



Muestras de código Python

ONTAP Select

NetApp
May 09, 2024

Tabla de contenidos

- Muestras de código Python 1
 - Script para crear un clúster 1
 - JSON para el script a fin de crear un clúster 8
 - Script para añadir una licencia de nodo. 12
 - Script para eliminar un clúster 16
 - Módulo de soporte común 18
 - Script para cambiar el tamaño de los nodos del clúster 22

Muestras de código Python

Script para crear un clúster

Puede utilizar el siguiente script para crear un clúster basado en los parámetros definidos en el script y en un archivo de entrada JSON.

```
#!/usr/bin/env python
##-----
#
# File: cluster.py
#
# (C) Copyright 2019 NetApp, Inc.
#
# This sample code is provided AS IS, with no support or warranties of
# any kind, including but not limited for warranties of merchantability
# or fitness of any kind, expressed or implied. Permission to use,
# reproduce, modify and create derivatives of the sample code is granted
# solely for the purpose of researching, designing, developing and
# testing a software application product for use with NetApp products,
# provided that the above copyright notice appears in all copies and
# that the software application product is distributed pursuant to terms
# no less restrictive than those set forth herein.
#
##-----

import traceback
import argparse
import json
import logging

from deploy_requests import DeployRequests

def add_vcenter_credentials(deploy, config):
    """ Add credentials for the vcenter if present in the config """
    log_debug_trace()

    vcenter = config.get('vcenter', None)
    if vcenter and not deploy.resource_exists('/security/credentials',
                                              'hostname', vcenter
                                              ['hostname']):
        log_info("Registering vcenter {} credentials".format(vcenter
                                                              ['hostname']))
        data = {k: vcenter[k] for k in ['hostname', 'username',
```

```

'password']]
    data['type'] = "vcenter"
    deploy.post('/security/credentials', data)

def add_standalone_host_credentials(deploy, config):
    """ Add credentials for standalone hosts if present in the config.
        Does nothing if the host credential already exists on the Deploy.
    """
    log_debug_trace()

    hosts = config.get('hosts', [])
    for host in hosts:
        # The presense of the 'password' will be used only for standalone
        hosts.
        # If this host is managed by a vcenter, it should not have a host
        'password' in the json.
        if 'password' in host and not deploy.resource_exists
        ('/security/credentials',
                                                'hostname',
        host['name']):
            log_info("Registering host {} credentials".format(host[
            'name']))
            data = {'hostname': host['name'], 'type': 'host',
                    'username': host['username'], 'password': host
            ['password']}
            deploy.post('/security/credentials', data)

def register_unkown_hosts(deploy, config):
    ''' Registers all hosts with the deploy server.
        The host details are read from the cluster config json file.

        This method will skip any hosts that are already registered.
        This method will exit the script if no hosts are found in the
        config.
    '''
    log_debug_trace()

    data = {"hosts": []}
    if 'hosts' not in config or not config['hosts']:
        log_and_exit("The cluster config requires at least 1 entry in the
        'hosts' list got {}".format(config))

    missing_host_cnt = 0
    for host in config['hosts']:

```

```

        if not deploy.resource_exists('/hosts', 'name', host['name']):
            missing_host_cnt += 1
            host_config = {"name": host['name'], "hypervisor_type": host
['type']}

            if 'mgmt_server' in host:
                host_config["management_server"] = host['mgmt_server']
                log_info(
                    "Registering from vcenter {mgmt_server}".format(**
host))

            if 'password' in host and 'user' in host:
                host_config['credential'] = {
                    "password": host['password'], "username": host[
'user']}

            log_info("Registering {type} host {name}".format(**host))
            data["hosts"].append(host_config)

# only post /hosts if some missing hosts were found
if missing_host_cnt:
    deploy.post('/hosts', data, wait_for_job=True)

def add_cluster_attributes(deploy, config):
    ''' POST a new cluster with all needed attribute values.
        Returns the cluster_id of the new config
    '''
    log_debug_trace()

    cluster_config = config['cluster']
    cluster_id = deploy.find_resource('/clusters', 'name', cluster_config
['name'])

    if not cluster_id:
        log_info("Creating cluster config named {name}".format(
**cluster_config))

        # Filter to only the valid attributes, ignores anything else in
the json
        data = {k: cluster_config[k] for k in [
            'name', 'ip', 'gateway', 'netmask', 'ontap_image_version',
'dns_info', 'ntp_servers']}

        num_nodes = len(config['nodes'])

        log_info("Cluster properties: {}".format(data))

