



Deployments HowTo: VM and Elasticsearch&Kibana

hands-on

Alessandro Costantini

alessandro.costantini@cnae.infn.it

VM@INFN-Cloud



- **VM**

The solution described in this guide consists on instantiation of a Virtual Machine instantiated on INFN-CLOUD infrastructure.

The instantiation of a VM comes with the responsibility of **maintaining it and all the services it hosts.**

Please read the [INFN Cloud AUP](#) in order to understand the responsibilities you have in managing this service.

- **Prerequisites**

The user has to be registered in the INDIGO-IAM system for INFN-Cloud
Sys-Admin nomination required

VM: How to deploy



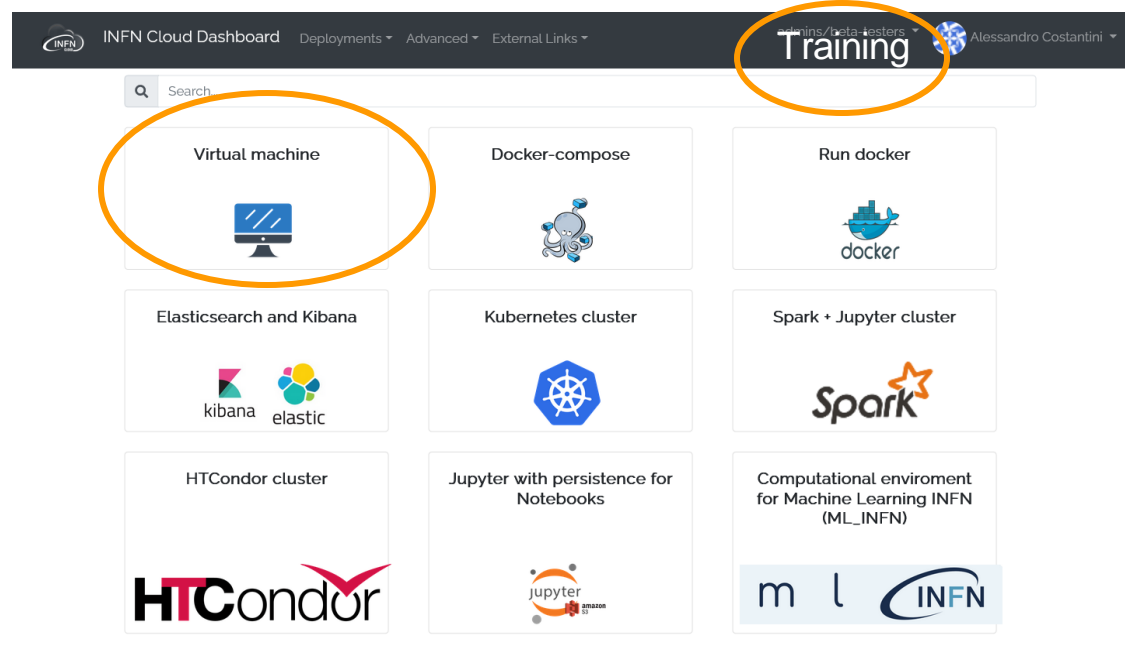
- **Step 1**

Authenticating to the INFN-Cloud dashboard <https://my.cloud.infn.it/>

- **Step 2**

Select the **VM deployment** and the right **group**

Use Chrome




Welcome to infn-cloud

Sign in with



Not a member?



