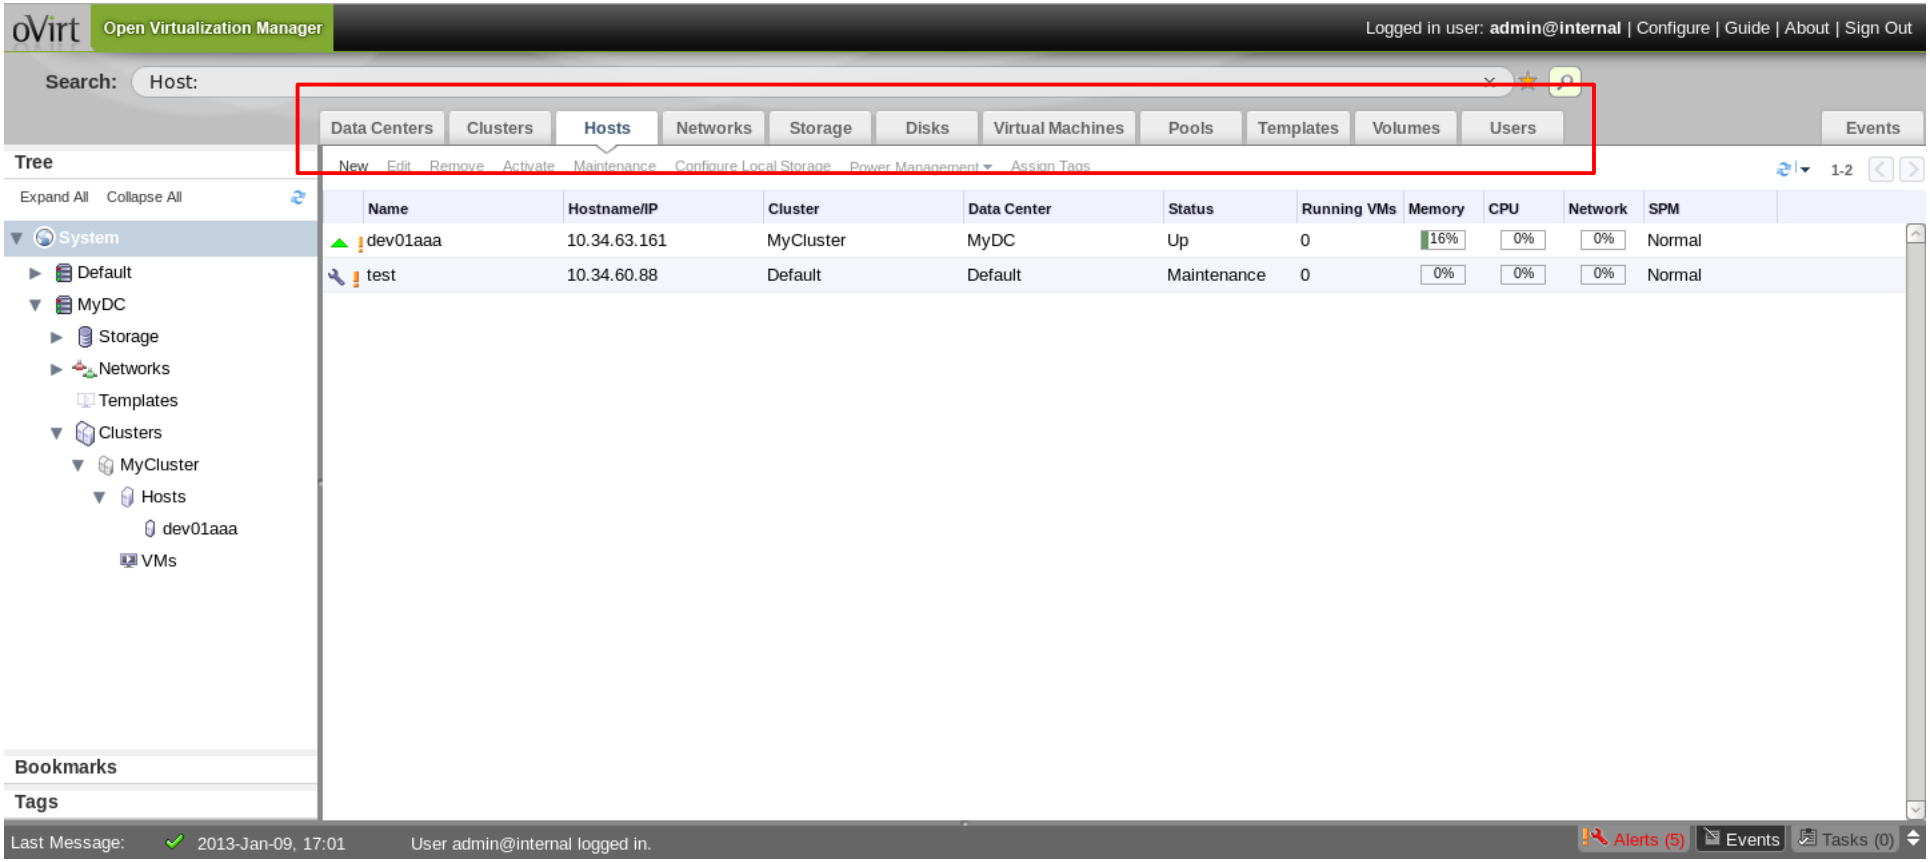

Web Admin user interface

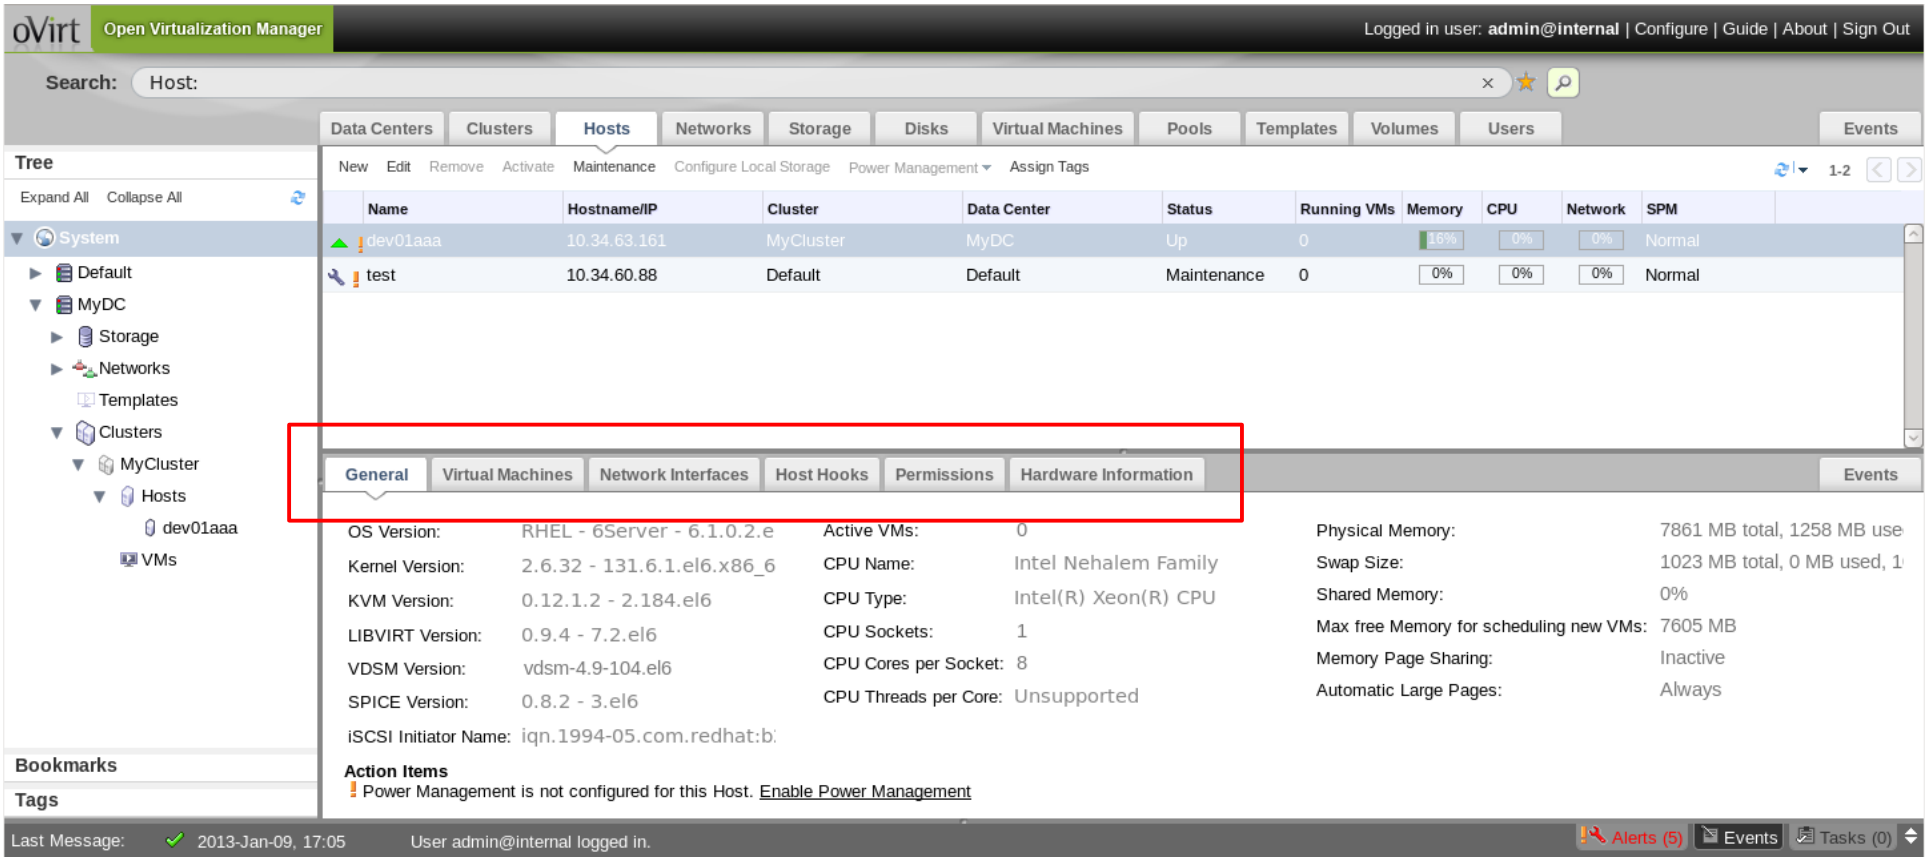

Web Admin user interface

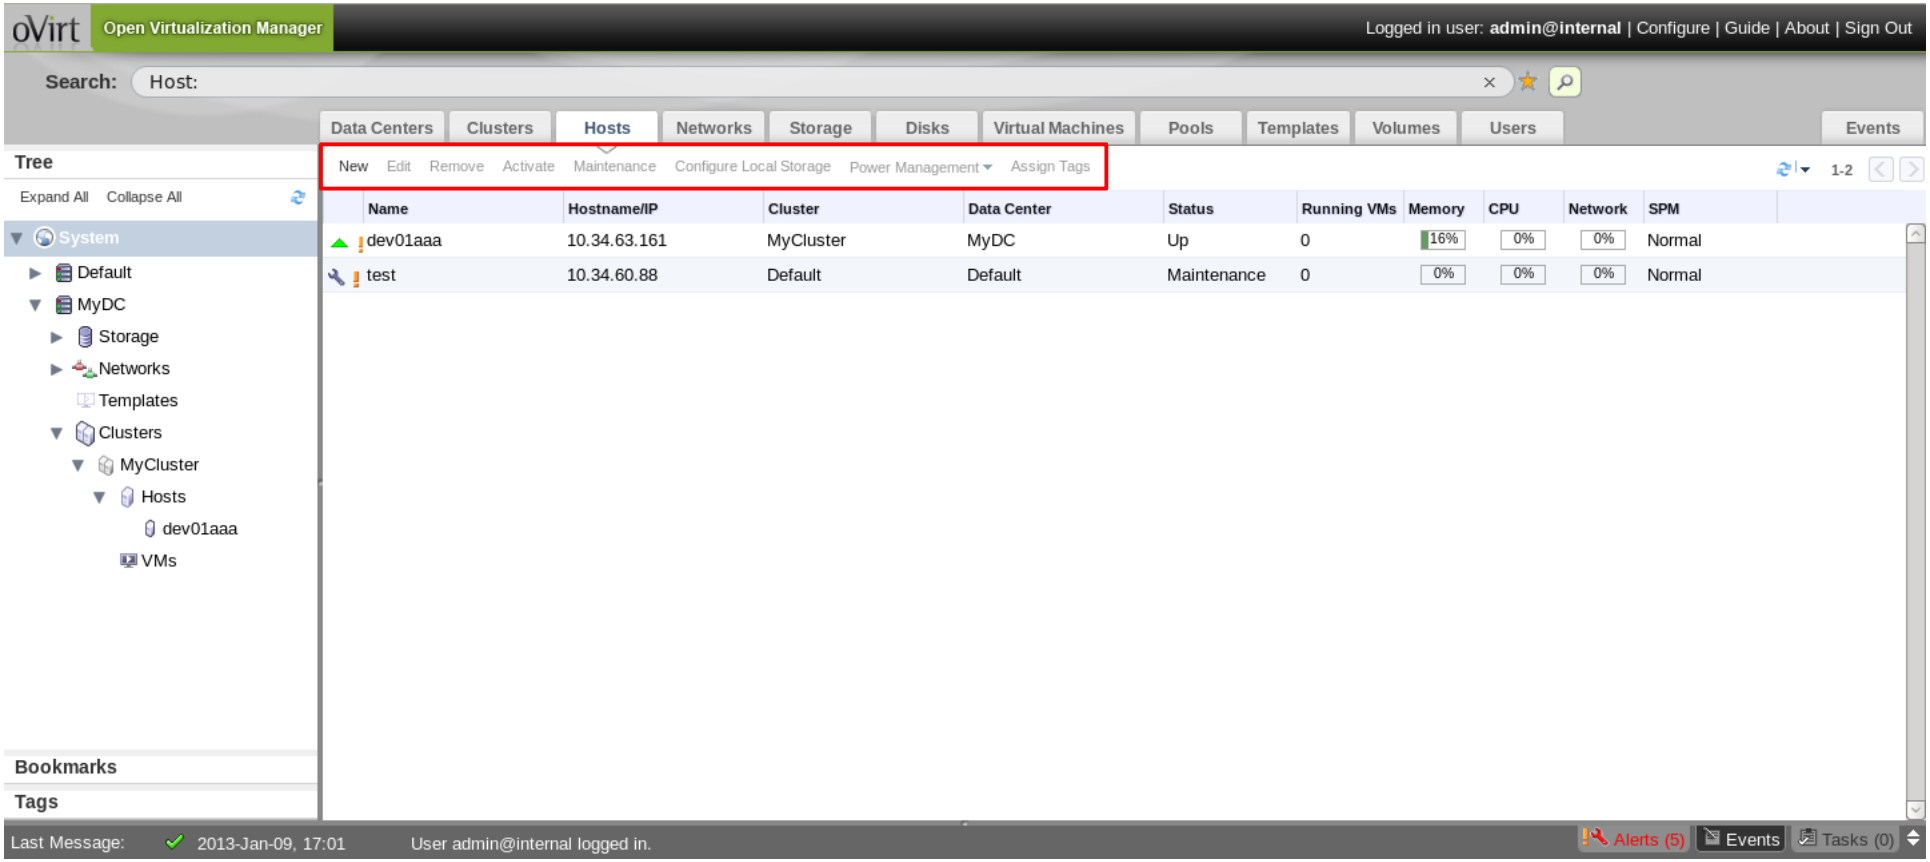

Web Admin user interface

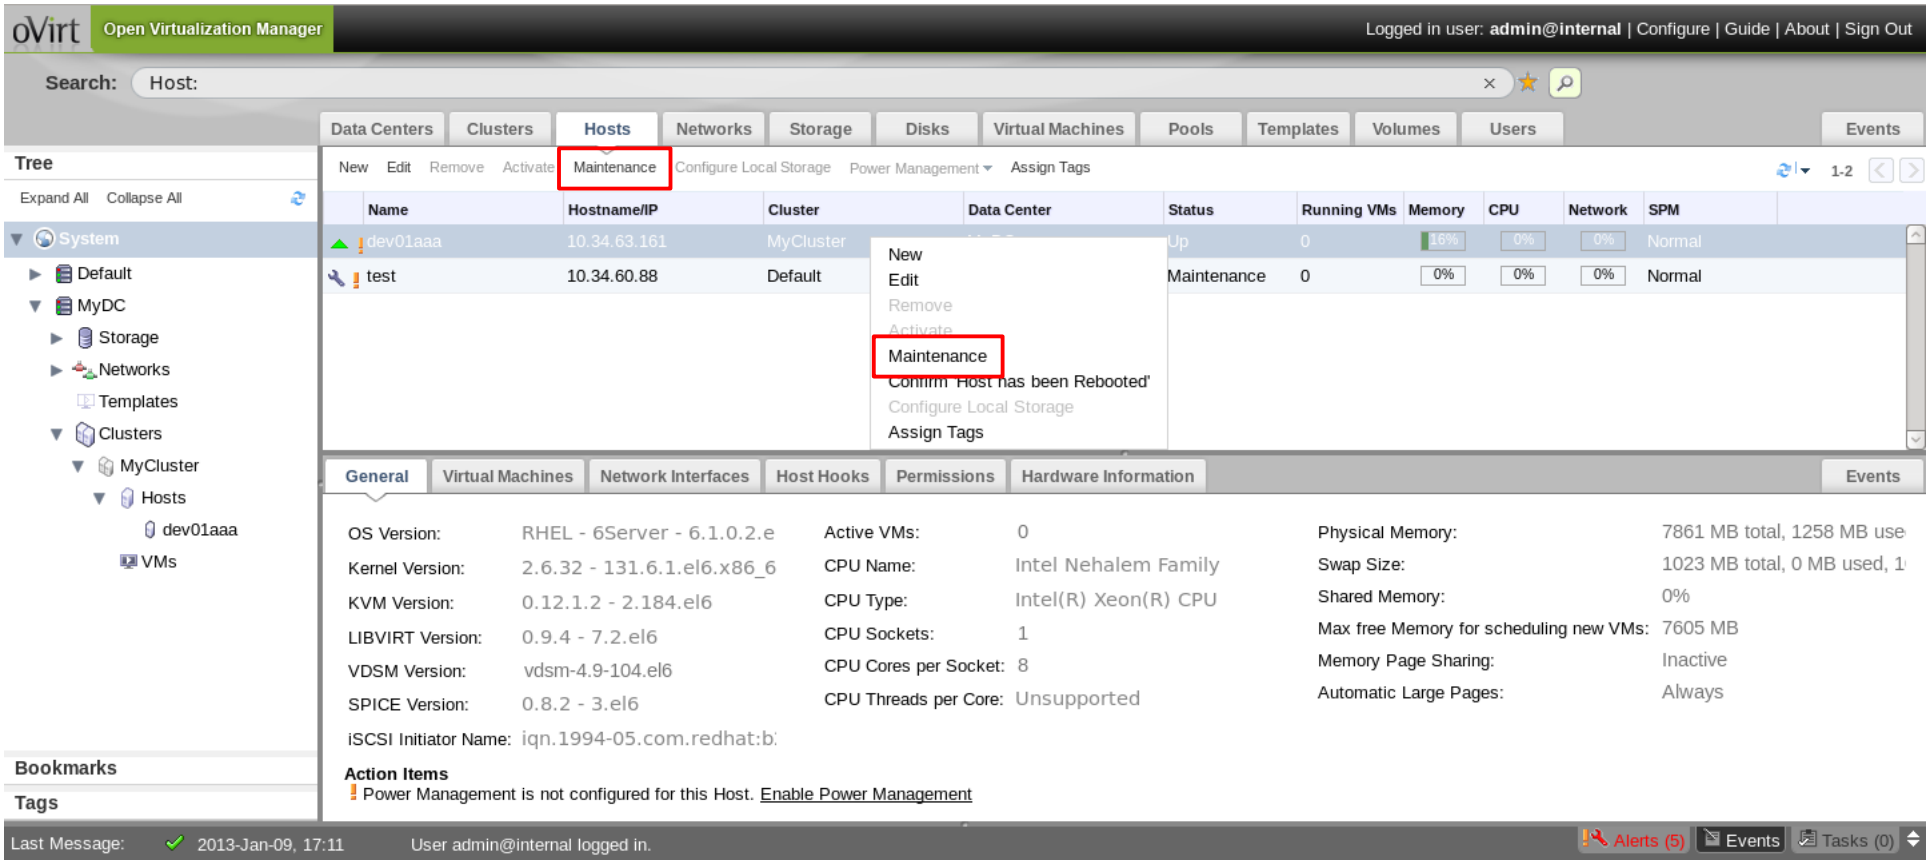

What's currently possible

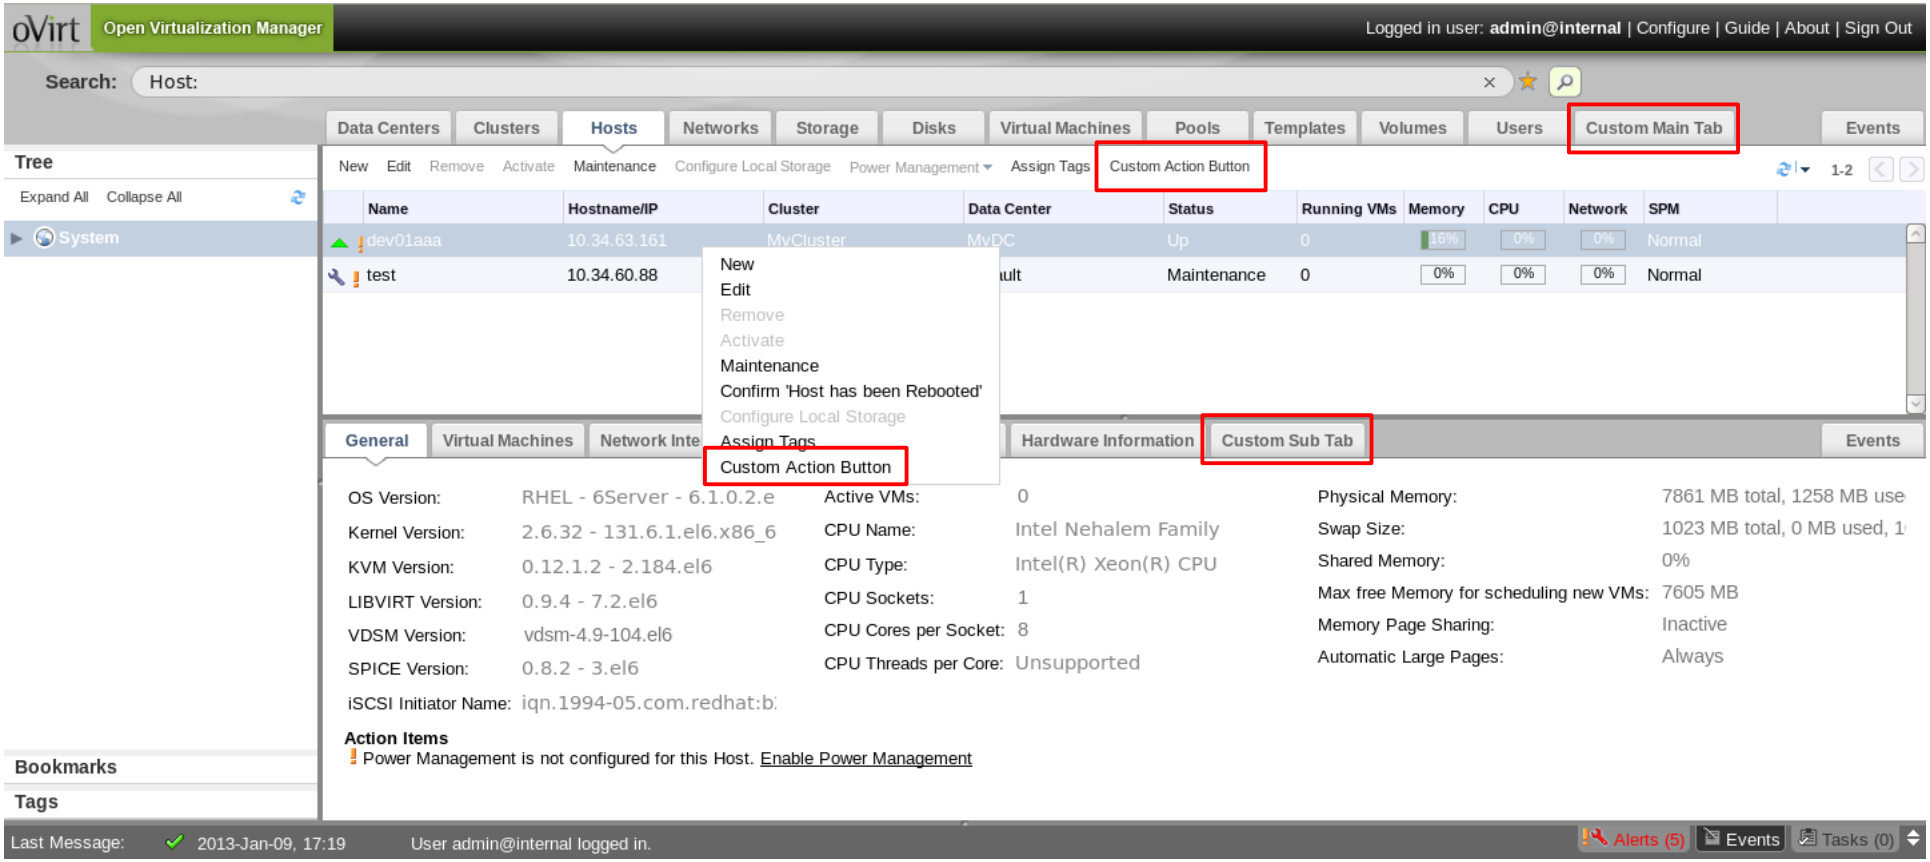

Writing plugins

```
<!DOCTYPE html>
<html>
<head>
    <!-- Fetch additional resources if necessary -->
    <script type="text/javascript" src="jquery-min.js"></script>

    <!-- Actual plugin code -->
    <script>
        // Access plugin API from iframe context
        var api = parent.pluginApi('myPlugin');

        // Register plugin event handler functions
        api.register({
            UiInit: function() {
                api.addMainTab('Foo Tab', 'foo-tab', 'http://foo.com/');
            }
        });

        // Tell plugin infrastructure that we are ready
        api.ready();
    </script>

</head>
<body> <!-- HTML body is intentionally empty --> </body>
</html>
```

UI plugin basics

oVirt

- Plugin host page
  - Hosts actual plugin code (JavaScript)
    `/usr/share/ovirt-engine/ui-plugins/<resourcePath>/<hostPage>.html`
- Plugin descriptor
  - Meta-data + default configuration
    `/usr/share/ovirt-engine/ui-plugins/<descriptorName>.json`
- Plugin user configuration
  - Override default configuration, tweak runtime behavior
    `/etc/ovirt-engine/ui-plugins/<descriptorName>-config.json`

Supported API functions

- `addMainTab(label, historyToken, contentUrl)`
- `addSubTab(entityTypeName, label, historyToken, contentUrl)`