```

```

        resp = deploy.post('/v3/clusters?node_count={}'.format(num_nodes),
data)
        cluster_id = resp.headers.get('Location').split('/')[-1]

        return cluster_id

def get_node_ids(deploy, cluster_id):
    ''' Get the the ids of the nodes in a cluster. Returns a list of
node_ids.'''
    log_debug_trace()

    response = deploy.get('/clusters/{}/nodes'.format(cluster_id))
    node_ids = [node['id'] for node in response.json().get('records')]
    return node_ids

def add_node_attributes(deploy, cluster_id, node_id, node):
    ''' Set all the needed properties on a node '''
    log_debug_trace()

    log_info("Adding node '{}' properties".format(node_id))

    data = {k: node[k] for k in ['ip', 'serial_number', 'instance_type',
                                'is_storage_efficiency_enabled'] if k in
node}
    # Optional: Set a serial_number
    if 'license' in node:
        data['license'] = {'id': node['license']}

    # Assign the host
    host_id = deploy.find_resource('/hosts', 'name', node['host_name'])
    if not host_id:
        log_and_exit("Host names must match in the 'hosts' array, and the
nodes.host_name property")

    data['host'] = {'id': host_id}

    # Set the correct raid_type
    is_hw_raid = not node['storage'].get('disks') # The presence of a
list of disks indicates sw_raid
    data['passthrough_disks'] = not is_hw_raid

    # Optionally set a custom node name
    if 'name' in node:
        data['name'] = node['name']

```

```

log_info("Node properties: {}".format(data))
deploy.patch('/clusters/{}/nodes/{}'.format(cluster_id, node_id),
data)

def add_node_networks(deploy, cluster_id, node_id, node):
    ''' Set the network information for a node '''
    log_debug_trace()

    log_info("Adding node '{}' network properties".format(node_id))

    num_nodes = deploy.get_num_records('/clusters/{}/nodes'.format
(cluster_id))

    for network in node['networks']:

        # single node clusters do not use the 'internal' network
        if num_nodes == 1 and network['purpose'] == 'internal':
            continue

        # Deduce the network id given the purpose for each entry
        network_id = deploy.find_resource(
'/clusters/{}/nodes/{}/networks'.format(cluster_id, node_id),
                                'purpose', network['purpose'])

        data = {"name": network['name']}
        if 'vlan' in network and network['vlan']:
            data['vlan_id'] = network['vlan']

        deploy.patch('/clusters/{}/nodes/{}/networks/{}'.format(
cluster_id, node_id, network_id), data)

def add_node_storage(deploy, cluster_id, node_id, node):
    ''' Set all the storage information on a node '''
    log_debug_trace()

    log_info("Adding node '{}' storage properties".format(node_id))
    log_info("Node storage: {}".format(node['storage']['pools']))

    data = {'pool_array': node['storage']['pools']} # use all the json
properties
    deploy.post(
        '/clusters/{}/nodes/{}/storage/pools'.format(cluster_id, node_id),
data)

    if 'disks' in node['storage'] and node['storage']['disks']:
        data = {'disks': node['storage']['disks']}

```

```

        deploy.post(
            '/clusters/{}/nodes/{}/storage/disks'.format(cluster_id,
node_id), data)

def create_cluster_config(deploy, config):
    ''' Construct a cluster config in the deploy server using the input
json data '''
    log_debug_trace()

    cluster_id = add_cluster_attributes(deploy, config)

    node_ids = get_node_ids(deploy, cluster_id)
    node_configs = config['nodes']

    for node_id, node_config in zip(node_ids, node_configs):
        add_node_attributes(deploy, cluster_id, node_id, node_config)
        add_node_networks(deploy, cluster_id, node_id, node_config)
        add_node_storage(deploy, cluster_id, node_id, node_config)

    return cluster_id

def deploy_cluster(deploy, cluster_id, config):
    ''' Deploy the cluster config to create the ONTAP Select VMs. '''
    log_debug_trace()
    log_info("Deploying cluster: {}".format(cluster_id))

    data = {'ontap_credential': {'password': config['cluster']
['ontap_admin_password']}}
    deploy.post('/clusters/{}/deploy?inhibit_rollback=true'.format
(cluster_id),
                data, wait_for_job=True)

def log_debug_trace():
    stack = traceback.extract_stack()
    parent_function = stack[-2][2]
    logging.getLogger('deploy').debug('Calling %s()' % parent_function)

def log_info(msg):
    logging.getLogger('deploy').info(msg)

def log_and_exit(msg):
    logging.getLogger('deploy').error(msg)