VM: basic configuration



Step 3

- **Select storage**
 - VM with no additional storage
 - VM with block storage

Select

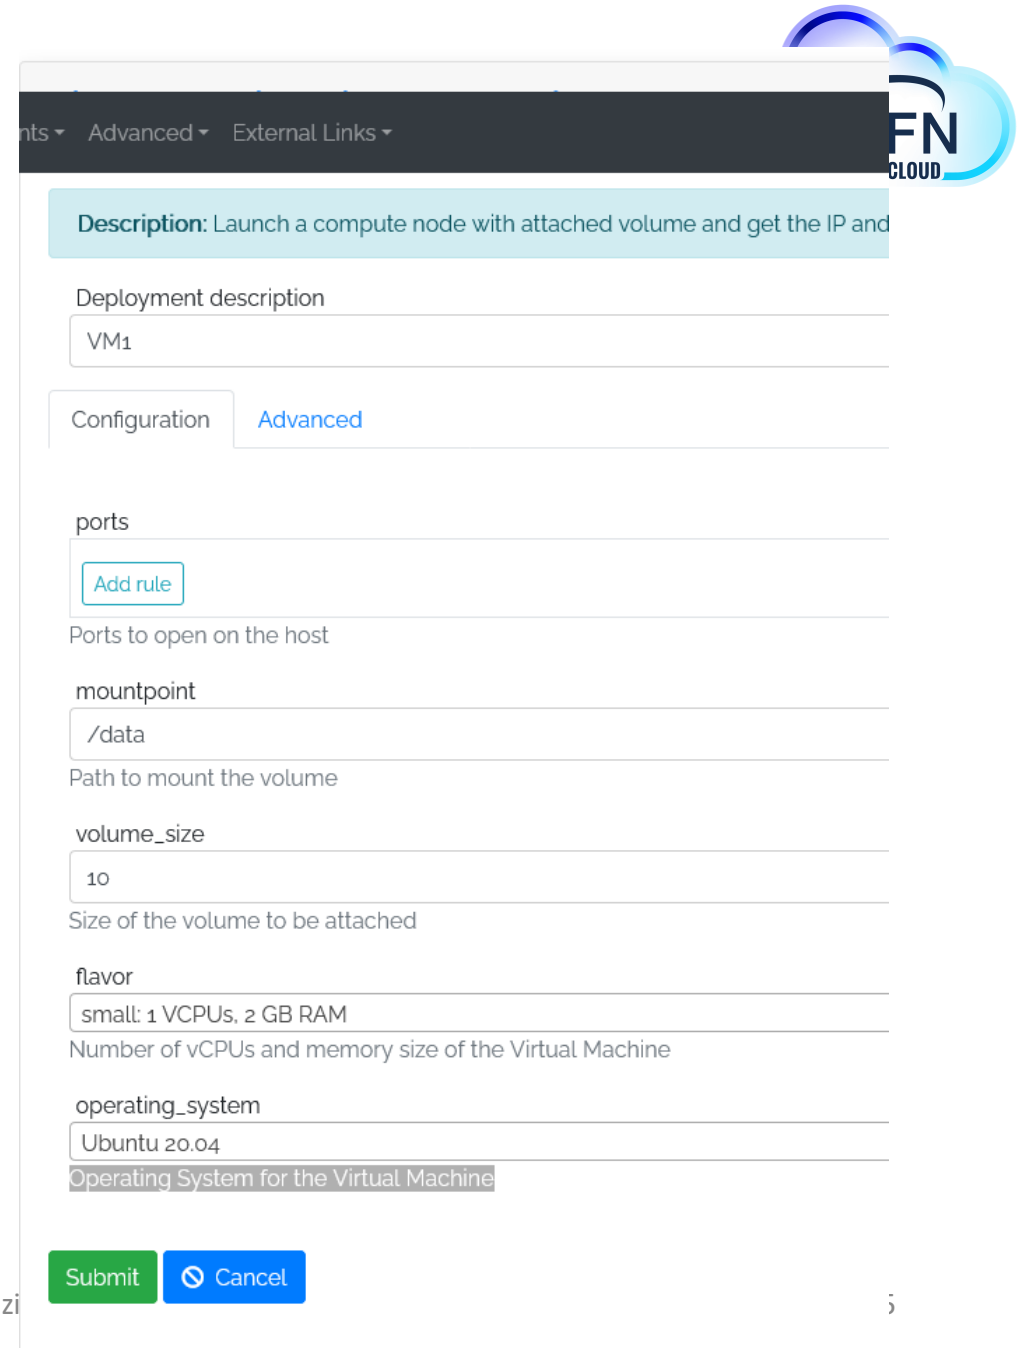
VM with no additional storage

VM with block storage
Attach a volume to the machine

VM: basic configuration

Step 4

- **Deployment description**
 - A brief description of the deployment
- **Configuration TAB**
 - Ports
 - Service ports to be open on the new VM.
Only the ssh port (22) is opened by default
 - `volume_size`
 - Size of the volume to be used to store the data (default 10GB)
 - `mountpoint`
 - Path to mount the data volume (default `/data`)
 - `volume_size`
 - Size of the volume to be used to store the data (default 10GB)
 - `flavor`
 - Number of vCPUs and memory size of the Virtual Machine
 - `operating_system`
 - Operating System for the Virtual Machine