- `setTabContentUrl(historyToken, contentUrl)`
- `setTabAccessible(historyToken, tabAccessible)`

```
String
Boolean
Number
Object
```

Supported API functions

- addMainTabActionButton(entityTypeName, label, buttonInterface, options)
  - Can add the button at the toolbar, context menu, or both
- addSubTabActionButton(mainTabEntityName, subTabEntityName, label, buttonInterface, options)
- showDialog(title, contentUrl, width, height)
- loginUserName()
- loginUserId()

```
String
Boolean
Number
Object
```

Supported API events

- UiInit
- {entity}SelectionChange(selectedItems[])
- UserLogin(fullUserName, userId)
- UserLogout()
- RestApiSessionAcquired(sessionId)
  - Using this one you can do anything I showed earlier
  - You can also use the new external events feature to get a more native integration
- MessageReceived (allows Plugin HTML to interact with the UI plugin)

```
String
Boolean
Number
Object
```

Plugin descriptor

- Meta-data + default configuration
  `/usr/share/ovirt-engine/ui-plugins/<descriptorName>.json`

```
{

    // A name that uniquely identifies the plugin (required)
    "name": "foo",

    // URL of plugin host page that invokes the plugin code (required)
    "url": "/webadmin/webadmin/plugin/foo/start.html",

    // Default configuration object associated with the plugin (optional)
    "config": { "band": "ZZ Top", "classic": true, "score": 10 },

    // Path to plugin static resources (optional)
    // Used when serving plugin files through PluginResourceServlet
    // This path is relative to /usr/share/ovirt-engine/ui-plugins
    "resourcePath": "foo-files"

}
```