```



```

exit(1)

def configure_logging(verbose):
    FORMAT = '%(asctime)-15s:%(levelname)s:%(name)s: %(message)s'
    if verbose:
        logging.basicConfig(level=logging.DEBUG, format=FORMAT)
    else:
        logging.basicConfig(level=logging.INFO, format=FORMAT)
        logging.getLogger('requests.packages.urllib3.connectionpool')
        .setLevel(
            logging.WARNING)

def main(args):
    configure_logging(args.verbose)
    deploy = DeployRequests(args.deploy, args.password)

    with open(args.config_file) as json_data:
        config = json.load(json_data)

        add_vcenter_credentials(deploy, config)

        add_standalone_host_credentials(deploy, config)

        register_unkown_hosts(deploy, config)

        cluster_id = create_cluster_config(deploy, config)

        deploy_cluster(deploy, cluster_id, config)

def parseArgs():
    parser = argparse.ArgumentParser(description='Uses the ONTAP Select
    Deploy API to construct and deploy a cluster.')
    parser.add_argument('-d', '--deploy', help='Hostname or IP address of
    Deploy server')
    parser.add_argument('-p', '--password', help='Admin password of Deploy
    server')
    parser.add_argument('-c', '--config_file', help='Filename of the
    cluster config')
    parser.add_argument('-v', '--verbose', help='Display extra debugging
    messages for seeing exact API calls and responses',
                        action='store_true', default=False)
    return parser.parse_args()

if __name__ == '__main__':

```

```
args = parseArgs()
main(args)
```

JSON para el script a fin de crear un clúster

Cuando crea o elimina un clúster de ONTAP Select con ejemplos de código Python, debe proporcionar un archivo JSON como entrada del script. Puede copiar y modificar la muestra JSON adecuada en función de sus planes de implementación.

Clúster de un solo nodo en ESXi

```
{
  "hosts": [
    {
      "password": "mypassword1",
      "name": "host-1234",
      "type": "ESX",
      "username": "admin"
    }
  ],
  "cluster": {
    "dns_info": {
      "domains": ["lab1.company-demo.com", "lab2.company-demo.com",
        "lab3.company-demo.com", "lab4.company-demo.com"],
    },

    "dns_ips": ["10.206.80.135", "10.206.80.136"]
  },
  "ontap_image_version": "9.7",
  "gateway": "10.206.80.1",
  "ip": "10.206.80.115",
  "name": "mycluster",
  "ntp_servers": ["10.206.80.183", "10.206.80.142"],
  "ontap_admin_password": "mypassword2",
  "netmask": "255.255.254.0"
},
  "nodes": [
    {
      "serial_number": "3200000nn",
      "ip": "10.206.80.114",
      "name": "node-1",
      "networks": [
        {
```

```

        "name": "ontap-external",
        "purpose": "mgmt",
        "vlan": 1234
    },
    {
        "name": "ontap-external",
        "purpose": "data",
        "vlan": null
    },
    {
        "name": "ontap-internal",
        "purpose": "internal",
        "vlan": null
    }
],
"host_name": "host-1234",
"is_storage_efficiency_enabled": false,
"instance_type": "small",
"storage": {
    "disk": [],
    "pools": [
        {
            "name": "storage-pool-1",
            "capacity": 4802666790125
        }
    ]
}
}
]
}

```

Clúster de un solo nodo en ESXi mediante vCenter

```

{
  "hosts": [
    {
      "name": "host-1234",
      "type": "ESX",
      "mgmt_server": "vcenter-1234"
    }
  ],

  "cluster": {
    "dns_info": { "domains": ["lab1.company-demo.com", "lab2.company-
demo.com"],

```

```

        "lab3.company-demo.com", "lab4.company-demo.com"
    ],
    "dns_ips": ["10.206.80.135", "10.206.80.136"]
},

"ontap_image_version": "9.7",
"gateway": "10.206.80.1",
"ip": "10.206.80.115",
"name": "mycluster",
"ntp_servers": ["10.206.80.183", "10.206.80.142"],
"ontap_admin_password": "mypassword2",
"netmask": "255.255.254.0"
},

"vcenter": {
    "password": "mypassword2",
    "hostname": "vcenter-1234",
    "username": "selectadmin"
},

"nodes": [
    {
        "serial_number": "3200000nn",
        "ip": "10.206.80.114",
        "name": "node-1",
        "networks": [
            {
                "name": "ONTAP-Management",
                "purpose": "mgmt",
                "vlan": null
            },
            {
                "name": "ONTAP-External",
                "purpose": "data",
                "vlan": null
            },
            {
                "name": "ONTAP-Internal",
                "purpose": "internal",
                "vlan": null
            }
        ]
    },
    {
        "host_name": "host-1234",
        "is_storage_efficiency_enabled": false,
        "instance_type": "small",
    }
]

```

```

    "storage": {
      "disk": [],
      "pools": [
        {
          "name": "storage-pool-1",
          "capacity": 5685190380748
        }
      ]
    }
  }
}

