

# Atlas On Demand Use

## Cloud Execution Environment

### OPERATING INSTRUCTIONS

**Copyright**

© Ericsson AB 2016–2018. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

**Disclaimer**

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

**Trademark List**

All trademarks mentioned herein are the property of their respective owners. These are shown in the document Trademark Information.



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Description	1
1.2	Prerequisites	1
<b>2</b>	<b>Atlas On Demand</b>	<b>2</b>
2.1	Deployment	2
2.2	Check Application VM	3
2.3	Create Backup	4
2.4	Delete Atlas VM	5
2.5	Verify Application VM	5
<b>3</b>	<b>Restore Atlas VM</b>	<b>5</b>





# 1 Introduction

Single server configuration of Cloud Execution Environment (CEE) requires low amount of management resources, this makes it possible to disable the Atlas virtual machine (VM) after tenant VMs are deployed.

This document describes how to disable Atlas in a single server CEE installation.

**Note:** In this document, memory and storage quantities are represented according to the JESD100B.01 standard:

- KB refers to  $2^{10}$  bytes
- MB refers to  $2^{20}$  bytes
- GB refers to  $2^{30}$  bytes

## 1.1 Description

Single server is a special, low cost platform of CEE. It is installed on a single compute host server that runs:

- Single vCIC with low resource consumption, including vCPU cores, memory huge pages, local storage
- Tenant VMs
- Atlas for the time of tenant VM provisioning

This CEE platform does not include any external storage connectivity.

Set up Atlas On Demand by deploying an Atlas VM and VNF. Create a backup of the Atlas VM to save the Atlas VM configuration. If Atlas is no longer required, delete the Atlas VM or image.

**Note:** If Heat is used to deploy an application, the state is saved in the Atlas database. If OpenStack command-line client is used to manage (stop, delete) resources, the database in Atlas becomes unsynchronized and Atlas cannot manage these resources.

## 1.2 Prerequisites

### 1.2.1 Documents

Before starting the procedure, ensure that the following documents are available:

- Atlas SW Installation
- Atlas Backup



- Atlas Restore

### 1.2.2 Tools and Equipment

The following tools and equipment are needed:

- Access to the CEE REST APIs
- Secure Shell (SSH) logon to CEE

### 1.2.3 Conditions

Before starting this procedure, ensure that the following conditions are met:

- Credentials for logging in to the host are available.
- No other maintenance activities are taking place at the same time.

## 2 Atlas On Demand

This section describes how Atlas On Demand can be achieved.

### 2.1 Deployment

Deploy an Atlas VM by following the instructions in [Atlas SW Installation](#).

Use an Atlas image file with the following ID syntax:

`ecs-atlas-x86_64-<target_atlas_version>-<build_number>.qcow2`  
and the `<name_of_atlas>` variable, **Atlas\_for\_backup**.

For a minimum configuration, see: Section 2.1.1 on page 2.

Deploy an application using Atlas.

#### 2.1.1 Atlas Minimum Configuration

Atlas can be launched with a minimum configuration flavor of

Disk: 10 GB

RAM: 1024 MB

VCPU: 2

Ephemeral disk/Attached volume: 10 GB



**Note:** Although the minimum size of ephemeral disk is 10 GB, it must be proportional to the size of uploaded images and packages.

Images and packages are stored in /tmp, which is a fixed 30% of the overall disk. So the ephemeral disk must have the size about 4 times bigger than the images and packages altogether.

Update the following parameters in localrc:

# Flavor specifications

Disk=10

RAM=1024

VCPU=2

# Size of data volume

readonly DATA\_IMAGE\_SIZE=10

DATA\_IMAGE\_SIZE is defined per use (size of application, or image to upload). Sizes are divided in percentage of total DATA\_IMAGE\_SIZE attached:

```
/dev/data-vg/internal-data-lv 15%
/dev/data-vg/log-collector 10%
/dev/data-vg/var/core 10%
/dev/data-vg/var/archives 10%
/dev/data-vg/var/lib/mysql 10%
/dev/data-vg/var/log 10%
/dev/data-vg/tmp 30%
/dev/data-vg/var/tmp 5%
```

