GETTING STARTED GUIDE:
RED HAT ENTERPRISE LINUX ON AMAZON EC2

INTRODUCTION

Red Hat Enterprise Linux on Amazon EC2 is the certified and supported operating platform purchasable by the hour—delivered by Amazon and supported by Red Hat. As this is Red Hat Enterprise Linux, most operational and management tasks are the same or very similar to managing Red Hat Enterprise Linux for on-premise deployments. Basic knowledge of management tasks is recommended, as well as afamiliarity with the Red Hat Enterprise Linux System Administrator Guide (→ http://www.redhat.com/docs/manuals/enterprise/RHEL-4-Manual/en-US/System_Administration_Guide/) or Red Hat Enterprise Linux Deployment Guide (→ http://www.redhat.com/docs/manuals/enterprise/RHEL-5-manual/en-US/RHEL510/Deployment_Guide/index.html) as well as different courses offered by Red Hat on the subject.

This guide concentrates on differences specific to operating within the Amazon EC2 (Elastic Compute Cloud). Only basic EC2 operations are explained in this document, if you want to learn more and understand the full EC2 API, read the EC2 technical documentation, which is available in the Resource Center at Amazon Web Services, (→ www.amazon.com/aws).

Most major management tasks, for example running and stopping instances, are simplified if you install and use the EC2 Firefox extension that is available through the Amazon Web Service Resource Center. Although not required for operation, many of the operations available within the EC2 API tools (for example "ec2run") are also available and easily managed via the EC2 Firefox extension.

EC2 CONSTRAINTS

One important feature of EC2 based virtual instances during the beta period is that they are always **stateless.** This means that whenever you restart an instance, it picks up the state of the underlying AMI (Amazon Machine Image). For many tasks suitable for virtual instances, this is a favorable model as you always know exactly what is the initial state of an instance, and can simplify the management of multiple servers processing same tasks or offering same services.

However, stateless nature **requires you to closely follow the management procedure** described below, to make sure that you don't lose any configuration changes or updates and additional installations that you might perform.

Additionally, Red Hat Network (RHN) stores the profile of all connected servers. If you restart an instance of your profile on RHN, the corresponding profile becomes invalid and unusable. More about this requirement and the reasoning behind it is described in the Red Hat NetworkChannel Management Guide (→ https://rhn.redhat.com/rhn/help/channel-mgmt/).

While running instances themselves do not contain any persistent storage, Amazon offers the process of rebundling, which provides a method for storing the data and configuration of an instance and makes it available for multiple future executions of a virtual instance. Rebundling utilizes Amazon's S3 (Simple Storage Service). The management process described below utilizes rebundling to store the configuration of an instance after managing the instance with tools such as RHN.

The illustration on the next page shows the flowchart of the process you need to follow when managing Enterprise Linux Hosted images and instances.

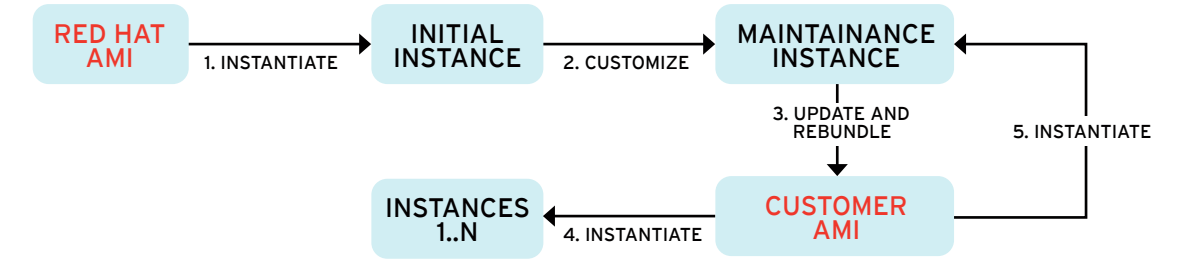

GETTING STARTED OVERVIEW

The following steps presume that you have purchased a subscription to Red Hat Enterprise Linux on Amazon EC2, you have access to the AMI's that are provided by the subscription, have activated your Red Hat Network entitlement, and have accounts at both Amazon Web Services and Red Hat Network. If you have not completed these steps – please review the process at www.redhat.com/solutions/cloud.

- Start & Configuring the Initial Instance
- Storing the Configured Image
- Future Configuration and Maintenance of Images

START & CONFIGURING THE INITIAL INSTANCE

Using 'ec2run', start a new instance and wait for it to become available. Once the image is up and running, you can connect to it using an ssh client.

On the first log in, you are presented with a configuration interface, similar to the one used in the first boot step on physical servers. Use it to configure your virtual instance as well as register for RHN access. Once the first boot step is done, you can update the instance and install any additional required software, as well as perform any additional configurations.

Note: do not restart the instance before completing the next step.

STORING THE CONFIGURED IMAGE

Once you're satisfied with the state of the running instance, you can prepare for restart. To avoid losing all configuration changes to the instantiated AMI, create your own AMI, using the provided **ec2-rhel-bundle** script, which should be accessible from the root command prompt. This script creates your custom image, storing it in S3 persistent storage, as well as disabling RHN access for any regular virtual instances that you might start in the future. Use the resulting custom image for all future virtual instances.

The script requires you to provide your account number, your access keys, your private key, and your certificate. All of this information can be located at the "AWS Access Identifiers" section of the Amazon Web Services web site. The private key and certificate are RSA encoded files that need to be uploaded to the instance via scp. Amazon suggests you place these files into /tmp or /mnt so they are not bundled into your maintenance image.

FUTURE CONFIGURATION AND MAINTENANCE OF IMAGES

As previously mentioned, if you update or somehow change a running instance, these changes are lost after the instance restarts. Furthermore, if you perform an update, you risk making your RHN profile unusable in the future. Therefore, the proper way to do those changes is to designate one running instance as the maintenance instance, and perform all tasks on it. After completion, build a new custom image, as described above, and use it for all your future virtual instances. To designate an instance as maintenance, use the provided script **ec2-rhel-maintenance**, which is also accessible from the root command prompt inside the virtual instance. The script is very simple, and it enables the previously disabled RHN access. It is highly recommended that you restart your running virtual instances using the newly built image, to propagate the changes made in the maintenance instance.

KERNEL UPDATE NOTICE

EC2 currently doesn't permit changing kernels associated with the available image, rather a new AMI is provided that is associated with appropriate kernels.

Part of the service of Red Hat Enterprise Linux Hosted is to provide new AMIs whenever a relevant updated kernel is available. During the BETA period, whenever a new kernel becomes available to Red Hat Enterprise Linux through RHN, Red Hat Enterprise Linux Hosted Subscribers also receive access to a new AMI contained within EC2. While users of the EC2 provided AMIs do **not** have the ability to upgrade kernels of running instances or stored AMIs, customers are encouraged to migrate to the latest version of the AMI whenever a kernel and resulting new AMI is available.