```

Clúster de un solo nodo en KVM



- A partir de ONTAP Select 9.10.1, ya no es posible poner en marcha un clúster nuevo en el hipervisor KVM.
- A partir de ONTAP Select 9.11.1, toda la funcionalidad de capacidad de gestión ya no está disponible para los clústeres y hosts KVM existentes, excepto para las funciones de desconectar y eliminar.

```

{
  "hosts": [
    {
      "password": "mypassword1",
      "name": "host-1234",
      "type": "KVM",
      "username": "root"
    }
  ],
  "cluster": {
    "dns_info": {
      "domains": ["lab1.company-demo.com", "lab2.company-demo.com",
        "lab3.company-demo.com", "lab4.company-demo.com"]
    },
    "dns_ips": ["10.206.80.135", "10.206.80.136"]
  },
  "ontap_image_version": "9.7",
  "gateway": "10.206.80.1",
  "ip": "10.206.80.115",
  "name": "CBF4ED97",
  "ntp_servers": ["10.206.80.183", "10.206.80.142"],

```

```

    "ontap_admin_password": "mypassword2",
    "netmask": "255.255.254.0"
  },
  "nodes": [
    {
      "serial_number": "3200000nn",
      "ip": "10.206.80.115",
      "name": "node-1",
      "networks": [
        {
          "name": "ontap-external",
          "purpose": "mgmt",
          "vlan": 1234
        },
        {
          "name": "ontap-external",
          "purpose": "data",
          "vlan": null
        },
        {
          "name": "ontap-internal",
          "purpose": "internal",
          "vlan": null
        }
      ]
    },
    {
      "host_name": "host-1234",
      "is_storage_efficiency_enabled": false,
      "instance_type": "small",
      "storage": {
        "disk": [],
        "pools": [
          {
            "name": "storage-pool-1",
            "capacity": 4802666790125
          }
        ]
      }
    }
  ]
}

```

Script para añadir una licencia de nodo

Se puede usar el siguiente script para añadir una licencia de un nodo ONTAP Select.

```

#!/usr/bin/env python
##-----
#
# File: add_license.py
#
# (C) Copyright 2019 NetApp, Inc.
#
# This sample code is provided AS IS, with no support or warranties of
# any kind, including but not limited for warranties of merchantability
# or fitness of any kind, expressed or implied. Permission to use,
# reproduce, modify and create derivatives of the sample code is granted
# solely for the purpose of researching, designing, developing and
# testing a software application product for use with NetApp products,
# provided that the above copyright notice appears in all copies and
# that the software application product is distributed pursuant to terms
# no less restrictive than those set forth herein.
#
##-----

import argparse
import logging
import json

from deploy_requests import DeployRequests

def post_new_license(deploy, license_filename):
    log_info('Posting a new license: {}'.format(license_filename))

    # Stream the file as multipart/form-data
    deploy.post('/licensing/licenses', data={},
                files={'license_file': open(license_filename, 'rb')})

    # Alternative if the NLF license data is converted to a string.
    # with open(license_filename, 'rb') as f:
    #     nlf_data = f.read()
    #     r = deploy.post('/licensing/licenses', data={},
    #                     files={'license_file': (license_filename,
    nlf_data)})

def put_license(deploy, serial_number, data, files):
    log_info('Adding license for serial number: {}'.format(serial_number))

    deploy.put('/licensing/licenses/{}'.format(serial_number), data=data,
               files=files)

```

```

def put_used_license(deploy, serial_number, license_filename,
ontap_username, ontap_password):
    ''' If the license is used by an 'online' cluster, a username/password
must be given. '''

    data = {'ontap_username': ontap_username, 'ontap_password':
ontap_password}
    files = {'license_file': open(license_filename, 'rb')}

    put_license(deploy, serial_number, data, files)

def put_free_license(deploy, serial_number, license_filename):
    data = {}
    files = {'license_file': open(license_filename, 'rb')}

    put_license(deploy, serial_number, data, files)

def get_serial_number_from_license(license_filename):
    ''' Read the NLF file to extract the serial number '''
    with open(license_filename) as f:
        data = json.load(f)

        statusResp = data.get('statusResp', {})
        serialNumber = statusResp.get('serialNumber')
        if not serialNumber:
            log_and_exit("The license file seems to be missing the
serialNumber")

        return serialNumber

def log_info(msg):
    logging.getLogger('deploy').info(msg)

def log_and_exit(msg):
    logging.getLogger('deploy').error(msg)
    exit(1)

def configure_logging():
    FORMAT = '%(asctime)-15s:%(levelname)s:%(name)s: %(message)s'
    logging.basicConfig(level=logging.INFO, format=FORMAT)