The screenshot shows a web interface for configuring a VM. At the top right, there is a logo for 'FN CLOUD'. Below it, there are navigation tabs: 'Configuration' (selected) and 'Advanced'. The main content area is titled 'Description: Launch a compute node with attached volume and get the IP and...'. Below this, there are several input fields:

- Deployment description:** A text input field containing 'VM1'.
- ports:** A section with an 'Add rule' button and a label 'Ports to open on the host'.
- mountpoint:** A text input field containing '/data' with the label 'Path to mount the volume'.
- volume_size:** A text input field containing '10' with the label 'Size of the volume to be attached'.
- flavor:** A dropdown menu showing 'small: 1 VCPUs, 2 GB RAM' with the label 'Number of vCPUs and memory size of the Virtual Machine'.
- operating_system:** A dropdown menu showing 'Ubuntu 20.04' with the label 'Operating System for the Virtual Machine'.

At the bottom, there are two buttons: a green 'Submit' button and a blue 'Cancel' button.

VM: Advanced configuration



Step 5

Advanced TAB

Scheduling

Automatic (Default)

Manual

The following extra-settings can be set

Deployment creation timeout (minutes)

Do not delete the deployment in case of failure

Send a confirmation email when complete

Virtual machine with block device

Description: Launch a compute node with attached volume and get th

Deployment description

VM1

Configuration

Advanced

Configure scheduling:

Auto Manual

Set deployment creation timeout (minutes) 720

Do not delete the deployment in case of failure

Send a confirmation email when complete

Submit

Cancel

VM: Operate with the VM deployment



- **Step 6**

Once successful deployment completion, the user can reach the deployed VM configuration and endpoint

The screenshot displays the INFN Cloud Dashboard interface. At the top, the navigation bar includes 'INFN Cloud Dashboard', 'Deployments', 'Advanced', and 'External Links'. The user is logged in as 'admins/beta-testers' with the name 'Alessandro Costantini'. The main section is titled 'My deployments' and features a 'Refresh' button and a '+ New deployment' button. Below this, there is a search bar and a table of deployments. The table has columns for 'Description', 'Deployment identifier', and 'Status'. One deployment, 'VM1', is highlighted with an orange circle. Its deployment identifier is '11ed326d-4576-2ca9-b185-0242a79ac9f5' and its status is 'CREATE_COMPLETE'. To the right of the table, a detailed view for the selected deployment is shown, also circled in orange. It includes a 'Back' button and tabs for 'Overview', 'Input values', and 'Output values'. The 'Overview' tab is active, showing the following details: 'STATUS: CREATE_COMPLETE', 'CREATED AT: 2022-09-12 07:34:00', 'UPDATED AT: 2022-09-12 07:38:00', and 'DEPLOYED AT: BACKBONE-CNAF'.

| Description | Deployment identifier | Status |
|-------------|--------------------------------------|-----------------|
| VM1 | 11ed326d-4576-2ca9-b185-0242a79ac9f5 | CREATE_COMPLETE |

11ed326d-4576-2ca9-b185-0242a79ac9f5

← Back

Description: VM1

Overview Input values Output values

STATUS: CREATE_COMPLETE

CREATED AT: 2022-09-12 07:34:00

UPDATED AT: 2022-09-12 07:38:00

DEPLOYED AT: BACKBONE-CNAF

VM: Operate with the VM deployment



- **Step 7**

VM can be accessed by using the information available in the “Input values” and “Output values” fields of the deployment

The VM is now available

Full guide available at
- [VM_deployment@INFN-Cloud](#)

A screenshot of a web interface showing the details of a VM deployment. At the top, there is a header bar with the deployment ID '11ed326d-4576-2ca9-b185-0242a79ac9f5' and a 'Back' button. Below the header, the 'Description' is 'VM1'. There are three tabs: 'Overview' (selected), 'Input values', and 'Output values'. The main content area lists several parameters: 'additional_description: VM1', 'mem_size: 2 GB', 'mountpoint: /data', 'num_cpus: 1', 'os_distribution: ubuntu', 'os_version: 20.04', and 'users: [{"os_user_add_to_sudoers": True, "os_user_name": "acostantini", "os_user_ssh_public_key": "ssh-rsa'}]'.