Plugin user configuration

- Override default configuration, tweak runtime behavior
  `/etc/ovirt-engine/ui-plugins/<descriptorName>-config.json`

```
{

    // Custom configuration object associated with the plugin (optional)
    // This overrides the default plugin descriptor configuration, if any
    "config": { "band": "AC/DC" },

    // Whether the plugin should be loaded on WebAdmin startup (optional)
    // Default value is 'true'
    "enabled": true,

    // Relative order in which the plugin will be loaded (optional)
    // Default value is Integer.MAX_VALUE (lowest order)
    "order": 0

}
```

Runtime plugin configuration

- Merge user configuration (if any)
  on top of default configuration (if any)

```
{ "band": "ZZ Top", "classic": true, "score": 10 }
```

+

```
{ "band": "AC/DC" }
```

=

```
{ "band": "AC/DC", "classic": true, "score": 10 }
```

Main steps in plugin development

oVirt

(1) Write plugin descriptor

(2) Write plugin host page

(3) See plugin in action

UI Plugin - Live Demo

http://www.ovirt.org/Features/UIPlugins
http://ovedou.blogspot.co.il/
http://www.ovirt.org/Features/UIPlugins#UI_Plugins_Cras
git://gerrit.ovirt.org/samples-uiplugins

Plugin Initialization

oVirt

```
var init = function() {
    // Add new action button to Data Center main tab
    api.addMainTabActionButton('DataCenter', 'Protect DataCenter from Alien Invasion'
            // Callbacks and extra options for addMainTabActionButton function
            {
        isEnabled: function() {
            // Enable button only when selecting single Data Center
            return arguments.length == 1;
        },
        onClick: function() {
                                // Reset victory flag and open new game dialog
            victory = false;
            openDialog();
        },
        location: 'OnlyFromContext' // Make button available only from context menu
    });

    // Add new sub tab under Data Center main tab
    api.addSubTab('DataCenter', 'Space Shooter Score', subTabToken, subTabUrl,
                  {
                          alignRight: true
                  });
};
```

Opening a Dialog

oVirt

```
var openDialog = function() {
    // Show modal dialog with actual game content
    api.showDialog(dataCenter.name + ' under attack',
        dialogToken, dialogUrl, '340px', '560px',
        // Extra options for showDialog function
        {
            closeIconVisible: false,
            closeOnEscKey: false,
            resizeEnabled: true,
            buttons:
            [
                {
                    label: 'Cheat',
                    onClick: cheatGame
                },
                {
                    label: 'Get me outta here',
                    onClick: closeDialog
                }
            ]
        });
};
```

Event Registration

oVirt

```
// Register event handler functions that serve as callbacks on specific events
api.register({
    // Plugin init callback, ideal for one-time UI extensions
    UiInit: function() {
        if (browserRocks()) {
            init();
        }
    },
    // Data Center main tab item selection callback, useful for keeping track
    // of currently selected Data Center entity in the corresponding main tab
    DataCenterSelectionChange: function() {
        if (arguments.length == 1) {
            // If there is single entity selected, remember it
            // and make sure the custom sub tab is visible
            dataCenter = arguments[0];
            api.setTabAccessible(subTabToken, true);
        } else {
            // Otherwise, forget the entity
                            // and hide the custom sub tab
            dataCenter = null;
            api.setTabAccessible(subTabToken, false);
        }
    },
```

Event Registration – Cross-Window Communication

oVirt

```
// HTML5 message event callback, ideal for cross-window communication that
// works across different origins, i.e. unconstrained by Same-Origin Policy
MessageReceived: function(data, sourceWindow) {
    // If we get here, we already passed allowed source origin check
    switch (data) {
        // Game content HTML just loaded in the dialog
        case 'GameInit':
            gameContentWindow = sourceWindow;
            break;
        // User just won the game
        case 'GameWin':
            victory = true;
            incScore();
            updateSubTab();
            break;
        // Custom sub tab asks for current user score
        case 'GetDataCenterScore':
            subTabWindow = sourceWindow;
            updateSubTab();
            break;
    }
}
```

Cross-Window Communication

oVirt

```
var winGame = function() {
  Game.setBoard(3,new TitleScreen("You win!",
                                  "Press fire to play again",
                                  playGame));
  parent.postMessage('GameWin', '*');
};
```

Let's start!