```



```

logging.getLogger('requests.packages.urllib3.connectionpool').
setLevel(logging.WARNING)

def main(args):
    configure_logging()
    serial_number = get_serial_number_from_license(args.license)

    deploy = DeployRequests(args.deploy, args.password)

    # First check if there is already a license resource for this serial-
    number
    if deploy.find_resource('/licensing/licenses', 'id', serial_number):

        # If the license already exists in the Deploy server, determine if
        its used
        if deploy.find_resource('/clusters', 'nodes.serial_number',
        serial_number):

            # In this case, requires ONTAP creds to push the license to
            the node
            if args.ontap_username and args.ontap_password:
                put_used_license(deploy, serial_number, args.license,
                                args.ontap_username, args.ontap_password)
            else:
                print("ERROR: The serial number for this license is in
                use. Please provide ONTAP credentials.")
            else:
                # License exists, but its not used
                put_free_license(deploy, serial_number, args.license)
        else:
            # No license exists, so register a new one as an available license
            for later use
            post_new_license(deploy, args.license)

def parseArgs():
    parser = argparse.ArgumentParser(description='Uses the ONTAP Select
    Deploy API to add or update a new or used NLF license file.')
    parser.add_argument('-d', '--deploy', required=True, type=str, help
    ='Hostname or IP address of ONTAP Select Deploy')
    parser.add_argument('-p', '--password', required=True, type=str, help
    ='Admin password of Deploy server')
    parser.add_argument('-l', '--license', required=True, type=str, help
    ='Filename of the NLF license data')
    parser.add_argument('-u', '--ontap_username', type=str,

```

```

        help='ONTAP Select username with privelege to add
the license. Only provide if the license is used by a Node.')
        parser.add_argument('-o', '--ontap_password', type=str,
                            help='ONTAP Select password for the
ontap_username. Required only if ontap_username is given.')
        return parser.parse_args()

if __name__ == '__main__':
    args = parseArgs()
    main(args)

```

Script para eliminar un clúster

Es posible utilizar el siguiente script de la CLI para eliminar un clúster existente.

```

#!/usr/bin/env python
##-----
#
# File: delete_cluster.py
#
# (C) Copyright 2019 NetApp, Inc.
#
# This sample code is provided AS IS, with no support or warranties of
# any kind, including but not limited for warranties of merchantability
# or fitness of any kind, expressed or implied. Permission to use,
# reproduce, modify and create derivatives of the sample code is granted
# solely for the purpose of researching, designing, developing and
# testing a software application product for use with NetApp products,
# provided that the above copyright notice appears in all copies and
# that the software application product is distributed pursuant to terms
# no less restrictive than those set forth herein.
#
##-----

import argparse
import json
import logging

from deploy_requests import DeployRequests

def find_cluster(deploy, cluster_name):
    return deploy.find_resource('/clusters', 'name', cluster_name)

def offline_cluster(deploy, cluster_id):

```

```

# Test that the cluster is online, otherwise do nothing
response = deploy.get('/clusters/{}'.format(cluster_id))
cluster_data = response.json()['record']
if cluster_data['state'] == 'powered_on':
    log_info("Found the cluster to be online, modifying it to be
powered_off.")
    deploy.patch('/clusters/{}'.format(cluster_id), {'availability':
'powered_off'}, True)

def delete_cluster(deploy, cluster_id):
    log_info("Deleting the cluster({}).".format(cluster_id))
    deploy.delete('/clusters/{}'.format(cluster_id), True)
    pass

def log_info(msg):
    logging.getLogger('deploy').info(msg)

def configure_logging():
    FORMAT = '%(asctime)-15s:%(levelname)s:%(name)s: %(message)s'
    logging.basicConfig(level=logging.INFO, format=FORMAT)
    logging.getLogger('requests.packages.urllib3.connectionpool').
setLevel(logging.WARNING)

def main(args):
    configure_logging()
    deploy = DeployRequests(args.deploy, args.password)

    with open(args.config_file) as json_data:
        config = json.load(json_data)

        cluster_id = find_cluster(deploy, config['cluster']['name'])

        log_info("Found the cluster {} with id: {}".format(config
['cluster']['name'], cluster_id))

        offline_cluster(deploy, cluster_id)

        delete_cluster(deploy, cluster_id)

def parseArgs():
    parser = argparse.ArgumentParser(description='Uses the ONTAP Select
Deploy API to delete a cluster')

```

```

    parser.add_argument('-d', '--deploy', required=True, type=str, help
                        ='Hostname or IP address of Deploy server')
    parser.add_argument('-p', '--password', required=True, type=str, help
                        ='Admin password of Deploy server')
    parser.add_argument('-c', '--config_file', required=True, type=str,
                        help='Filename of the cluster json config')
    return parser.parse_args()

if __name__ == '__main__':
    args = parseArgs()
    main(args)

```

Módulo de soporte común

Todos los scripts de Python utilizan una clase Python común en un único módulo.

```

#!/usr/bin/env python
##-----
#
# File: deploy_requests.py
#
# (C) Copyright 2019 NetApp, Inc.
#
# This sample code is provided AS IS, with no support or warranties of
# any kind, including but not limited for warranties of merchantability
# or fitness of any kind, expressed or implied. Permission to use,
# reproduce, modify and create derivatives of the sample code is granted
# solely for the purpose of researching, designing, developing and
# testing a software application product for use with NetApp products,
# provided that the above copyright notice appears in all copies and
# that the software application product is distributed pursuant to terms
# no less restrictive than those set forth herein.
#
##-----

import json
import logging
import requests

requests.packages.urllib3.disable_warnings()

class DeployRequests(object):
    '''
    Wrapper class for requests that simplifies the ONTAP Select Deploy
    path creation and header manipulations for simpler code.