## 2.2 Check Application VM

Check the application from the deployed stack by using **nova list**, **ovft capp-list**, and **heat stack-list**.

```
atlasadm@atlas:~$ ovft capp-list
```

id	name	status	type	created_at
8251da77-332f-41b5-8f60-887f99513419	package	active	ovf	2016-02-12T06:26:03.000000

```
atlasadm@atlas:~$ heat stack-list
```

id	stack_name	stack_status	creation_time	updated_time
a5f047d8-75ba-4fc9-b93d-f73220e0a85a	stack	CREATE_COMPLETE	2018-03-20T07:03:29Z	2018-03-20T12:05:22Z



```
atlasadm@atlas:~$ nova list
```

ID	Name	Status	Task State	Power State	Networks
6845f0f2-6b4b-42ed-bff5-1883a006347c	Atlas	ACTIVE	None	Running	tenant_3583=10.33.190.38
38a57214-c639-4e21-9042-91542052a6e1	Controller node1	ACTIVE	None	Running	BP_1=127.3.0.5
ebbf9bc7-eda7-4c9b-8d48-b87c798a431a	Payload node1	ACTIVE	None	Running	BP_2=127.13.0.6

## 2.3 Create Backup

Create a backup of the Atlas configuration and store in local storage. See [Atlas Backup](#) for more information. To upload the backup to local storage do the following:

1. Switch user from atlasadm to root with the below command:  
atlasadm@atlas:~ \$ **sudo -i**
2. Ensure that the current directory is set to /var/archives:  
  
root@atlas:~ # **cd /var/archives**
3. Verify that the created backup directory and files are available:  
root@atlas:/var/archives# **ls <atlas\_backup\_directory>**

An example of the command is:

```
root@atlas:/var/archives# ls AtlasBackup1472730328
```

An example of the printout is:

```
AtlasBackup.1472730328-all-mysql-databases.sql.bz2.enc
AtlasBackup.1472730328-root.master.tar.gz.enc
AtlasBackup.1472730328-etc-puppet-hiera-data-passwords.yaml.master.tar.gz.enc
AtlasBackup.1472730328.sha256.enc
AtlasBackup.1472730328-home-atlasadm.master.tar.gz.enc
```

**Note:** The backup consists of the above listed files.

4. Upload the backup for a specified ID:

```
root@atlas:~# scp -r *<id>* <user@laptop>:<storage_path>
```





**Note:** This step is performed if Swift is not present. If Swift is present, the created backup is automatically uploaded to Swift storage.

## 2.4 Delete Atlas VM

Deleting the Atlas VM saves CPU, memory and disk size, deleting the Atlas image saves storage space.

Delete the Atlas VM using the following command:

```
nova delete Atlas_for_backup
```

Delete the Atlas image using the following command:

```
root@cic-1:~# rm /root/artifacts/ecs-atlas-x86_64-<target_atlas_version>-<build_number>.qcow2
```

## 2.5 Verify Application VM

Verify that the application is running and the resources previously allocated to the Atlas VM are free:

```
root@cic-1:~$ nova list
```

ID	Name	Status	Task State	Power State	Networks
38a57214-c639-4e21-9042-91542052a6e1	Controller node1	ACTIVE	None	Running	BP_1=127.3.0.5
ebbf9bc7-eda7-4c9b-8d48-b87c798a431a	Payload node1	ACTIVE	None	Running	BP_2=127.13.0.6

To verify free space after deleting the Atlas VM, check the `nova-compute.log` for available resources.

## 3 Restore Atlas VM

The Atlas VM configuration can be restored using the previously created backup.



There is no need to restore the Atlas VM if Atlas was shut down and not deleted. The image and atlas\_install.sh transferred to the vCIC in the original installation can be used if it is still available in the vCIC.

1. Deploy an Atlas VM in the same way as in Section 2.1 on page 2.
2. Restore the configuration of the Atlas VM from the backup stored in local storage by following the instructions in [Atlas Restore](#).
3. Check the application as described in Section 2.2 on page 3 and compare the results.