oVirt

- Call api.ready() to tell the oVirt engine to start initializing your plugin

oVirt

oVirt-Foreman plugin

oVirt Plugin – Dashboard

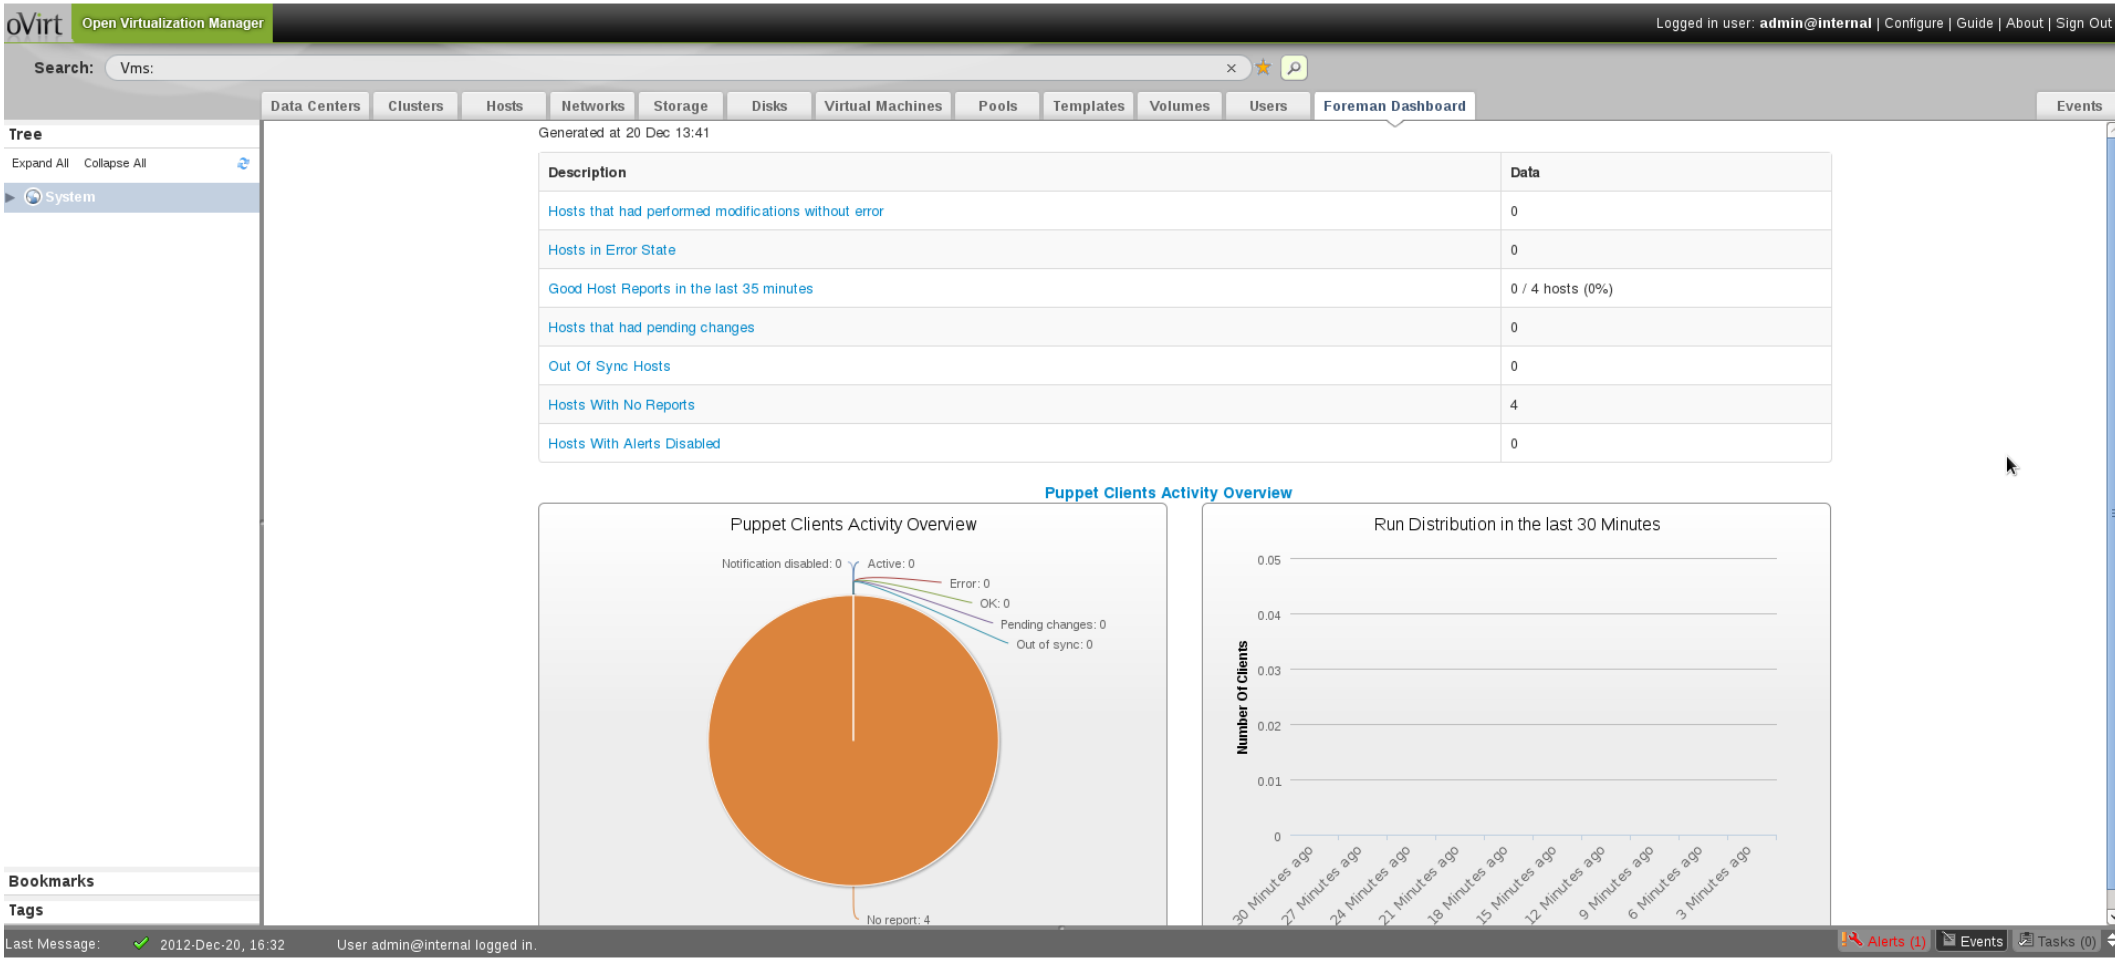

oVirt Plugin – VM Foreman Details

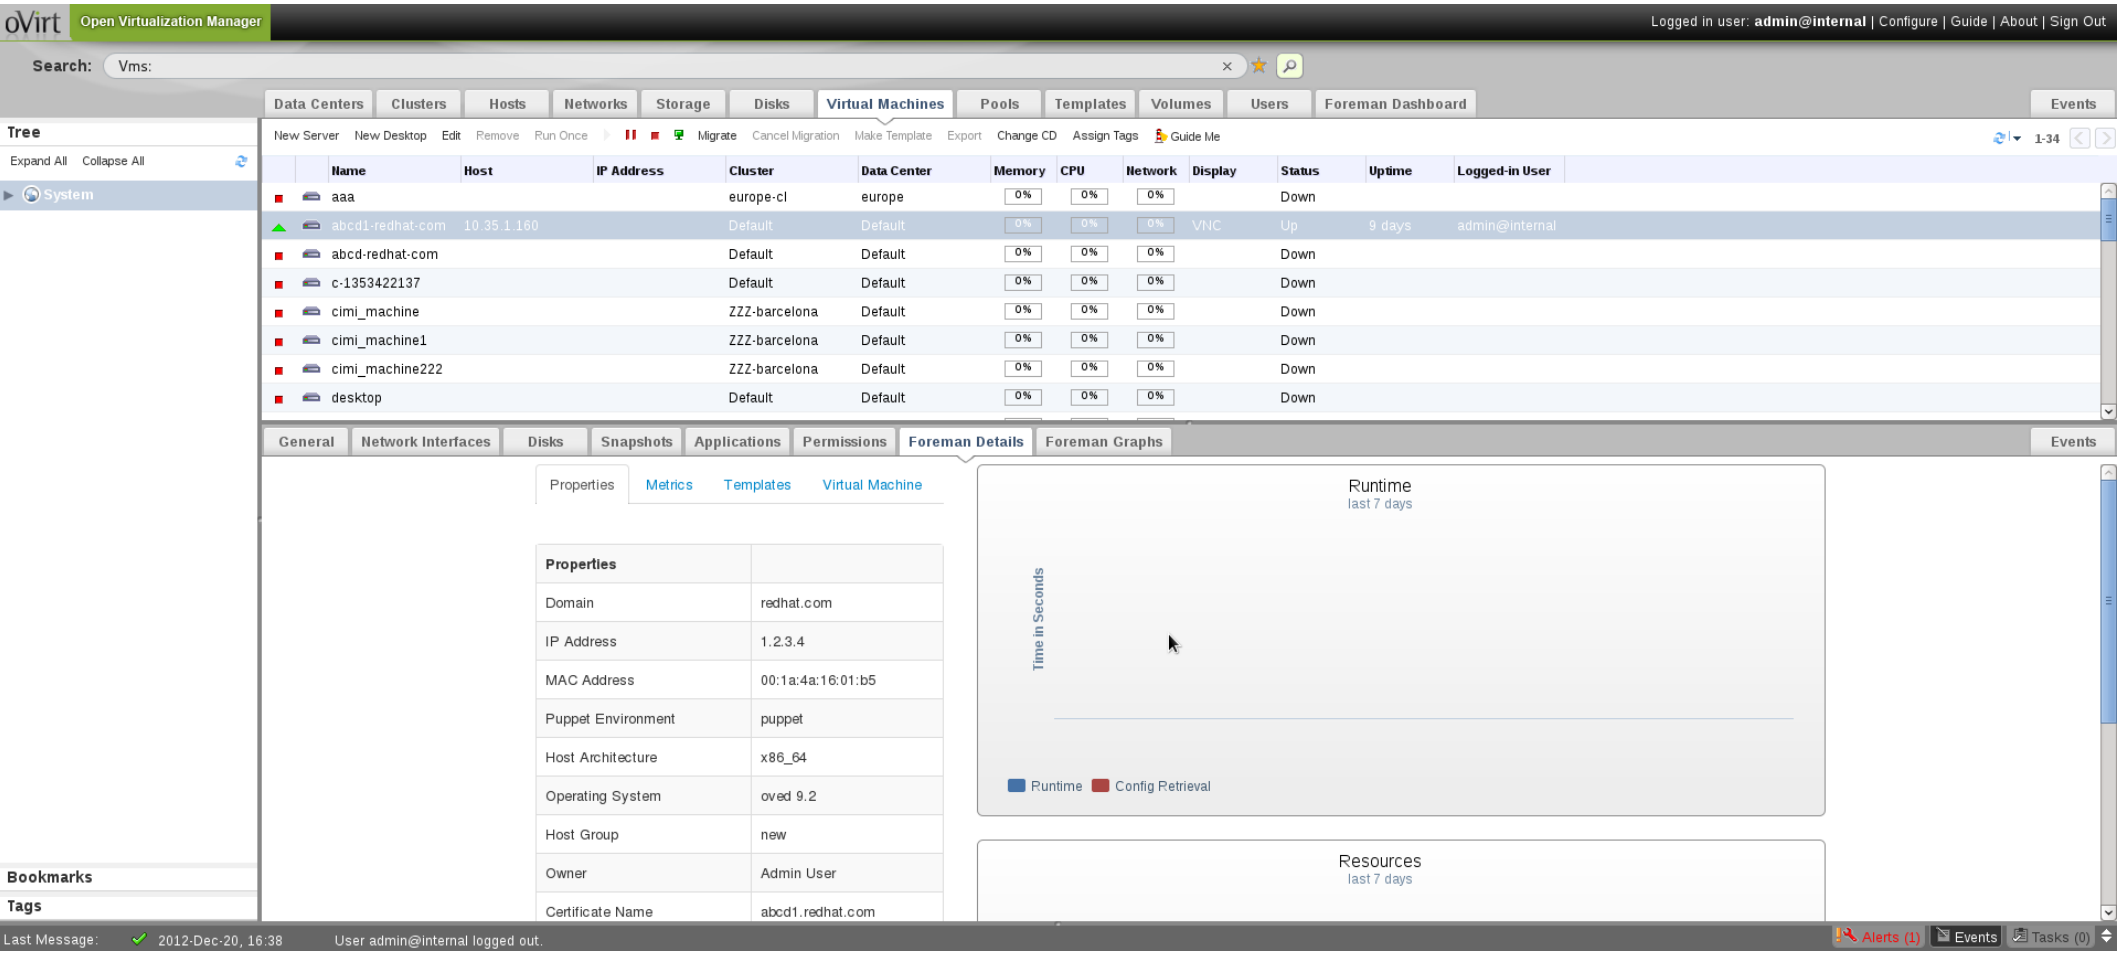

oVirt Plugin – VM Foreman Graphs

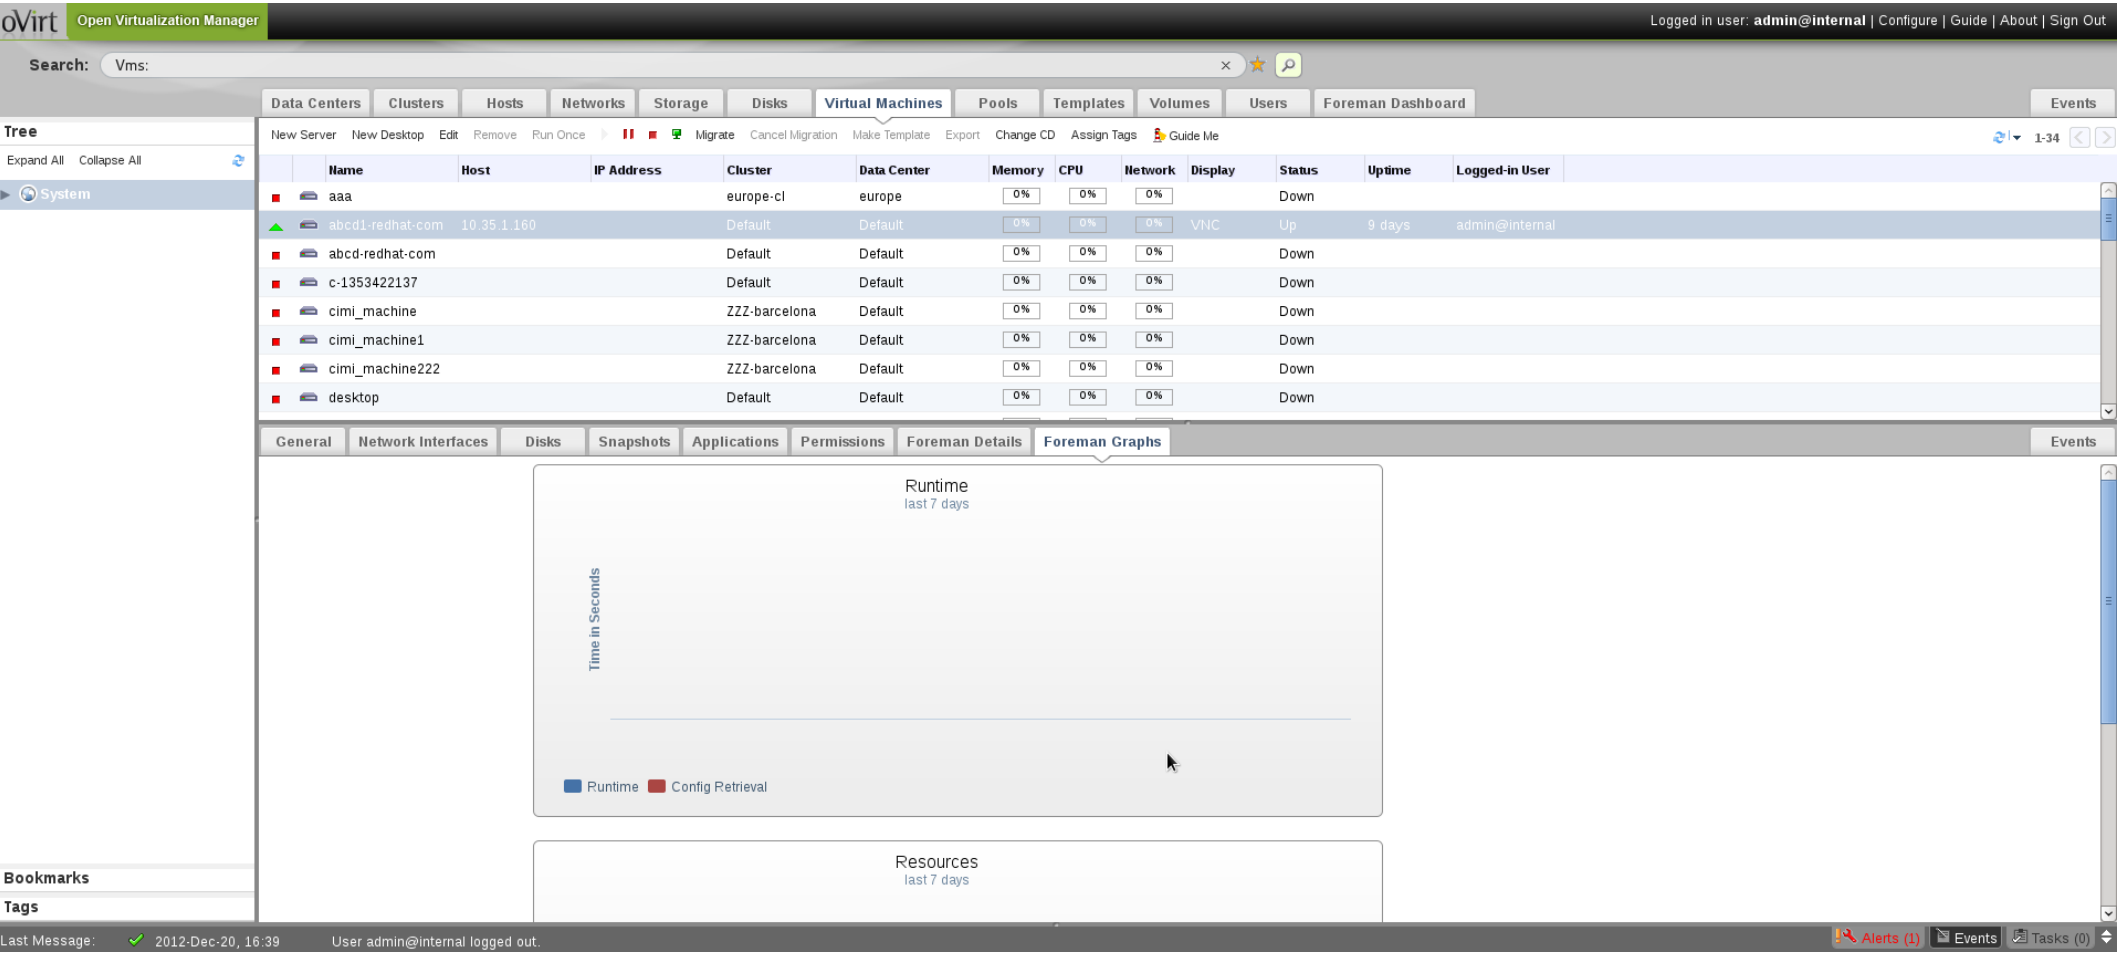

oVirt Plugin – start.html file

- Register section - register the following event handlers
  - UiInit – add the main/sub tabs

```
UiInit: function() {
    // Foreman Details will show the details of the Foreman Host that matches the VM
    // URL will be change upon VM selection event
    api.addSubTab('VirtualMachine', conf.foremanDetailsLabel, 'foreman-details', '');

    // Foreman Graphs will show the different Graphs of the Foreman Host that matches the VM
    // URL will be change upon VM selection event
    api.addSubTab('VirtualMachine', conf.foremanGraphsLabel, 'foreman-graphs', '');

    // We hide both sub-tabs until some VM is selected
    setAccessibleVmForemanSubTabs(false);

    // Dashboard Main Tab
    api.addMainTab(conf.foremanDashboardLabel, 'foreman', conf.url + '/dashboard/ovirt');
},
```

oVirt Plugin – start.html file

- RestApiSessionAcquired – relogin to Foreman

```
RestApiSessionAcquired: function(sessionId) {
        var userName = conf.ovirtUserNamePrefix + api.loginUserName();
        relogin(conf.logoutUrl, conf.loginUrl, userName, api.loginUserId(), sessionId);
},
```

- VirtualMachineSelectionChange – set sub-tabs URL if the selected VM exists in Foreman

```
VirtualMachineSelectionChange: function() {
    setAccessibleVmForemanSubTabs(false);
    if (arguments.length == 1) {
        var vmId = arguments[0].id;
        var vmName = arguments[0].name;
        var foremanSearchHostUrl = conf.url + '/hosts?&search=' +
                                                encodeURIComponent('uuid=' + vmId) + '+or+' +
                                                encodeURIComponent('name=' + vmName) + '&format=json';
        // Get the relevant host URLs and set up the Sub-Tabs accordingly
        $.getJSON(foremanSearchHostUrl, function (data) {
            try {
                // If a host was returned we set the show the Sub-Tabs
                if (data[0] != null) {
                    showVmForemanSubTabs(data[0].host.name);
                }
            } catch (err) {
                // We do nothing here. The tab will remain inaccessible
            }
        });
    }
},
```

oVirt Plugin – start.html file

- UserLogout – logout from Foreman

```
UserLogout: function() {
    // Logging out from Foreman as well
    $.get(conf.logoutUrl, { });
}
```

oVirt Plugin – start.html file

- Helper functions

```
var api = parent.pluginApi('foreman');
var conf = api.configObject();

function setAccessibleVmForemanSubTabs(shouldShowSubTabs) {
    api.setTabAccessible('foreman-details', shouldShowSubTabs);
    api.setTabAccessible('foreman-graphs', shouldShowSubTabs);
}

function showVmForemanSubTabs(hostName) {
    try {
        setAccessibleVmForemanSubTabs(false);
        var foremanUrl = conf.url + '/hosts/' + hostName;
        var foremanVmDetailsUrl = foremanUrl + '/ovirt';
        var foremanVmGraphsUrl = foremanUrl + '/graphs/ovirt';
        api.setTabContentUrl('foreman-details', foremanVmDetailsUrl);
        api.setTabContentUrl('foreman-graphs', foremanVmGraphsUrl);
        setAccessibleVmForemanSubTabs(true);
    } catch (err) {
        // We do nothing here. The tab will remain inaccessible
    }
}
```

oVirt

Scheduling API

Introduction

- The need - construct user-defined scheduling policy

  Re: [Users] How to define max number of running VMs on a host?

  ....

  I have 4 graphic workstations with 3 graphic cards on each. I wanna passthrough graphic cards to the VMs one by one, since one workstation has only 3 cards, I must limit the number of running VM on a host to 3.