```

```

'''

def __init__(self, ip, admin_password):
    self.base_url = 'https://{}/api'.format(ip)
    self.auth = ('admin', admin_password)
    self.headers = {'Accept': 'application/json'}
    self.logger = logging.getLogger('deploy')

def post(self, path, data, files=None, wait_for_job=False):
    if files:
        self.logger.debug('POST FILES:')
        response = requests.post(self.base_url + path,
                                auth=self.auth, verify=False,
                                files=files)
    else:
        self.logger.debug('POST DATA: %s', data)
        response = requests.post(self.base_url + path,
                                auth=self.auth, verify=False,
                                json=data,
                                headers=self.headers)

    self.logger.debug('HEADERS: %s\nBODY: %s', self.filter_headers
(response), response.text)
    self.exit_on_errors(response)

    if wait_for_job and response.status_code == 202:
        self.wait_for_job(response.json())
    return response

def patch(self, path, data, wait_for_job=False):
    self.logger.debug('PATCH DATA: %s', data)
    response = requests.patch(self.base_url + path,
                              auth=self.auth, verify=False,
                              json=data,
                              headers=self.headers)

    self.logger.debug('HEADERS: %s\nBODY: %s', self.filter_headers
(response), response.text)
    self.exit_on_errors(response)

    if wait_for_job and response.status_code == 202:
        self.wait_for_job(response.json())
    return response

def put(self, path, data, files=None, wait_for_job=False):
    if files:
        print('PUT FILES: {}'.format(data))

```

```

        response = requests.put(self.base_url + path,
                                auth=self.auth, verify=False,
                                data=data,
                                files=files)

    else:
        self.logger.debug('PUT DATA:')
        response = requests.put(self.base_url + path,
                                auth=self.auth, verify=False,
                                json=data,
                                headers=self.headers)

        self.logger.debug('HEADERS: %s\nBODY: %s', self.filter_headers
(response), response.text)
        self.exit_on_errors(response)

        if wait_for_job and response.status_code == 202:
            self.wait_for_job(response.json())
        return response

    def get(self, path):
        """ Get a resource object from the specified path """
        response = requests.get(self.base_url + path, auth=self.auth,
verify=False)
        self.logger.debug('HEADERS: %s\nBODY: %s', self.filter_headers
(response), response.text)
        self.exit_on_errors(response)
        return response

    def delete(self, path, wait_for_job=False):
        """ Delete's a resource from the specified path """
        response = requests.delete(self.base_url + path, auth=self.auth,
verify=False)
        self.logger.debug('HEADERS: %s\nBODY: %s', self.filter_headers
(response), response.text)
        self.exit_on_errors(response)

        if wait_for_job and response.status_code == 202:
            self.wait_for_job(response.json())
        return response

    def find_resource(self, path, name, value):
        ''' Returns the 'id' of the resource if it exists, otherwise None
'''
        resource = None
        response = self.get('{path}?{field}={value}'.format(
                                path=path, field=name, value=value))

```

```

        if response.status_code == 200 and response.json().get(
('num_records') >= 1:
            resource = response.json().get('records')[0].get('id')
            return resource

    def get_num_records(self, path, query=None):
        ''' Returns the number of records found in a container, or None on
error '''
        resource = None
        query_opt = '?{}'.format(query) if query else ''
        response = self.get('{path}{query}'.format(path=path, query
=query_opt))
        if response.status_code == 200 :
            return response.json().get('num_records')
        return None

    def resource_exists(self, path, name, value):
        return self.find_resource(path, name, value) is not None

    def wait_for_job(self, response, poll_timeout=120):
        last_modified = response['job']['last_modified']
        job_id = response['job']['id']

        self.logger.info('Event: ' + response['job']['message'])

        while True:
            response = self.get('/jobs/{}?fields=state,message&'
                                'poll_timeout={}&last_modified=>={}'
                                .format(
                                    job_id, poll_timeout, last_modified))

            job_body = response.json().get('record', {})

            # Show interesting message updates
            message = job_body.get('message', '')
            self.logger.info('Event: ' + message)

            # Refresh the last modified time for the poll loop
            last_modified = job_body.get('last_modified')

            # Look for the final states
            state = job_body.get('state', 'unknown')
            if state in ['success', 'failure']:
                if state == 'failure':
                    self.logger.error('FAILED background job.\nJOB: %s',
job_body)

                    exit(1) # End the script if a failure occurs