11ed326d-4576-2ca9-b185-0242a79ac9f5 ← Back

Description: VM1

[Overview](#) [Input values](#) [Output values](#)

additional_description: VM1

mem_size: 2 GB

mountpoint: /data

num_cpus: 1

os_distribution: ubuntu

os_version: 20.04

users: [{"os_user_add_to_sudoers": True, "os_user_name": "acostantini", "os_user_ssh_public_key": "ssh-rsa"}]

Elasticsearch and Kibana @INFN-Cloud



- **Elasticsearch and Kibana**

Elasticsearch and **Kibana** are open source software

Elasticsearch is a search and analytics engine

Kibana lets users visualize data with charts and graphs in Elasticsearch

Missing

Logstash is a server-side data processing pipeline that ingests data from multiple sources simultaneously, transforms it, and then sends it to a "stash" like Elasticsearch

Elasticsearch and Kibana @INFN-Cloud



- **Enabled by INFN-Cloud**

Deploy a **virtual machine** pre-configured with the **Elasticsearch** search and analytics engine and with **Kibana** for simple visualization of data with charts and graphs in Elasticsearch

The instantiation of a VM comes with the responsibility of **maintaining it and all the services it hosts.**

Please read the [INFN Cloud AUP](#) in order to understand the responsibilities you have in managing this service.