Previous Scheduling Mechanism

oVirt

- Executes the selected distribution algorithm on the Cluster
  - Evenly Distributed
  - Power Saving
- Scheduling
  - Selects a host to run/migrate VM
- Load balancing
  - Selects a VM to migrate and Host to migrate to
- Two distribution algorithms, taking into consideration only CPU usage
- No way to construct user-defined policy

New Scheduling Mechanism

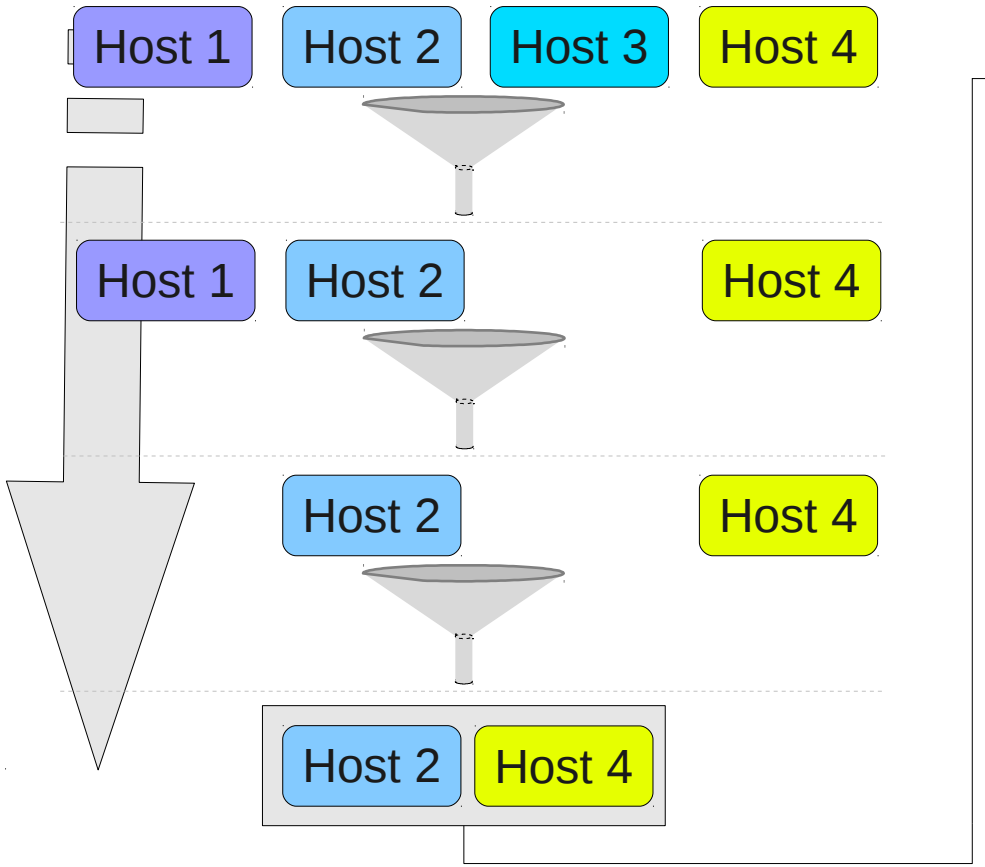

	func 1	func 2	sum
Factor	5	2	
Host 2	10	2	54
Host 4	3	12	**39***

***Host 4 sum: 3*5+12*2 = 39**

Filter Module

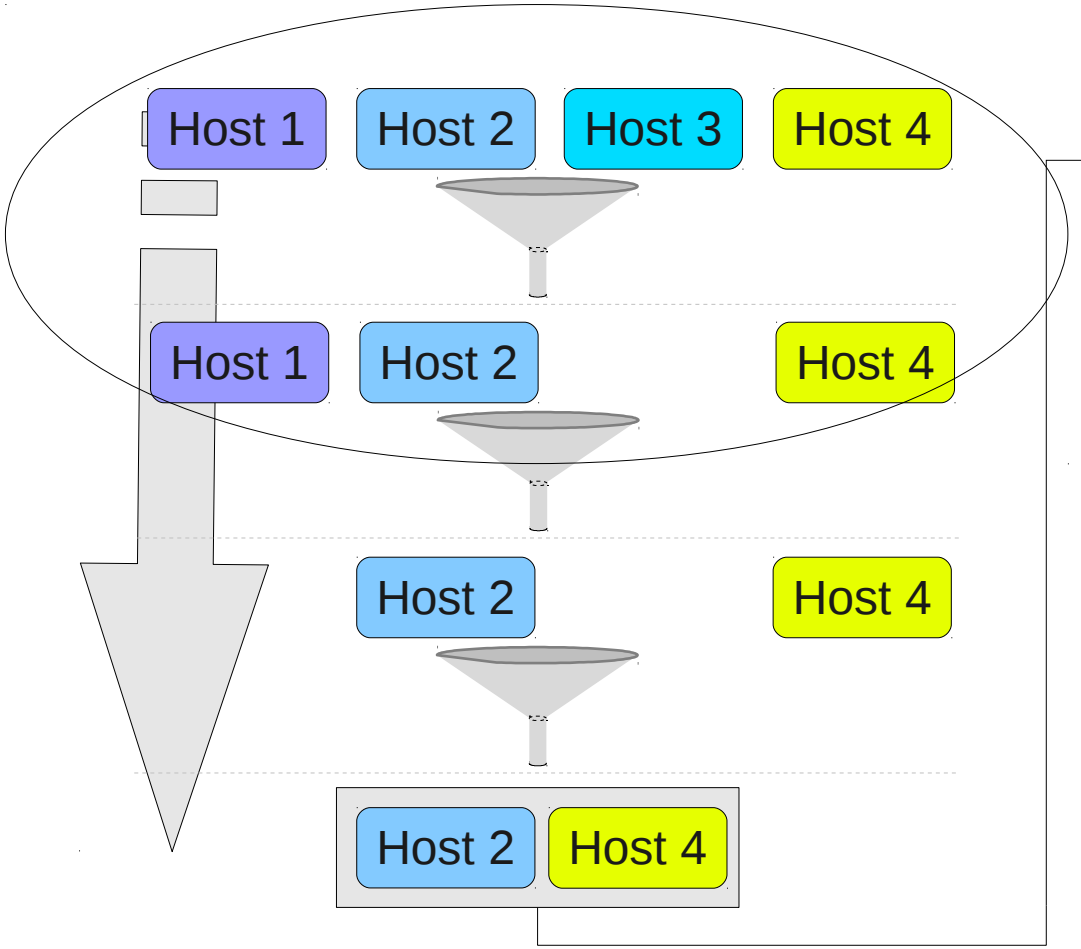

	func 1	func 2	sum
Factor	5	2	
Host 2	10	2	54
Host 4	3	12	**39***

***Host 4 sum: 3*5+12*2 = 39**

Weight Module

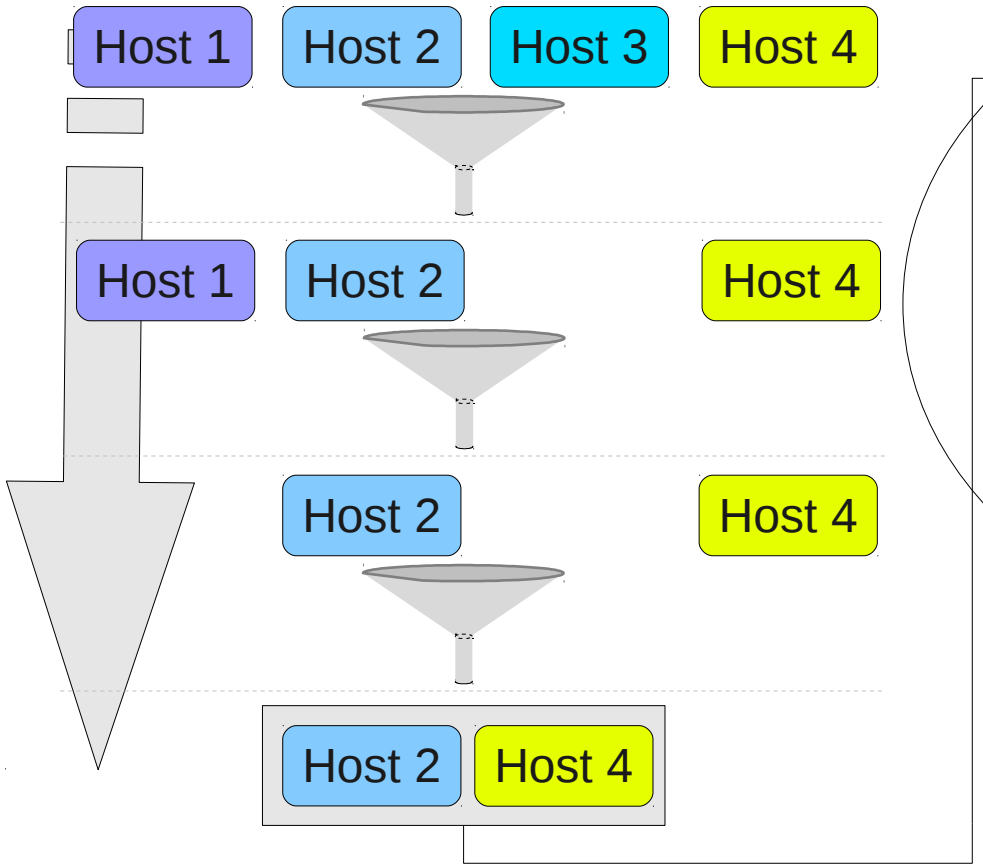

	func 1	func 2	sum
Factor	5	2	
Host 2	10	2	54
Host 4	3	12	**39***

***Host 4 sum: 3*5+12*2 = 39**

Filter Module

oVirt

- Logical unit which filters out hosts
  - Clear cut logic
  - Easy to write and maintain
  - Chained up-dependently to allow complete filtering
  - Allows custom parameters
- Existing logic (pin-to-host, memory limitations, etc.) is translated into filters
- External filters written in python can be loaded into engine

Let's go back to the example

Re: [Users] How to define max number of running VMs on a host?

....

I have 4 graphic workstations with 3 graphic cards on each. I wanna passthrough graphic cards to the VMs one by one, since one workstation has only 3 cards, I must limit the number of running VM on a host to 3.

Filter: filters out hosts with number running of vms > 3

Filter Example

```
class max_vms():
    '''returns only hosts with less running vms then the maximum'''

    #What are the values this module will accept, used to present
    #the user with options
    properties_validation = 'maximum_vm_count=[0-9]*'

    def do_filter(self, hosts_ids, vm_id, args_map):
        #open a connection to the rest api
        try:
            connection = API(url='http://host:port',
                             username='user@domain', password='')
        except BaseException as ex:
            #letting the external proxy know there was an error
            print >> sys.stderr, ex
            return

        #get our parameters from the map
        maximum_vm_count = int(args_map.get('maximum_vm_count', 100))

        #get all the hosts with the given ids
        engine_hosts = \
            connection.hosts.list(
                query=" or ".join(["id=%s" % u for u in hosts_ids]))

        #iterate over them and decide which to accept
        accepted_host_ids = []
        for engine_host in engine_hosts:
            if(engine_host and
                    engine_host.summary.active < maximum_vm_count):
                accepted_host_ids.append(engine_host.id)
        print accepted_host_ids
```

Filter Example

```
class max_vms():
    '''returns only hosts with less running vms then the maximum'''

    ##What are the values this module will accept, used to present
    #the user with options
    properties_validation = 'maximum_vm_count=[0-9]*'

    def do_filter(self, hosts_ids, vm_id, args_map):
        #open a connection to the rest api
        try:
            connection = API(url='http://host:port',
                             username='user@domain', password='')
        except BaseException as ex:
            #letting the external proxy know there was an error
            print >> sys.stderr, ex
            return

        #get our parameters from the map
        maximum_vm_count = int(args_map.get('maximum_vm_count', 100))

        #get all the hosts with the given ids
        engine_hosts = \
            connection.hosts.list(
                query=" or ".join(["id=%s" % u for u in hosts_ids]))

        #iterate over them and decide which to accept
        accepted_host_ids = []
        for engine_host in engine_hosts:
            if(engine_host and
                    engine_host.summary.active < maximum_vm_count):
                accepted_host_ids.append(engine_host.id)
        print accepted_host_ids
```

Filter Example

```
class max_vms():
    '''returns only hosts with less running vms then the maximum'''

    #What are the values this module will accept, used to present
    #the user with options
    properties_validation = 'maximum_vm_count=[0-9]*'

    def do_filter(self, hosts_ids, vm_id, args_map):
        #open a connection to the rest api
        try:
            connection = API(url='http://host:port',
                             username='user@domain', password='')
        except BaseException as ex:
            #letting the external proxy know there was an error
            print >> sys.stderr, ex
            return

        #get our parameters from the map
        maximum_vm_count = int(args_map.get('maximum_vm_count', 100))

        #get all the hosts with the given ids
        engine_hosts = \
            connection.hosts.list(
                query=" or ".join(["id=%s" % u for u in hosts_ids]))

        #iterate over them and decide which to accept
        accepted_host_ids = []
        for engine_host in engine_hosts:
            if(engine_host and
                    engine_host.summary.active < maximum_vm_count):
                accepted_host_ids.append(engine_host.id)
        print accepted_host_ids
```

Filter Example

```
class max_vms():
    '''returns only hosts with less running vms then the maximum'''

    #What are the values this module will accept, used to present
    #the user with options
    properties_validation = 'maximum_vm_count=[0-9]*'

    def do_filter(self, hosts_ids, vm_id, args_map):
        #open a connection to the rest api
        try:
            connection = API(url='http://host:port',
                             username='user@domain', password='')
        except BaseException as ex:
            #letting the external proxy know there was an error
            print >> sys.stderr, ex
            return

        #get our parameters from the map
        maximum_vm_count = int(args_map.get('maximum_vm_count', 100))

        #get all the hosts with the given ids
        engine_hosts = \
            connection.hosts.list(
                query=" or ".join(["id=%s" % u for u in hosts_ids]))

        #iterate over them and decide which to accept
        accepted_host_ids = []
        for engine_host in engine_hosts:
            if(engine_host and
                    engine_host.summary.active < maximum_vm_count):
                accepted_host_ids.append(engine_host.id)
        print accepted_host_ids
```

Filter Example

```
class max_vms():
    '''returns only hosts with less running vms then the maximum'''

    #What are the values this module will accept, used to present
    #the user with options
    properties_validation = 'maximum_vm_count=[0-9]*'

    def do_filter(self, hosts_ids, vm_id, args_map):
        #open a connection to the rest api
        try:
            connection = API(url='http://host:port',
                             username='user@domain', password='')
        except BaseException as ex:
            #letting the external proxy know there was an error
            print >> sys.stderr, ex
            return

        #get our parameters from the map
        maximum_vm_count = int(args_map.get('maximum_vm_count', 100))

        #get all the hosts with the given ids
        engine_hosts = \
            connection.hosts.list(
                query=" or ".join(["id=%s" % u for u in hosts_ids]))

        #iterate over them and decide which to accept
        accepted_host_ids = []
        for engine_host in engine_hosts:
            if(engine_host and
                    engine_host.summary.active < maximum_vm_count):
                accepted_host_ids.append(engine_host.id)
        print accepted_host_ids
```