```

```

        break

    def exit_on_errors(self, response):
        if response.status_code >= 400:
            self.logger.error('FAILED request to URL: %s\nHEADERS: %s\nRESPONSE BODY: %s',
                             response.request.url,
                             self.filter_headers(response),
                             response.text)
            response.raise_for_status() # Displays the response error, and
            exits the script

    @staticmethod
    def filter_headers(response):
        ''' Returns a filtered set of the response headers '''
        return {key: response.headers[key] for key in ['Location',
        'request-id'] if key in response.headers}

```

Script para cambiar el tamaño de los nodos del clúster

Puede usar el siguiente script para cambiar el tamaño de los nodos de un clúster de ONTAP Select.

```

#!/usr/bin/env python
##-----
#
# File: resize_nodes.py
#
# (C) Copyright 2019 NetApp, Inc.
#
# This sample code is provided AS IS, with no support or warranties of
# any kind, including but not limited for warranties of merchantability
# or fitness of any kind, expressed or implied. Permission to use,
# reproduce, modify and create derivatives of the sample code is granted
# solely for the purpose of researching, designing, developing and
# testing a software application product for use with NetApp products,
# provided that the above copyright notice appears in all copies and
# that the software application product is distributed pursuant to terms
# no less restrictive than those set forth herein.
#
##-----

import argparse
import logging
import sys

```



```

from deploy_requests import DeployRequests

def _parse_args():
    """ Parses the arguments provided on the command line when executing
    this
        script and returns the resulting namespace. If all required
    arguments
        are not provided, an error message indicating the mismatch is
    printed and
        the script will exit.
    """

    parser = argparse.ArgumentParser(description=(
        'Uses the ONTAP Select Deploy API to resize the nodes in the
    cluster.'
        ' For example, you might have a small (4 CPU, 16GB RAM per node) 2
    node'
        ' cluster and wish to resize the cluster to medium (8 CPU, 64GB
    RAM per'
        ' node). This script will take in the cluster details and then
    perform'
        ' the operation and wait for it to complete.'
    ))
    parser.add_argument('--deploy', required=True, help=(
        'Hostname or IP of the ONTAP Select Deploy VM.'
    ))
    parser.add_argument('--deploy-password', required=True, help=(
        'The password for the ONTAP Select Deploy admin user.'
    ))
    parser.add_argument('--cluster', required=True, help=(
        'Hostname or IP of the cluster management interface.'
    ))
    parser.add_argument('--instance-type', required=True, help=(
        'The desired instance size of the nodes after the operation is
    complete.'
    ))
    parser.add_argument('--ontap-password', required=True, help=(
        'The password for the ONTAP administrative user account.'
    ))
    parser.add_argument('--ontap-username', default='admin', help=(
        'The username for the ONTAP administrative user account. Default:
    admin.'
    ))
    parser.add_argument('--nodes', nargs='+', metavar='NODE_NAME', help=(

```

```

        'A space separated list of node names for which the resize
operation'
        ' should be performed. The default is to apply the resize to all
nodes in'
        ' the cluster. If a list of nodes is provided, it must be provided
in HA'
        ' pairs. That is, in a 4 node cluster, nodes 1 and 2 (partners)
must be'
        ' resized in the same operation.'
    ))
    return parser.parse_args()

def _get_cluster(deploy, parsed_args):
    """ Locate the cluster using the arguments provided """

    cluster_id = deploy.find_resource('/clusters', 'ip', parsed_args
.cluster)
    if not cluster_id:
        return None
    return deploy.get('/clusters/%s?fields=nodes' % cluster_id).json
()['record']

def _get_request_body(parsed_args, cluster):
    """ Build the request body """

    changes = {'admin_password': parsed_args.ontap_password}

    # if provided, use the list of nodes given, else use all the nodes in
the cluster
    nodes = [node for node in cluster['nodes']]
    if parsed_args.nodes:
        nodes = [node for node in nodes if node['name'] in parsed_args
.nodes]

    changes['nodes'] = [
        {'instance_type': parsed_args.instance_type, 'id': node['id']} for
node in nodes]

    return changes

def main():
    """ Set up the resize operation by gathering the necessary data and
then send
        the request to the ONTAP Select Deploy server.

