



Python-Codebeispiele

ONTAP Select

NetApp
February 03, 2026

Inhalt

- Python-Codebeispiele 1
 - Skript zum Erstellen eines Clusters 1
 - JSON für das Erstellen eines Clusters 8
 - Single-Node-Cluster in ESXi 8
 - Single-Node-Cluster in ESXi über vCenter 9
 - Single-Node-Cluster auf KVM 11
 - Skript zum Hinzufügen einer Node-Lizenz 12
 - Skript zum Löschen eines Clusters 16
- Common Support-Modul 18
- Skript zur Größenanpassung der Cluster-Nodes 22

Python-Codebeispiele

Skript zum Erstellen eines Clusters

Sie können mit dem folgenden Skript auf der Grundlage von Parametern, die in dem Skript definiert sind, und einer JSON-Eingabedatei einen Cluster erstellen.

```
#!/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 für das Erstellen eines Clusters

Wenn Sie ein ONTAP Select-Cluster mithilfe der Python-Codebeispiele erstellen oder löschen, müssen Sie eine JSON-Datei als Eingabe für das Skript bereitstellen. Sie können das entsprechende JSON-Muster basierend auf Ihren Implementierungsplänen kopieren und ändern.

Single-Node-Cluster in 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
      }
    ]
  }
}
]
}

```

Single-Node-Cluster in ESXi über 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
        }
      ]
    }
  ]
}

```

Single-Node-Cluster auf KVM

```

{
  "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
    }
  ]
}
]
}

```

Skript zum Hinzufügen einer Node-Lizenz

Mit dem folgenden Skript können Sie eine Lizenz für einen ONTAP Select-Knoten hinzufügen.

```

#!/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)

```

Skript zum Löschen eines Clusters

Sie können das folgende CLI-Skript verwenden, um einen vorhandenen Cluster zu löschen.

```

#!/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/{}?fields=state'.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)

```

Common Support-Modul

Alle Python-Skripte verwenden eine gemeinsame Python-Klasse in einem einzigen Modul.

```

#!/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/{id}?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}

```

Skript zur Größenanpassung der Cluster-Nodes

Mit dem folgenden Skript können Sie die Größe der Nodes in einem ONTAP Select Cluster ändern.

```

#!/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())

```

Copyright-Informationen

Copyright © 2026 NetApp. Alle Rechte vorbehalten. Gedruckt in den USA. Dieses urheberrechtlich geschützte Dokument darf ohne die vorherige schriftliche Genehmigung des Urheberrechtinhabers in keiner Form und durch keine Mittel – weder grafische noch elektronische oder mechanische, einschließlich Fotokopieren, Aufnehmen oder Speichern in einem elektronischen Abrufsystem – auch nicht in Teilen, vervielfältigt werden.

Software, die von urheberrechtlich geschütztem NetApp Material abgeleitet wird, unterliegt der folgenden Lizenz und dem folgenden Haftungsausschluss:

DIE VORLIEGENDE SOFTWARE WIRD IN DER VORLIEGENDEN FORM VON NETAPP ZUR VERFÜGUNG GESTELLT, D. H. OHNE JEGLICHE EXPLIZITE ODER IMPLIZITE GEWÄHRLEISTUNG, EINSCHLIESSLICH, JEDOCH NICHT BESCHRÄNKT AUF DIE STILLSCHWEIGENDE GEWÄHRLEISTUNG DER MARKTGÄNGIGKEIT UND EIGNUNG FÜR EINEN BESTIMMTEN ZWECK, DIE HIERMIT AUSGESCHLOSSEN WERDEN. NETAPP ÜBERNIMMT KEINERLEI HAFTUNG FÜR DIREKTE, INDIREKTE, ZUFÄLLIGE, BESONDERE, BEISPIELHAFT SCHÄDEN ODER FOLGESCHÄDEN (EINSCHLIESSLICH, JEDOCH NICHT BESCHRÄNKT AUF DIE BESCHAFFUNG VON ERSATZWAREN ODER -DIENSTLEISTUNGEN, NUTZUNGS-, DATEN- ODER GEWINNVERLUSTE ODER UNTERBRECHUNG DES GESCHÄFTSBETRIEBS), UNABHÄNGIG DAVON, WIE SIE VERURSACHT WURDEN UND AUF WELCHER HAFTUNGSTHEORIE SIE BERUHEN, OB AUS VERTRAGLICH FESTGELEGTER HAFTUNG, VERSCHULDENSUNABHÄNGIGER HAFTUNG ODER DELIKTSHAFTUNG (EINSCHLIESSLICH FAHRLÄSSIGKEIT ODER AUF ANDEREM WEGE), DIE IN IRGEND EINER WEISE AUS DER NUTZUNG DIESER SOFTWARE RESULTIEREN, SELBST WENN AUF DIE MÖGLICHKEIT DERARTIGER SCHÄDEN HINGEWIESEN WURDE.

NetApp behält sich das Recht vor, die hierin beschriebenen Produkte jederzeit und ohne Vorankündigung zu ändern. NetApp übernimmt keine Verantwortung oder Haftung, die sich aus der Verwendung der hier beschriebenen Produkte ergibt, es sei denn, NetApp hat dem ausdrücklich in schriftlicher Form zugestimmt. Die Verwendung oder der Erwerb dieses Produkts stellt keine Lizenzierung im Rahmen eines Patentrechts, Markenrechts oder eines anderen Rechts an geistigem Eigentum von NetApp dar.

Das in diesem Dokument beschriebene Produkt kann durch ein oder mehrere US-amerikanische Patente, ausländische Patente oder anhängige Patentanmeldungen geschützt sein.

ERLÄUTERUNG ZU „RESTRICTED RIGHTS“: Nutzung, Vervielfältigung oder Offenlegung durch die US-Regierung unterliegt den Einschränkungen gemäß Unterabschnitt (b)(3) der Klausel „Rights in Technical Data – Noncommercial Items“ in DFARS 252.227-7013 (Februar 2014) und FAR 52.227-19 (Dezember 2007).

Die hierin enthaltenen Daten beziehen sich auf ein kommerzielles Produkt und/oder einen kommerziellen Service (wie in FAR 2.101 definiert) und sind Eigentum von NetApp, Inc. Alle technischen Daten und die Computersoftware von NetApp, die unter diesem Vertrag bereitgestellt werden, sind gewerblicher Natur und wurden ausschließlich unter Verwendung privater Mittel entwickelt. Die US-Regierung besitzt eine nicht ausschließliche, nicht übertragbare, nicht unterlizenzierbare, weltweite, limitierte unwiderrufliche Lizenz zur Nutzung der Daten nur in Verbindung mit und zur Unterstützung des Vertrags der US-Regierung, unter dem die Daten bereitgestellt wurden. Sofern in den vorliegenden Bedingungen nicht anders angegeben, dürfen die Daten ohne vorherige schriftliche Genehmigung von NetApp, Inc. nicht verwendet, offengelegt, vervielfältigt, geändert, aufgeführt oder angezeigt werden. Die Lizenzrechte der US-Regierung für das US-Verteidigungsministerium sind auf die in DFARS-Klausel 252.227-7015(b) (Februar 2014) genannten Rechte beschränkt.

Markeninformationen

NETAPP, das NETAPP Logo und die unter <http://www.netapp.com/TM> aufgeführten Marken sind Marken von NetApp, Inc. Andere Firmen und Produktnamen können Marken der jeweiligen Eigentümer sein.