



Workflow processes

ONTAP Select

NetApp
January 31, 2025

Table of Contents

- Workflow processes 1
- Before you use the API workflows 1
- Workflow 1: Create a single-node evaluation cluster on ESXi 1

Workflow processes

Before you use the API workflows

You should prepare to review and use the workflow processes.

Understand the API calls used in the workflows

The ONTAP Select online documentation page includes the details of every REST API call. Rather than repeat those details here, each API call used in the workflow samples includes only the information you need to locate the call on the documentation page. After locating a specific API call, you can review the complete details of the call, including the input parameters, output formats, HTTP status codes, and request processing type.

The following information is included for each API call within a workflow to help locate the call on the documentation page:

- **Category**
The API calls are organized on the documentation page into functionally related areas or categories. To locate a specific API call, scroll to the bottom of the page and click the applicable API category.
- **HTTP verb**
The HTTP verb identifies the action performed on a resource. Each API call is executed through a single HTTP verb.
- **Path**
The path determines the specific resource which the action applies to as part of performing a call. The path string is appended to the core URL to form the complete URL identifying the resource.

Construct a URL to directly access the REST API

In addition to the ONTAP Select documentation page, you can also access the Deploy REST API directly through a programming language such as Python. In this case, the core URL is slightly different than the URL used when accessing the online documentation page. When accessing the API directly, you must append /api to the domain and port string. For example:

```
http://deploy.mycompany.com/api
```

Workflow 1: Create a single-node evaluation cluster on ESXi

You can deploy a single-node ONTAP Select cluster on a VMware ESXi host managed by vCenter. The cluster is created with an evaluation license.

The cluster creation workflow differs in the following situations:

- The ESXi host is not managed by vCenter (standalone host)
- Multiple nodes or hosts are used within the cluster
- Cluster is deployed in a production environment with a purchased license
- The KVM hypervisor is used instead of VMware ESXi

1. Register vCenter server credential

When deploying to an ESXi host managed by a vCenter server, you must add a credential before registering the host. The Deploy administration utility can then use the credential to authenticate to vCenter.

Category	HTTP verb	Path
Deploy	POST	/security/credentials

Curl

```
curl -iX POST -H 'Content-Type: application/json' -u admin:<password> -k -d @step01 'https://10.21.191.150/api/security/credentials'
```

JSON input (step01)

```
{
  "hostname": "vcenter.company-demo.com",
  "type": "vcenter",
  "username": "misteradmin@vsphere.local",
  "password": "mypassword"
}
```

Processing type

Asynchronous

Output

- Credential ID in the location response header
- Job object

2. Register a hypervisor host

You must add a hypervisor host where the virtual machine containing the ONTAP Select node will run.

Category	HTTP verb	Path
Cluster	POST	/hosts

Curl

```
curl -iX POST -H 'Content-Type: application/json' -u admin:<password> -k -d @step02 'https://10.21.191.150/api/hosts'
```

JSON input (step02)

```

{
  "hosts": [
    {
      "hypervisor_type": "ESX",
      "management_server": "vcenter.company-demo.com",
      "name": "esx1.company-demo.com"
    }
  ]
}

```

Processing type

Asynchronous

Output

- Host ID in the location response header
- Job object

3. Create a cluster

When you create an ONTAP Select cluster, the basic cluster configuration is registered and the node names are automatically generated by Deploy.

Category	HTTP verb	Path
Cluster	POST	/clusters

Curl

The query parameter `node_count` should be set to 1 for a single-node cluster.

```

curl -iX POST -H 'Content-Type: application/json' -u admin:<password> -k
-d @step03 'https://10.21.191.150/api/clusters? node_count=1'

```

JSON input (step03)

```

{
  "name": "my_cluster"
}

```

Processing type

Synchronous

Output

- Cluster ID in the location response header

4. Configure the cluster

There are several attributes you must provide as part of configuring the cluster.