```

```

"""

logging.basicConfig(
    format='[%(asctime)s] [%(levelname)5s] %(message)s', level=
logging.INFO,)

logging.getLogger('requests.packages.urllib3').setLevel(logging
.WARNING)

parsed_args = _parse_args()
deploy = DeployRequests(parsed_args.deploy, parsed_args
.deploy_password)

cluster = _get_cluster(deploy, parsed_args)
if not cluster:
    deploy.logger.error(
        'Unable to find a cluster with a management IP of %s' %
parsed_args.cluster)
    return 1

changes = _get_request_body(parsed_args, cluster)
deploy.patch('/clusters/%s' % cluster['id'], changes, wait_for_job
=True)

if __name__ == '__main__':
    sys.exit(main())

```

Información de copyright

Copyright © 2024 NetApp, Inc. Todos los derechos reservados. Imprimido en EE. UU. No se puede reproducir este documento protegido por copyright ni parte del mismo de ninguna forma ni por ningún medio (gráfico, electrónico o mecánico, incluidas fotocopias, grabaciones o almacenamiento en un sistema de recuperación electrónico) sin la autorización previa y por escrito del propietario del copyright.

El software derivado del material de NetApp con copyright está sujeto a la siguiente licencia y exención de responsabilidad:

ESTE SOFTWARE LO PROPORCIONA NETAPP «TAL CUAL» Y SIN NINGUNA GARANTÍA EXPRESA O IMPLÍCITA, INCLUYENDO, SIN LIMITAR, LAS GARANTÍAS IMPLÍCITAS DE COMERCIALIZACIÓN O IDONEIDAD PARA UN FIN CONCRETO, CUYA RESPONSABILIDAD QUEDA EXIMIDA POR EL PRESENTE DOCUMENTO. EN NINGÚN CASO NETAPP SERÁ RESPONSABLE DE NINGÚN DAÑO DIRECTO, INDIRECTO, ESPECIAL, EJEMPLAR O RESULTANTE (INCLUYENDO, ENTRE OTROS, LA OBTENCIÓN DE BIENES O SERVICIOS SUSTITUTIVOS, PÉRDIDA DE USO, DE DATOS O DE BENEFICIOS, O INTERRUPCIÓN DE LA ACTIVIDAD EMPRESARIAL) CUALQUIERA SEA EL MODO EN EL QUE SE PRODUJERON Y LA TEORÍA DE RESPONSABILIDAD QUE SE APLIQUE, YA SEA EN CONTRATO, RESPONSABILIDAD OBJETIVA O AGRAVIO (INCLUIDA LA NEGLIGENCIA U OTRO TIPO), QUE SURJAN DE ALGÚN MODO DEL USO DE ESTE SOFTWARE, INCLUSO SI HUBIEREN SIDO ADVERTIDOS DE LA POSIBILIDAD DE TALES DAÑOS.

NetApp se reserva el derecho de modificar cualquiera de los productos aquí descritos en cualquier momento y sin aviso previo. NetApp no asume ningún tipo de responsabilidad que surja del uso de los productos aquí descritos, excepto aquello expresamente acordado por escrito por parte de NetApp. El uso o adquisición de este producto no lleva implícita ninguna licencia con derechos de patente, de marcas comerciales o cualquier otro derecho de propiedad intelectual de NetApp.

Es posible que el producto que se describe en este manual esté protegido por una o más patentes de EE. UU., patentes extranjeras o solicitudes pendientes.

LEYENDA DE DERECHOS LIMITADOS: el uso, la copia o la divulgación por parte del gobierno están sujetos a las restricciones establecidas en el subpárrafo (b)(3) de los derechos de datos técnicos y productos no comerciales de DFARS 252.227-7013 (FEB de 2014) y FAR 52.227-19 (DIC de 2007).

Los datos aquí contenidos pertenecen a un producto comercial o servicio comercial (como se define en FAR 2.101) y son propiedad de NetApp, Inc. Todos los datos técnicos y el software informático de NetApp que se proporcionan en este Acuerdo tienen una naturaleza comercial y se han desarrollado exclusivamente con fondos privados. El Gobierno de EE. UU. tiene una licencia limitada, irrevocable, no exclusiva, no transferible, no sublicenciable y de alcance mundial para utilizar los Datos en relación con el contrato del Gobierno de los Estados Unidos bajo el cual se proporcionaron los Datos. Excepto que aquí se disponga lo contrario, los Datos no se pueden utilizar, desvelar, reproducir, modificar, interpretar o mostrar sin la previa aprobación por escrito de NetApp, Inc. Los derechos de licencia del Gobierno de los Estados Unidos de América y su Departamento de Defensa se limitan a los derechos identificados en la cláusula 252.227-7015(b) de la sección DFARS (FEB de 2014).

Información de la marca comercial

NETAPP, el logotipo de NETAPP y las marcas que constan en <http://www.netapp.com/TM> son marcas comerciales de NetApp, Inc. El resto de nombres de empresa y de producto pueden ser marcas comerciales de sus respectivos propietarios.