Weight Module

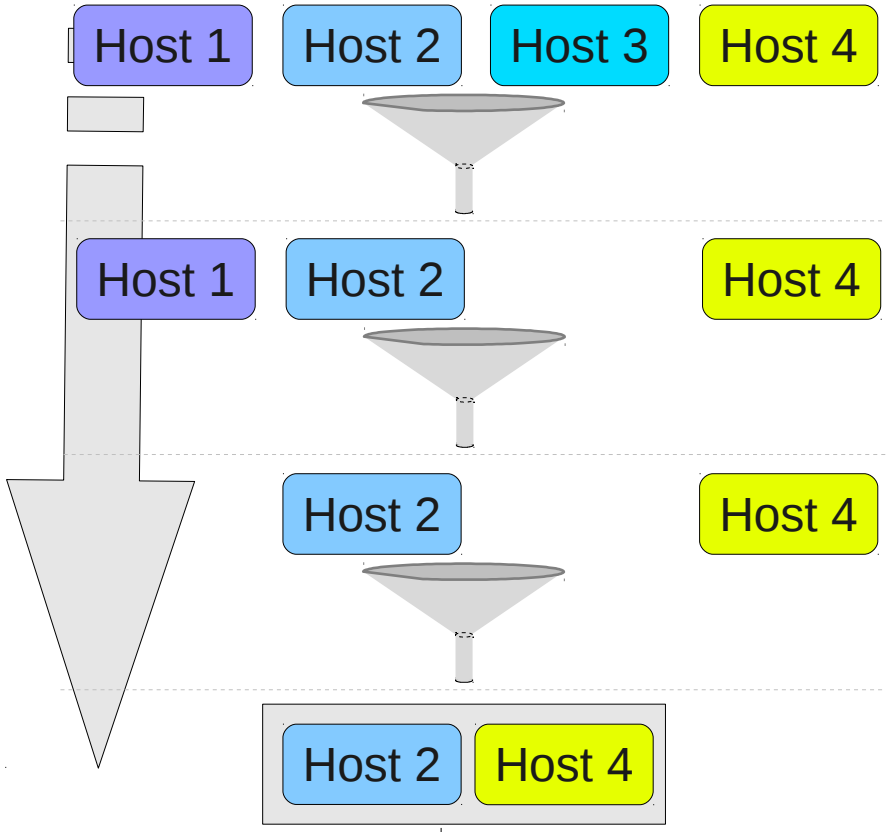

	func 1	func 2	sum
Factor	5	2	
Host 2	10	2	54
Host 4	3	12	**39***

***Host 4 sum: 3*5+12*2 = 39**

Weight Module

oVirt

- Logical unit which weights hosts
- The lower the better
- Weights can be prioritized using Factors (defaults to 1)
- The result is a score table, which will be taken into consideration when scheduling the VM

Weight Module

oVirt

- Predefined Weight Modules
  - Even Distribution
    - Each host weight will be scored according to CPU load
  - Power Saving
    - Define Max_Weight
    - if (no VMs on Host) → Max_Weight
    - Else (Max_Weight – Even_Distribution_Weight)
- External Weight Modules written in python can be loaded into the engine

Weight Module Example

```
class even_vm_distribution():
    '''rank hosts by the number of running vms on them, with the least first'''

    properties_validation = ''

    def do_score(self, hosts_ids, vm_id, args_map):
        #open a connection to the rest api
        try:
            connection = API(url='http://host:port',
                             username='user@domain', password='')
        except BaseException as ex:
            #letting the external proxy know there was an error
            print >> sys.stderr, ex
            return

        #get all the hosts with the given ids
        engine_hosts = \
            connection.hosts.list(
                query=" or ".join(["id=%s" % u for u in hosts_ids]))

        #iterate over them and score them based on the number of vms running
        host_scores = []
        for engine_host in engine_hosts:
            if(engine_host and
                    engine_host.summary):
                host_scores.append((engine_host.id, engine_host.summary.active))
        print host_scores
```

Weight Module Example

```
class even_vm_distribution():
    '''rank hosts by the number of running vms on them, with the least first'''

    properties_validation = ''

    def do_score(self, hosts_ids, vm_id, args_map):
        #open a connection to the rest api
        try:
            connection = API(url='http://host:port',
                             username='user@domain', password='')
        except BaseException as ex:
            #letting the external proxy know there was an error
            print >> sys.stderr, ex
            return

        #get all the hosts with the given ids
        engine_hosts = \
            connection.hosts.list(
                query=" or ".join(["id=%s" % u for u in hosts_ids]))

        #iterate over them and score them based on the number of vms running
        host_scores = []
        for engine_host in engine_hosts:
            if(engine_host and
                    engine_host.summary):
                host_scores.append((engine_host.id, engine_host.summary.active))
        print host_scores
```

Weight Module Example

```
class even_vm_distribution():
    '''rank hosts by the number of running vms on them, with the least first'''

    properties_validation = ''

    def do_score(self, hosts_ids, vm_id, args_map):
        #open a connection to the rest api
        try:
            connection = API(url='http://host:port',
                             username='user@domain', password='')
        except BaseException as ex:
            #letting the external proxy know there was an error
            print >> sys.stderr, ex
            return

        #get all the hosts with the given ids
        engine_hosts = \
            connection.hosts.list(
                query=" or ".join(["id=%s" % u for u in hosts_ids]))

        #iterate over them and score them based on the number of vms running
        host_scores = []
        for engine_host in engine_hosts:
            if(engine_host and
                    engine_host.summary):
                host_scores.append((engine_host.id, engine_host.summary.active))
        print host_scores
```

Load Balancing

- Triggers a scheduled task to determine which VM needs to be migrated
- A single load balancing logic is allowed per cluster

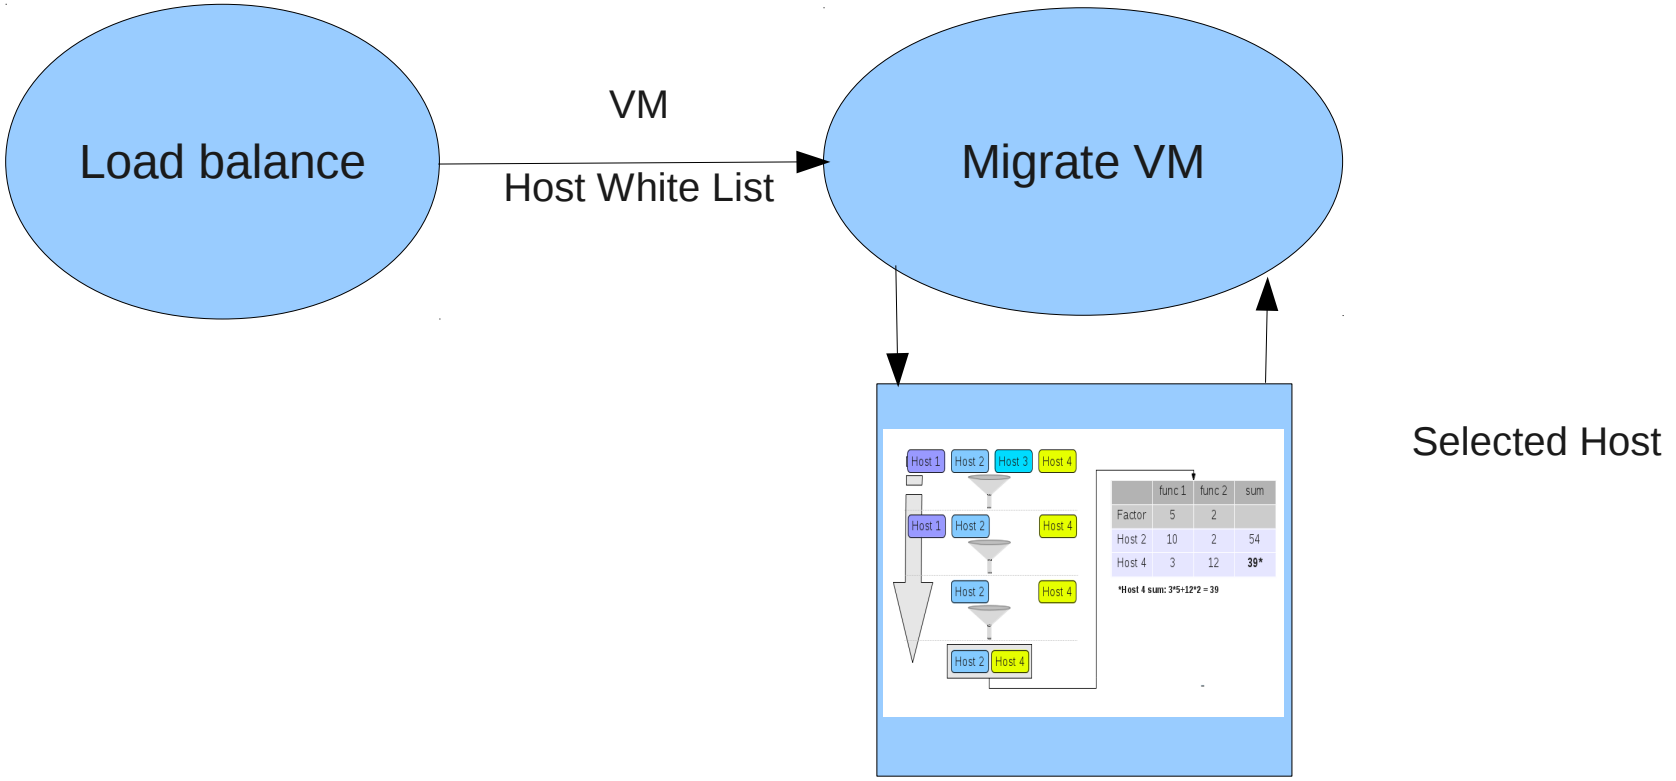

Load Balancing

oVirt

- For backward compatibility we have 2 predefined Load Balancing algorithms
  - Even Distribution
    - Calculates over-utilized and under-utilized hosts according to upper CPU load threshold
    - Select a VM out of the over-utilized hosts
    - Pass VM and under-utilized hosts to the scheduler
    - migrate VM to the host selected by the scheduler
  - Power Saving
    - Same as Even Distribution, but with a second threshold for low CPU load

Load Balancing Example

… same as previous

```
#iterate over them and decide which to balance from
over_loaded_host = None
white_listed_hosts = []
for engine_host in engine_hosts:
    if(engine_host):
        if (engine_host.summary.active < maximum_vm_count):
            white_listed_hosts.append(engine_host.id)
            continue
        if(not over_loaded_host or
                over_loaded_host.summary.active
                < engine_host.summary.active):
            over_loaded_host = engine_host

if(not over_loaded_host):
    return

selected_vm = None
#just pick the first we find
host_vms = connection.vms.list('host='+over_loaded_host.name)
if host_vms:
    selected_vm = host_vms[0].id
else:
    return

print (selected_vm, white_listed_hosts)
```

Load Balancing Example

… same as previous

```
#iterate over them and decide which to balance from
over_loaded_host = None
white_listed_hosts = []
for engine_host in engine_hosts:
    if(engine_host):
        if (engine_host.summary.active < maximum_vm_count):
            white_listed_hosts.append(engine_host.id)
            continue
        if(not over_loaded_host or
                over_loaded_host.summary.active
                < engine_host.summary.active):
            over_loaded_host = engine_host

if(not over_loaded_host):
    return

selected_vm = None
#just pick the first we find
host_vms = connection.vms.list('host='+over_loaded_host.name)
if host_vms:
    selected_vm = host_vms[0].id
else:
    return

print (selected_vm, white_listed_hosts)
```

Load Balancing Example

… same as previous

```
#iterate over them and decide which to balance from
over_loaded_host = None
white_listed_hosts = []
for engine_host in engine_hosts:
    if(engine_host):
        if (engine_host.summary.active < maximum_vm_count):
            white_listed_hosts.append(engine_host.id)
            continue
        if(not over_loaded_host or
                over_loaded_host.summary.active
                < engine_host.summary.active):
            over_loaded_host = engine_host

if(not over_loaded_host):
    return

selected_vm = None
#just pick the first we find
host_vms = connection.vms.list('host='+over_loaded_host.name)
if host_vms:
    selected_vm = host_vms[0].id
else:
    return

print (selected_vm, white_listed_hosts)
```

External Policy Units

- Scanning directory /usr/share/ovirt-scheduler-proxy/plugins for python source files
  - Analyze for filter / weight / balance functions
  - Cache results
  - Expose source files as external policy units