Category	HTTP verb	Path
Cluster	PATCH	/clusters/{cluster_id}

Curl

You must provide the cluster ID.

```
curl -iX PATCH -H 'Content-Type: application/json' -u admin:<password> -k -d @step04 'https://10.21.191.150/api/clusters/CLUSTERID'
```

JSON input (step04)

```
{
  "dns_info": {
    "domains": ["lab1.company-demo.com"],
    "dns_ips": ["10.206.80.135", "10.206.80.136"]
  },
  "ontap_image_version": "9.5",
  "gateway": "10.206.80.1",
  "ip": "10.206.80.115",
  "netmask": "255.255.255.192",
  "ntp_servers": {"10.206.80.183"}
}
```

Processing type

Synchronous

Output

None

5. Retrieve the node name

The Deploy administration utility automatically generates the node identifiers and names when a cluster is created. Before you can configure a node, you must retrieve the assigned ID.

Category	HTTP verb	Path
Cluster	GET	/clusters/{cluster_id}/nodes

Curl

You must provide the cluster ID.

```
curl -iX GET -u admin:<password> -k
'https://10.21.191.150/api/clusters/CLUSTERID/nodes?fields=id,name'
```

Processing type

Synchronous

Output

- Array records each describing a single node with the unique ID and name

6. Configure the nodes

You must provide the basic configuration for the node, which is the first of three API calls used to configure a node.

Category	HTTP verb	Path
Cluster	PATH	/clusters/{cluster_id}/nodes/{node_id}

Curl

You must provide the cluster ID and node ID.

```
curl -iX PATCH -H 'Content-Type: application/json' -u admin:<password> -k
-d @step06 'https://10.21.191.150/api/clusters/CLUSTERID/nodes/NODEID'
```

JSON input (step06)

You must provide the host ID where the ONTAP Select node will run.

```
{
  "host": {
    "id": "HOSTID"
  },
  "instance_type": "small",
  "ip": "10.206.80.101",
  "passthrough_disks": false
}
```

Processing type

Synchronous

Output

None

7. Retrieve the node networks

You must identify the data and management networks used by the node in the single-node cluster. The internal network is not used with a single-node cluster.

Category	HTTP verb	Path
Cluster	GET	/clusters/{cluster_id}/nodes/{node_id}/networks

Curl

You must provide the cluster ID and node ID.

```
curl -iX GET -u admin:<password> -k 'https://10.21.191.150/api/
clusters/CLUSTERID/nodes/NODEID/networks?fields=id,purpose'
```

Processing type

Synchronous

Output

- Array of two records each describing a single network for the node, including the unique ID and purpose

8. Configure the node networking

You must configure the data and management networks. The internal network is not used with a single-node cluster.



Issue the following API call two times, once for each network.

Category	HTTP verb	Path
Cluster	PATCH	/clusters/{cluster_id}/nodes/{node_id}/networks/{network_id}

Curl

You must provide the cluster ID, node ID, and network ID.

```
curl -iX PATCH -H 'Content-Type: application/json' -u admin:<password> -k
-d @step08 'https://10.21.191.150/api/clusters/
CLUSTERID/nodes/NODEID/networks/NETWORKID'
```

JSON input (step08)

You need to provide the name of the network.

```
{
  "name": "sDOT_Network"
}
```

Processing type

Synchronous

Output

None

9. Configure the node storage pool

The final step in configuring a node is to attach a storage pool. You can determine the available storage pools through the vSphere web client, or optionally through the Deploy REST API.

Category	HTTP verb	Path
Cluster	PATCH	/clusters/{cluster_id}/nodes/{node_id}/networks/{network_id}

Curl

You must provide the cluster ID, node ID, and network ID.

```
curl -iX PATCH -H 'Content-Type: application/json' -u admin:<password> -k -d @step09 'https://10.21.191.150/api/clusters/ CLUSTERID/nodes/NODEID'
```

JSON input (step09)

The pool capacity is 2 TB.

```
{
  "pool_array": [
    {
      "name": "sDOT-01",
      "capacity": 2147483648000
    }
  ]
}
```

Processing type

Synchronous

Output

None

10. Deploy the cluster

After the cluster and node have been configured, you can deploy the cluster.

Category	HTTP verb	Path
Cluster	POST	/clusters/{cluster_id}/deploy

Curl

You must provide the cluster ID.

```
curl -iX POST -H 'Content-Type: application/json' -u admin:<password> -k  
-d @step10 'https://10.21.191.150/api/clusters/CLUSTERID/deploy'
```

JSON input (step10)

You must provide the password for the ONTAP administrator account.

```
{  
  "ontap_credentials": {  
    "password": "mypassword"  
  }  
}
```

Processing type

Asynchronous

Output

- Job object

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.