

# Plugin API

## Functions

### **get\_versions()**

Returns all versions of Hadoop that could be used with the plugin. It is responsibility of the plugin to make sure that for each version it has available required images, configs, whatever else it needs to create the Hadoop cluster.

*Returns:* list of strings - Hadoop versions

*Example Return Value:* ("Apache Hadoop 1.1.1", "CDH 3", "HDP 1.2")

### **get\_configs(hadoop\_version)**

Lists all configs supported by plugin with descriptions, defaults and node process for which this config is applicable.

*Returns:* list of **configs**

*Example Return Value:* (("heap size", "512", true, "jt"))

### **get\_supported\_node\_types(hadoop\_version)**

Lists of all supported NodeTypes for a given Hadoop version.

*Returns:* list of strings - node types

*Example Return Value:* ("mgmt", "jt+nn", "tt+dn")

### **validate\_cluster(cluster\_description)**

Validates **user\_inputs** in a given **cluster\_description**. Returns empty list if all inputs are correct. Otherwise for each incorrect input function should return **validation\_error** with meaningful content.

For each **user\_input** the function must check that it is applied to the correct Node Type, i.e. that corresponding **config.applicable\_node\_processes** and Node Type has one of the node processes in common. Also function should check that the provided value is correct for the given **config**.

*Returns:* list of **validation\_errors**

*Example Return Value:* (("heap size"), ("Heap size is a required field and must be specified"), ("mapred.task.timeout", "The parameter must be int"))

### **get\_infra(cluster\_description)**

Gets specifications for VMs required for cluster. Basically cluster description already contains specification for cluster nodes - number of VMs for each flavor. And in general case, plugin should just copy specs from **cluster\_description** to its **vm\_requests**, but that is not a strict rule. When returning **vm\_requests**, plugin

- must specify image for VMs
- could change VMs specs in any way it needs. For instance, plugin can ask for additional VMs for the management tool.

*Returns:* list of **vm\_requests**

*Example Return Value:* [(jt+nn", "m1.medium", "ubuntu-12.10", 1, {"set\_root\_pswd"="qwerty", "generate\_keys"=True), ("tt+dn", "m1.small", "ubuntu-12.10", 10, {"set\_root\_pswd"="qwerty", "generate\_keys"=True}]

### **configure\_cluster(cluster\_description, vm\_specs)**

Configures cluster on provisioned by savanna VMs. In this function plugin should perform all actions like adjusting OS, installing required packages (including Hadoop, if needed), configuring Hadoop, etc.

*Returns:* None

### **start\_cluster(cluster\_description, vm\_specs)**

Start already configured cluster. This method is guaranteed to be called only on cluster which was already prepared with **configure\_cluster(...)** call.

*Returns:* None

### **scale\_cluster(cluster\_description, remaining\_vm\_specs, new\_vm\_specs, delete\_vm\_specs)**

Scales cluster - adds/removes nodes to/from active cluster. This function should configure new nodes and attach them to cluster. Additionally here plugin can make some cleanup on VMs that will be deleted.

*Returns:* None

### **on\_terminate\_cluster(cluster\_description, vm\_specs)**

When user terminates cluster, Savanna simply shuts down all the cluster VMs. This method is guaranteed to be invoked before that, allowing plugin to do some clean-up.

*Returns:* None

## Objects

All fields are strings unless specified otherwise

### **config**

Specifies a single key-value pair. Both key and value are strings.

config\_name  
default\_value  
is\_optional: boolean  
applicable\_node\_processes: list of strings

### **user\_input**

Value provided by user for a specific config.

config\_name  
value

### **vm\_group\_description**

Specifies group of VMs within a cluster.

node\_type  
flavor  
configs: list of **user\_inputs**  
count: int

### **cluster\_description**

Contains all relevant info about cluster. This object is provided to the plugin for both cluster creation and scaling.

cluster\_name  
cluster\_configs: list of **user\_inputs**  
hadoop\_version  
vm\_groups: list of **vm\_group\_descriptions**

It is guaranteed that cluster contains exactly one group with node\_type == 'mgmt'. It is also guaranteed that cluster either contains one group with node\_type == 'jt+nn', or one group 'jt' and one group 'nn'. In all these cases group's count is exactly 1.

### **validation\_error**

Describes what is wrong with one of the values provided by user.

config\_name

error\_message

### **vm\_request**

Specifies VMs which plugin requests from Savanna.

flavor

image

count: int

configs: list of **configs**

### **vm\_specs**

Specifies VMs created by Savanna for plugin

flavor

image

count: int

configs: list of **configs**

servers: list of **servers**

### **server**

Specifies a single VM created for cluster

ip

compute\_host

auth type (password / key) ?

credentials ?

## **Image Registry API**

A component helping plugin find suitable images by some criteria. All search is based on tags. A tag is just a string. Each image could have several tags simultaneously.

### **set\_description(image, os\_description, hadoop\_version, extra)**

Sets human-readable information for image, for example, "Ubuntu 13.04 x86\_64, Apache Hadoop 1.1.1, Java 1.7u21"

*Returns:* None

**tag\_image(image, tags)**

Adds tags to image

*Returns:* None

**untag\_image(image, tags)**

Removes tags from image

*Returns:* None

**get\_image\_tags(image)**

Queries all tags for the given image

*Returns:* list of strings - image tags

**get\_image\_by\_tags(image, tags)**

Queries images having all of the specified tags

*Returns:* list of strings - images' ids

## VM Manager

A pack of low-level helpers to help plugin interact with vms

**execute(command)**

Executes command in shell on VM via ssh (non-interactive)

**copy\_to\_vm(filename)**

Copies file to VM via ssh

**copy\_from\_vm(filename)**

Copies file from vm via ssh

Additionally we are thinking about adding some helpers for some high-level actions, for example:  
install package (using apt-get/yum ?)