Cluster Policy Management

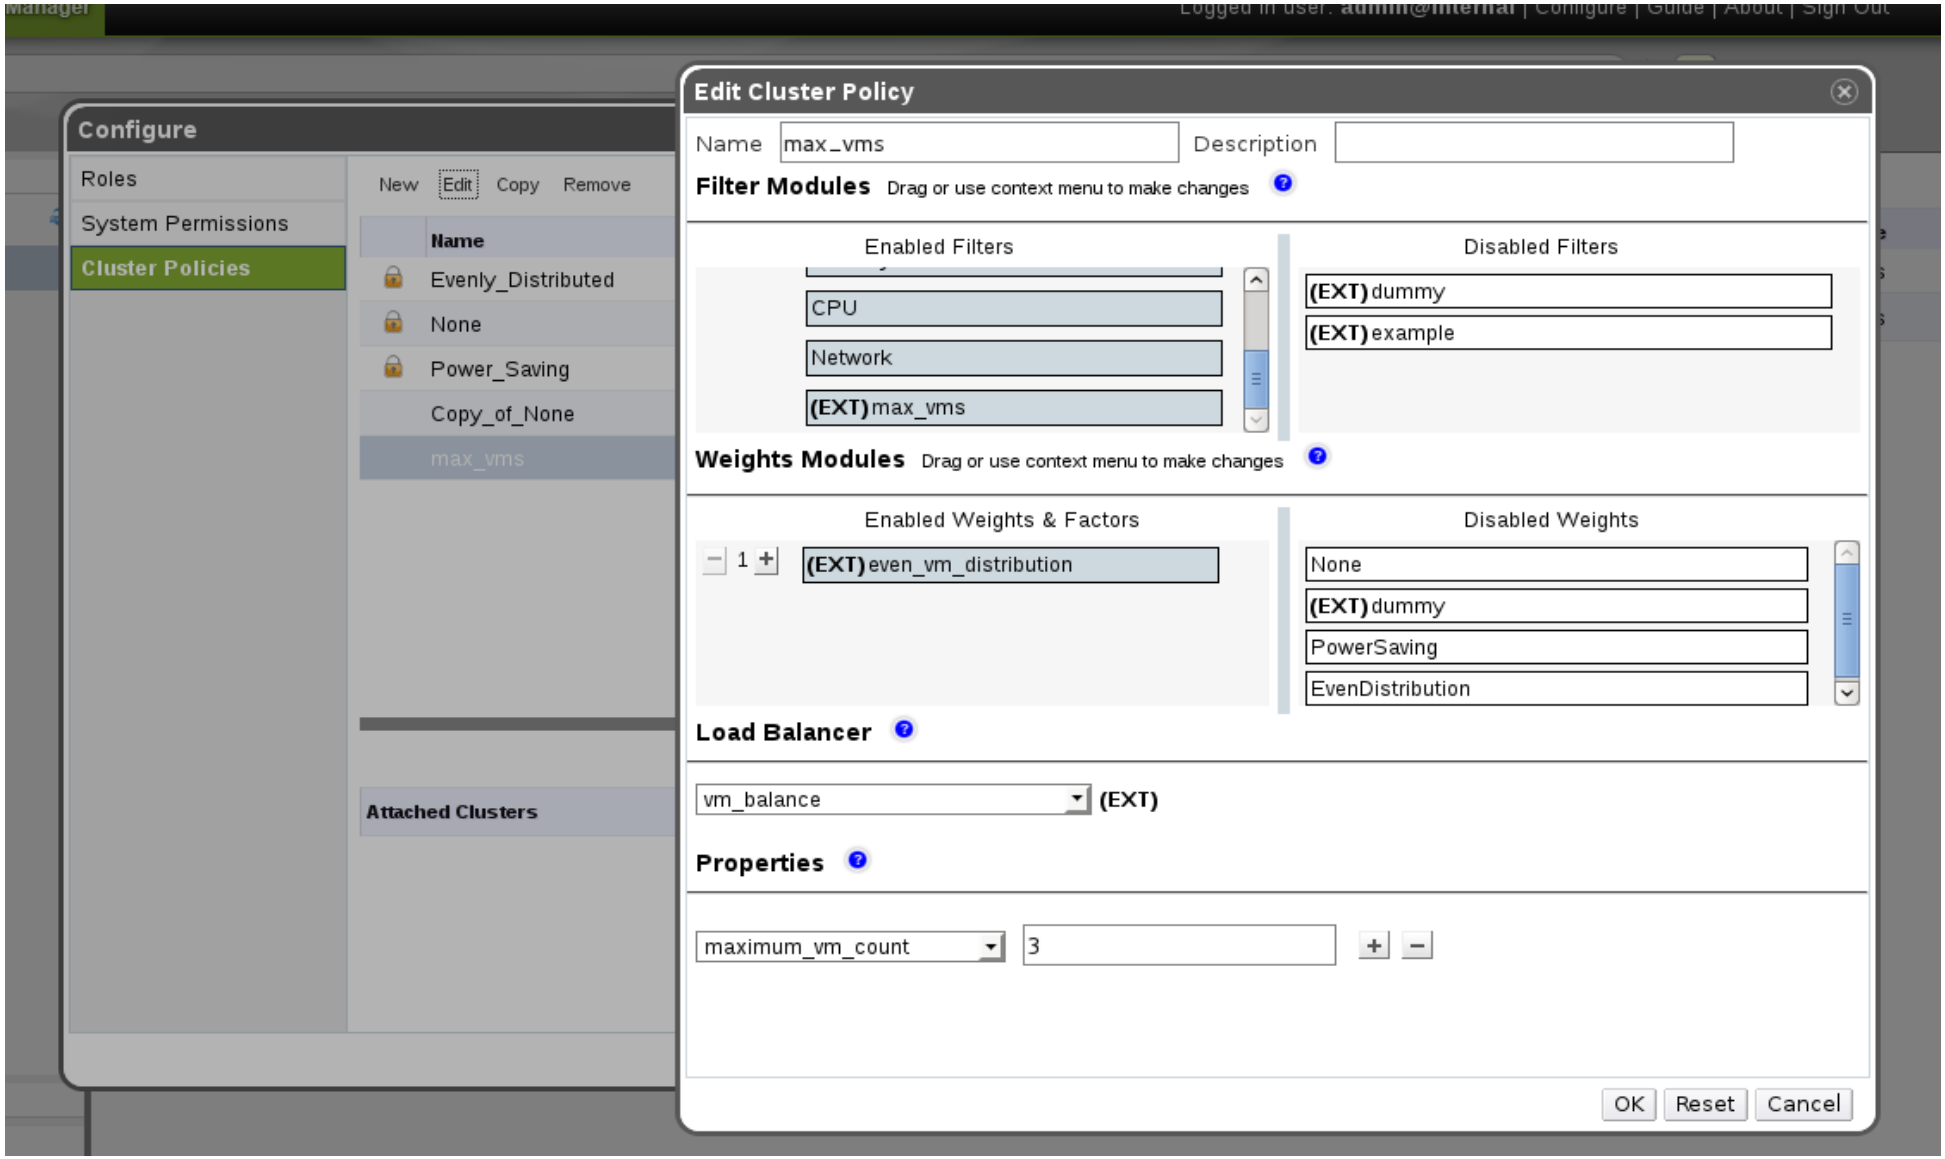

Apply Cluster Policy

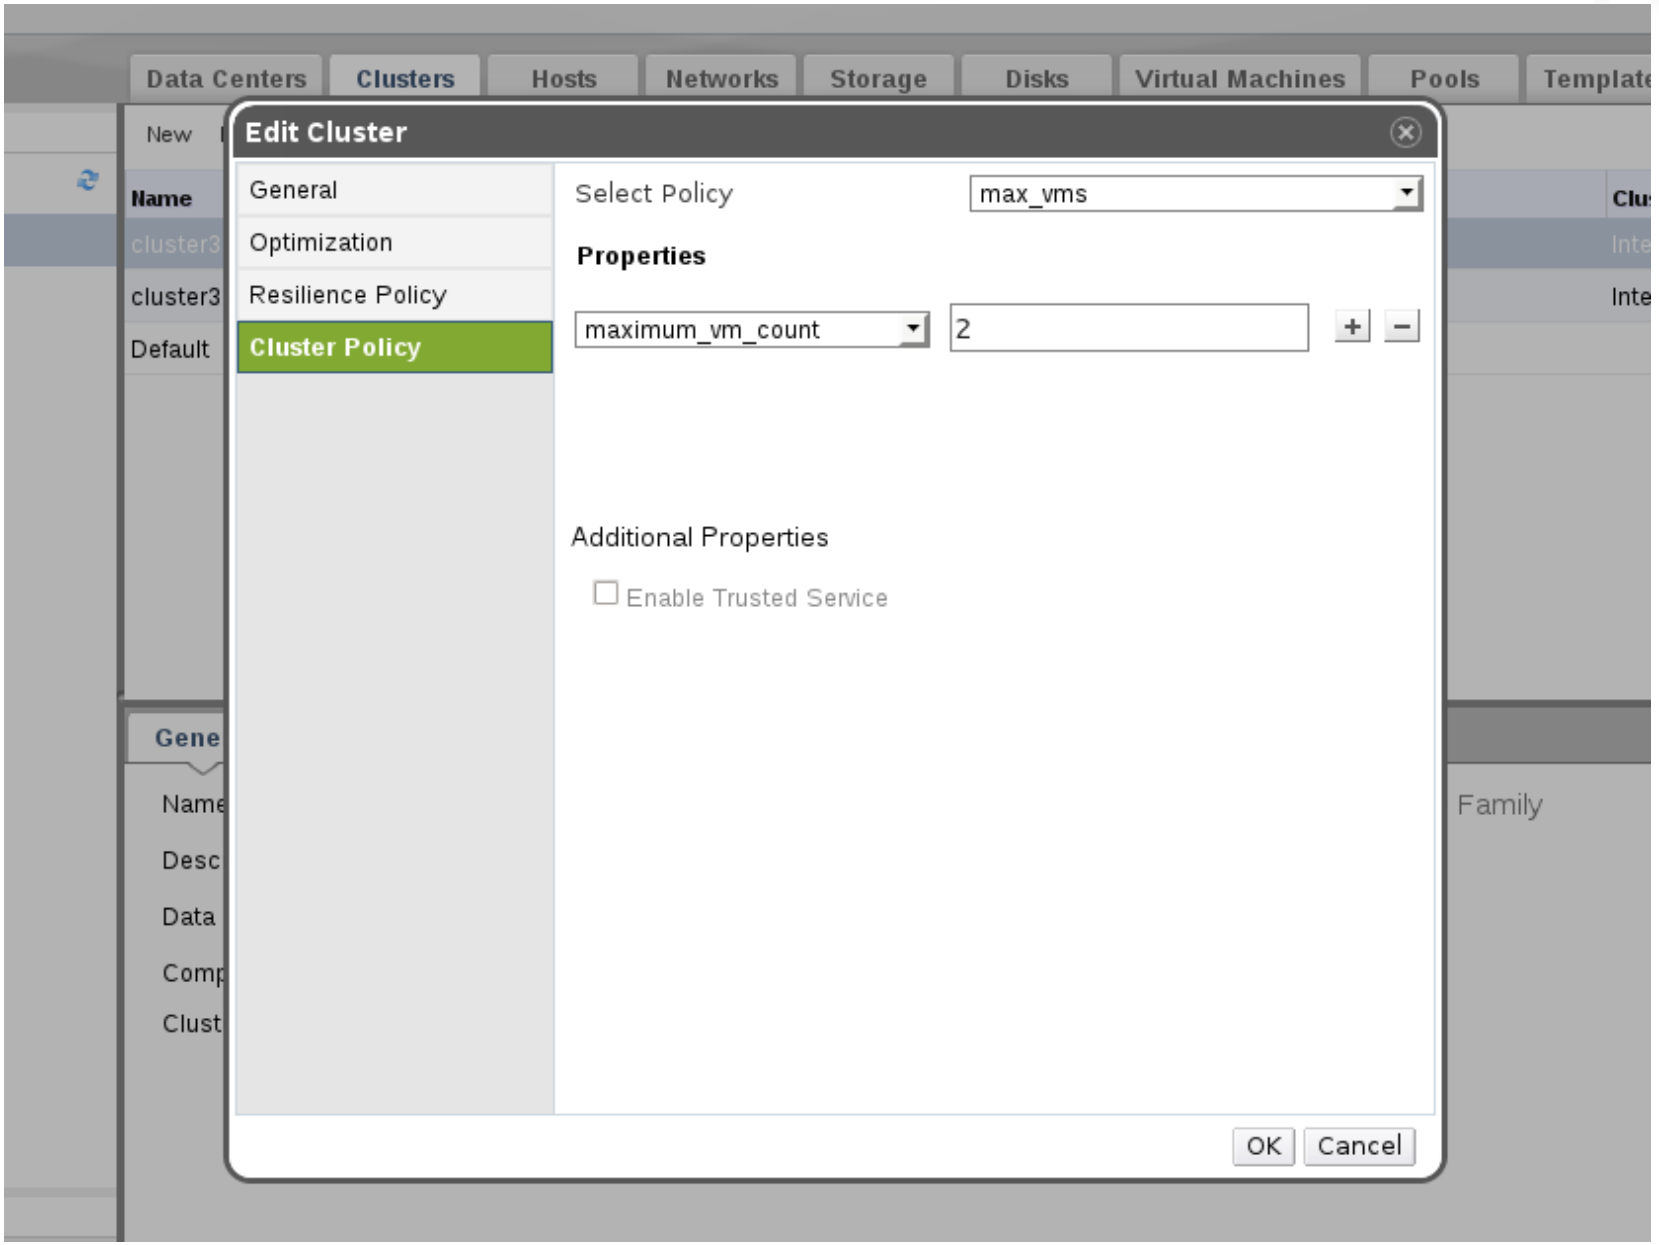

oVirt

VDSM Hooks

http://www.ovirt.org/Vdsm_Hooks
http://www.ovirt.org/VDSM-Hooks_Catalogue

Hooks

- “Hook” mechanism for customization
  - Allows administrator to define scripts to modify VM operation
    - Extend or modify VM configuration
    - Run different system scripts

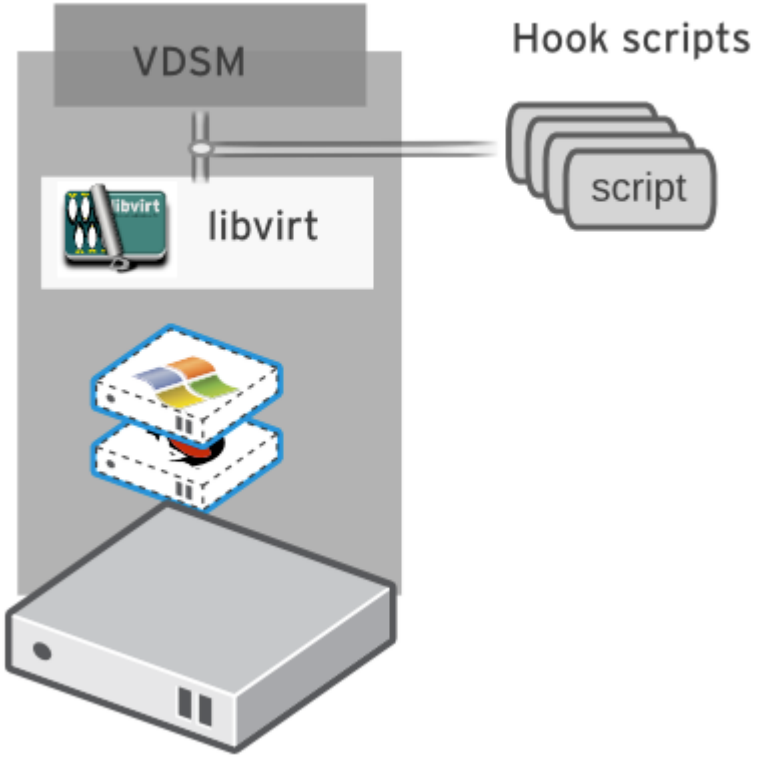

Hooks

oVirt

- Hook scripts are called at specific VM lifecycle events
- Hooks can modify a virtual machines XML definition before VM start
- Hooks can run system commands – eg. Apply firewall rule to VM

Hooks (cont.)

oVirt

- Lifecycle events where you can apply hooks
  - VDSM (management agent) Start
  - Before VM start
  - After VM start
  - Before VM migration in/out
  - After VM migration in/out
  - Before and After VM Pause
  - Before and After VM Continue
  - Before and After VM Hibernate
  - Before and After VM resume from hibernate
  - Before and After VM set ticket
  - On VM stop
  - On VDSM Stop

Hooks (cont.)