- **Prerequisites**

The user has to be registered in the IAM system for INFN-Cloud

Sys-Admin nomination required

E&K: How to deploy



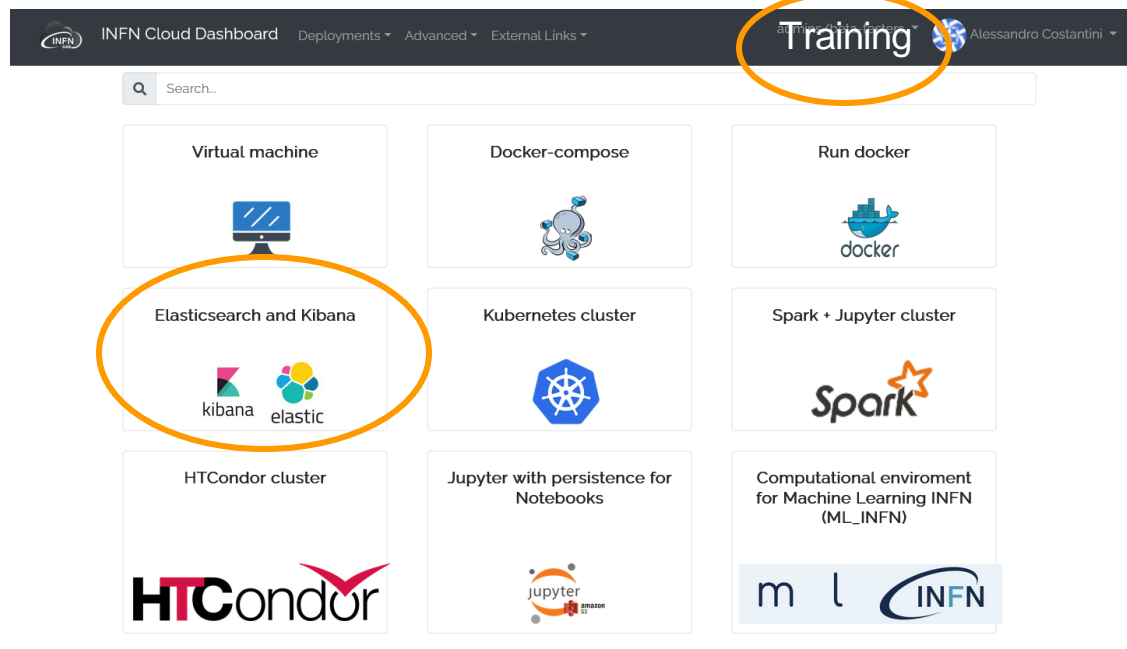

- **Step 1**

Authenticating to the INFN-Cloud dashboard <https://my.cloud.infn.it/>

- **Step 2**

Select the **EK deployment** and the right **group**

Use Chrome



Welcome to infn-cloud

Sign in with



Not a member?

Apply for an account

E&K: basic configuration



Step 3

- **Deployment description**
 - A brief description of the deployment
- **Configuration TAB**
 - contact_email
 - elastic_password
 - This password is created for the username elastic that allows you to log in to the dashboard.
 - kibana_password
 - This is a specific password for the internal user kibana_system. Kibana uses the connect and communicate with Elasticsearch
 - volume_size
 - Size of the volume to be used to store the data (default 10GB)
 - mountpoint
 - Path to mount the data volume (default /data)
 - flavor
 - Number of vCPUs and memory size of the Virtual Machine

INFN Cloud Dashboard Deployments ▾ Advanced ▾ External Links ▾

Elasticsearch and Kibana (version 8.1.3)

Description: Deploy a virtual machine pre-configured with the Elasticsearch search and analytics

Deployment description

EK1

Configuration **Advanced**

contact_email

acostantini@infn.it

Insert your Email for receiving notifications

elastic_password

.....

Password for user elastic

kibana_password

.....

Password for user kibana_system (internal user)

volume_size

10

Size of the volume to be used to store the data

mountpoint

/data

Path to mount the data volume

flavor

medium: 2 vCPUs, 4 GB RAM

Number of vCPUs and memory size of the Virtual Machine

Submit Cancel

E&K: Advanced configuration



Step 4

Advanced TAB

Scheduling

Automatic (Default)


Manual

The following extra-settings can be set

Deployment creation timeout (minutes)

Do not delete the deployment in case of failure

Send a confirmation email when complete



Elasticsearch and Kibana (version 8.1.3)

Description: Deploy a virtual machine pre-configured with the Elasticsearch

Deployment description

EK1

Configuration

Advanced

Configure scheduling:

Auto Manual

Set deployment creation timeout (minutes) 720

Do not delete the deployment in case of failure

Send a confirmation email when complete

Submit

Cancel

E&K: Operate with E&K deployment



- **Step 5**

Once successful deployment completion, the user can reach the deployed E&K configuration and endpoints

The screenshot displays the INFN Cloud Dashboard interface. At the top, the navigation bar includes 'INFN Cloud Dashboard', 'Deployments', 'Advanced', and 'External Links'. The user is logged in as 'admins/beta-testers' and 'Alessandro Costantini'. The main section is titled 'My deployments' and features a 'Refresh' button and a '+ New deployment' button. Below this, there is a table of deployments with columns for 'Description', 'Deployment identifier', 'Status', 'Creation time', 'Deployed at', and 'Actions'. One deployment is highlighted with an orange circle: 'EK1' with identifier '11ed31a4-338f-99c6-b185-0242a79ac9f5' and status 'CREATE_COMPLETE'. To the right, a detailed view of this deployment is shown, including a 'Back' button and tabs for 'Overview', 'Input values', and 'Output values'. The 'Output values' tab is selected, showing the following configuration and endpoints, which are also circled in orange:

- node_ip: 192.135.24.34
- kibana_username: elastic
- elasticsearch_endpoint: <https://elastic.192.135.24.34.myip.cloud.infn.it>
- kibana_endpoint: <https://kibana.192.135.24.34.myip.cloud.infn.it>
- ssh_account: acostantini

E&K: Operate with E&K deployment



- **Step 6**

Elasticsearch and Kibana dashboard that can be accessed by using the credential defined during the deployment configuration (**step 3**)

The operative E&K platform is now available

Full guide available at
- [Elasticsearch&Kibana](#)