oVirt

- Before NIC hotplug / hotunplug
- After NIC hotplug / hotunplug

Hooks

oVirt

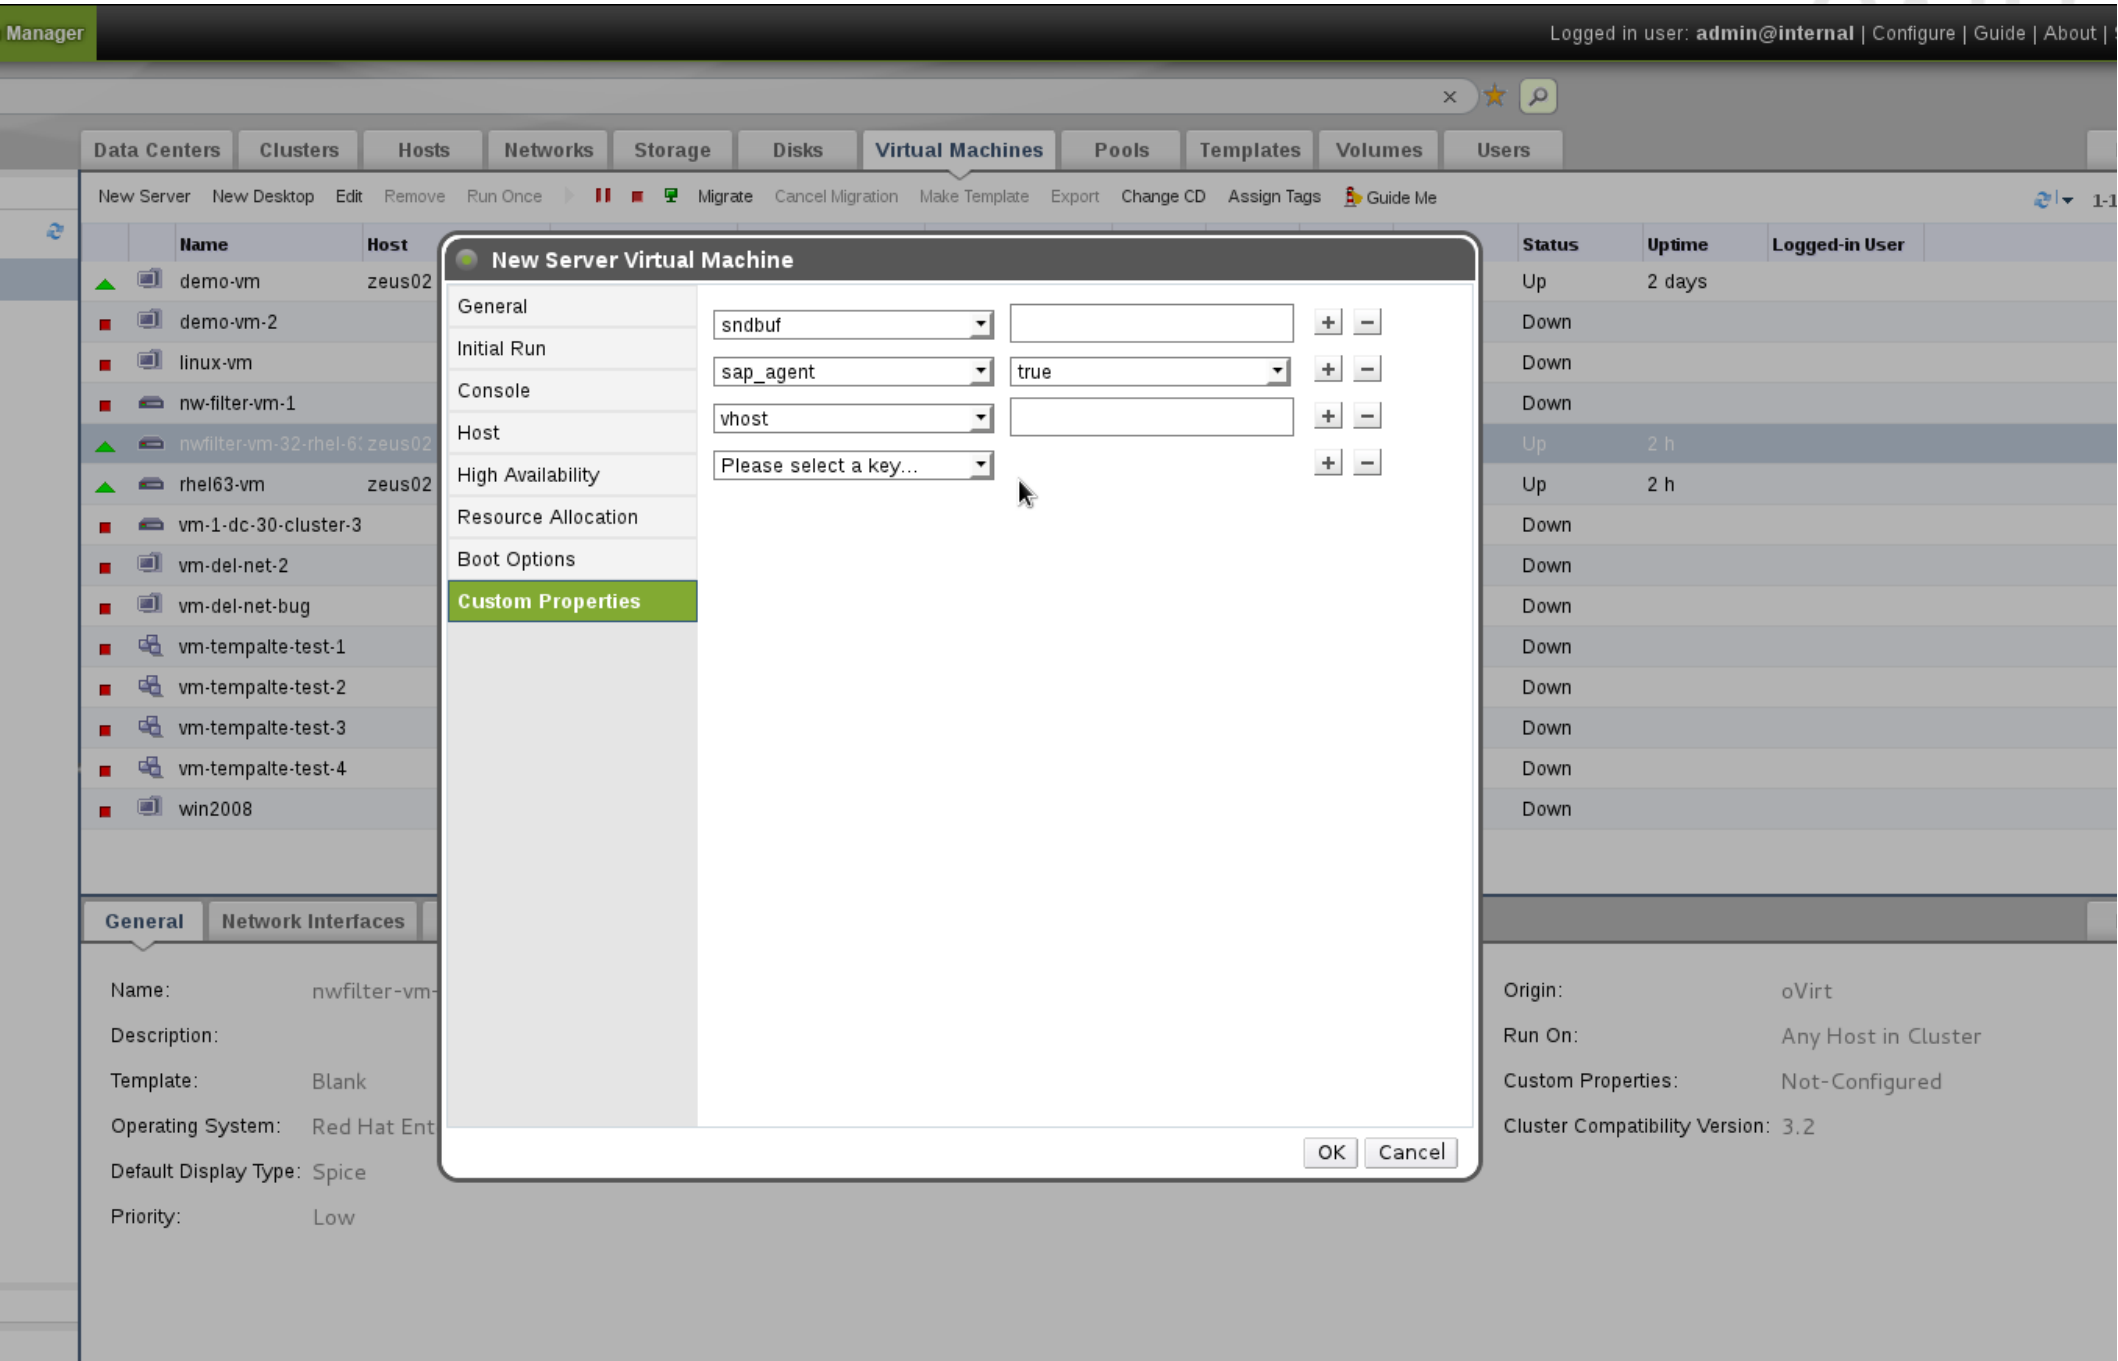
Manager
Logged in user: admin@internal | Configure | Guide | About |
Data Centers
Clusters
Hosts
Networks
Storage
Disks
Virtual Machines
Pools
Templates
Volumes
Users
New Server
New Desktop
Edit
Remove
Run Once
Migrate
Cancel Migration
Make Template
Export
Change CD
Assign Tags
Guide Me
Name
Host
Status
Uptime
Logged-in User
demo-vm
zeus02
Up
2 days
demo-vm-2
Down
linux-vm
Down
nw-filter-vm-1
Down
rhel63-vm
zeus02
Up
2 h
vm-1-dc-30-cluster-3
Down
vm-del-net-2
Down
vm-del-net-bug
Down
vm-tempalte-test-1
Down
vm-tempalte-test-2
Down
vm-tempalte-test-3
Down
vm-tempalte-test-4
Down
win2008
Down
New Server Virtual Machine
General
Initial Run
Console
Host
High Availability
Resource Allocation
Boot Options
Custom Properties
sndbuf
sap_agent
true
vhost
Please select a key...
OK
Cancel
General
Network Interfaces
Name:
Description:
Template:
Blank
Operating System:
Red Hat Ent
Default Display Type:
Spice
Priority:
Low
Origin:
oVirt
Run On:
Any Host in Cluster
Custom Properties:
Not-Configured
Cluster Compatibility Version:
3.2

Hooks

- Hooks installed in /usr/libexec/vdsm/hooks

```
/usr/libexec/vdsm/hooks
[root@dhcp-1-160 hooks]# ls -l
total 104
drwxr-xr-x 2 root root 4096 May 27 18:00 after_nic_hotplug
drwxr-xr-x 2 root root 4096 May 27 18:00 after_nic_hotplug_fail
drwxr-xr-x 2 root root 4096 May 27 18:00 after_nic_hotunplug
drwxr-xr-x 2 root root 4096 May 27 18:00 after_nic_hotunplug_fail
drwxr-xr-x 2 root root 4096 May 27 18:00 after_vdsm_stop
drwxr-xr-x 2 root root 4096 May 27 18:00 after_vm_cont
drwxr-xr-x 2 root root 4096 May 27 18:00 after_vm_dehibernate
drwxr-xr-x 2 root root 4096 May 27 18:00 after_vm_destroy
drwxr-xr-x 2 root root 4096 May 27 18:00 after_vm_hibernate
drwxr-xr-x 2 root root 4096 May 27 18:00 after_vm_migrate_destination
drwxr-xr-x 2 root root 4096 May 27 18:00 after_vm_migrate_source
drwxr-xr-x 2 root root 4096 May 27 18:00 after_vm_pause
drwxr-xr-x 2 root root 4096 May 27 18:00 after_vm_set_ticket
drwxr-xr-x 2 root root 4096 May 27 18:00 after_vm_start
drwxr-xr-x 2 root root 4096 May 27 18:00 before_nic_hotplug
drwxr-xr-x 2 root root 4096 May 27 18:00 before_nic_hotunplug
drwxr-xr-x 2 root root 4096 May 27 18:00 before_vdsm_start
drwxr-xr-x 2 root root 4096 May 27 18:00 before_vm_cont
drwxr-xr-x 2 root root 4096 May 27 18:00 before_vm_dehibernate
drwxr-xr-x 2 root root 4096 May 27 18:00 before_vm_destroy
drwxr-xr-x 2 root root 4096 May 27 18:00 before_vm_hibernate
drwxr-xr-x 2 root root 4096 May 27 18:00 before_vm_migrate_destination
drwxr-xr-x 2 root root 4096 May 27 18:00 before_vm_migrate_source
drwxr-xr-x 2 root root 4096 May 27 18:00 before_vm_pause
drwxr-xr-x 2 root root 4096 May 27 18:00 before_vm_set_ticket
drwxr-xr-x 2 root root 4096 May 27 18:00 before_vm_start
```

- In the Host sub-tab

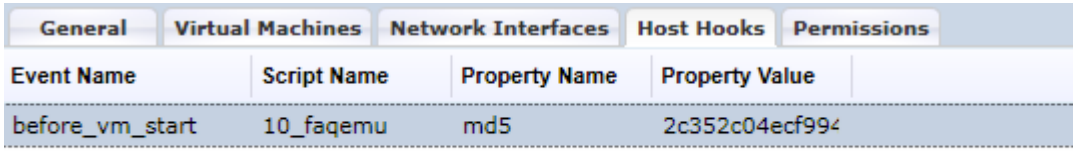

Hooks

```
#!/usr/bin/python

import os
import sys
import hooking
import traceback
from xml.dom import minidom

'''
watchdog vdsm hook
adding to domain xml
<watchdog model='i6300esb' action='reset'/>
'''

if os.environ.has_key('watchdog'):
    try:
        sys.stderr.write('watchdog: adding watchdog support\n')
        domxml = hooking.read_domxml()

        devices = domxml.getElementsByTagName('devices')[0]
        card = domxml.createElement('watchdog')
        card.setAttribute('model', 'i6300esb')
        card.setAttribute('action', 'reset')

        devices.appendChild(card)

        hooking.write_domxml(domxml)
    except:
        sys.stderr.write('watchdog: [unexpected error]: %s\n' % traceback.format_exc())
        sys.exit(2)
```

Hooks Catalog

directlun - Attach a LUN to a VM

faqemu - Fake QEMU emulation (the VM will "think" it has VT extensions enabled)

fileinject - Insert a file into the VM

floppy - Attach a floppy image to the VM

hostusb - Attach a hosts' USB device to the VM

hugepages - Enable hugepages for the VM

isolatedprivatevlan - Connect VM to an isolated private VLAN instead of a shared network

numa - Pin a VM to a NUMA node(s)

pincpu - Pin a VM to a set of CPU cores

promisc - Start VM on a promiscous mode enabled network

qemucmdline - Run an arbitrary QEMU command line

Hooks Catalog

qos - Set a specific network QoS for a VM

scratchpad - Start the VM with a disposable disk image, that will be removed when the VM is stopped

smartcard - Enable SPICE smartcard support for a VM

smbios - Alter a VMs smbios domain entry

sriov - Attach an SRIOV device to a VM

vhostmd - Enable vhostmd for a VM

vmdisk - Add an additional VM disk image

vmfex - Attach a Cisco VM-FEX device to a VM

oVirt

My Details:

ovedo@redhat.com

ovedo in #ovirt on irc.oftc.net

ovedou.blogspot.com

Mailing lists:

users@ovirt.org

engine-devel@ovirt.org

IRC:

#ovirt on irc.oftc.net

Thank you