



Automate using REST APIs

ONTAP tools for VMware vSphere 10.1

NetApp
June 21, 2024

Table of Contents

- Automate using REST APIs 1
 - Overview of REST APIs 1
 - How to access ONTAP tools for VMware vSphere REST API 1
 - Input variables controlling an API request 2
 - Access ONTAP tools for VMware vSphere API reference documentation through the Swagger UI 5
 - Get started with the REST API 6
 - Workflows 6

Automate using REST APIs

Overview of REST APIs

REST APIs can be used to perform several ONTAP tools for VMware vSphere management operations. REST APIs are exposed through the Swagger web page.

You can access the Swagger web page available at <https://loadbalancerIP:8443/> to display the REST API documentation, as well as to manually issue an API call.



All the APIs have request body and examples mentioned in the swagger page. The workflows and examples provided in this section are for reference purposes only.

How to access ONTAP tools for VMware vSphere REST API

You can access the ONTAP REST API in several different ways.

Network considerations

You can connect to the REST API through the following interfaces:

- Cluster management LIF
- Node management LIF
- SVM management LIF

The LIF you choose to use should be configured to support the HTTPS management protocol. Also, the firewall configuration in your network should allow the HTTPS traffic.



You should always use a cluster management LIF. This will load balance the API requests across all the nodes and avoid nodes that are offline or experiencing connectivity issues. If you have multiple cluster management LIFs configured, they are all equivalent regarding access to the REST API.

ONTAP tools for VMware vSphere API online documentation page

You can access the Swagger from the hyperlink in the support page of the NetAPP ONTAP tools for VMware vSphere plug-in.

The format of the URL used to access the documentation page for the most recent version of the API is:

```
`https://<loadbalancer_ip_address>/docs/api`
```

Custom software and tools

You can access ONTAP tools for VMware vSphere API using any of several different programming languages and tools. Popular choices include Python, Java, Curl, and PowerShell. A program, script, or tool that uses the API acts as a REST web services client. Using a programming language enables a deeper understanding of the API and provides an opportunity to automate ONTAP tools for VMware vSphere administration.

The format of the base URL used to directly access the most recent version of the API is:

```
`https://<loadbalancer_ip_address>/api`
```

To access a specific API version where multiple versions are supported, the format of the URL is:

```
`https://<loadbalancer_ip_address>/api/v1`
```

Input variables controlling an API request

You can control how an API call is processed through parameters and variables set in the HTTP request.

HTTP methods

The HTTP methods supported by ONTAP tools for VMware vSphere REST API are shown in the following table.



Not all the HTTP methods are available at each of the REST endpoints.

HTTP method	Description
GET	Retrieves object properties on a resource instance or collection.
POST	Creates a new resource instance based on the supplied input.
DELETE	Deletes an existing resource instance.
PUT	Modifies an existing resource instance.

Request headers

You should include several headers in the HTTP request.

Content-type

If the request body includes JSON, this header should be set to *application/json*.

Accept

This header should be set to *application/json*.

Authorization

Basic authentication should be set with the username and password encoded as a base64 string.

Request body

The content of the request body varies depending on the specific call. The HTTP request body consists of one of the following:

- JSON object with input variables
- Empty

Filtering objects

When issuing an API call that uses GET, you can limit or filter the returned objects based on any attribute. For example, you can specify an exact value to match:

```
<field>=<query value>
```

In addition to an exact match, other operators are available to return a set of objects over a range of values. ONTAP tools for VMware vSphere REST API supports the filtering operators shown in the table below.

Operator	Description
=	Equal to
<	Less than
>	Greater than
≤	Less than or equal to
≥	Greater than or equal to
UPDATE	Or
!	Not equal to
*	Greedy wildcard

You can also return a collection of objects based on whether a specific field is set or not set by using the **null** keyword or its negation **!null** as part of the query.



Any fields that are not set are generally excluded from matching queries.

Requesting specific object fields

By default, issuing an API call using GET returns only the attributes that uniquely identify the object or objects. This minimum set of fields acts as a key for each object and varies based on the object type. You can select additional object properties using the `fields` query parameter in the following ways:

Common or standard fields

Specify **fields=*** to retrieve the most commonly used object fields. These fields are typically maintained in local server memory or require little processing to access. These are the same properties returned for an object after using GET with a URL path key (UUID).

All fields

Specify **fields=**** to retrieve all the object fields, including those requiring additional server processing to access.

Custom field selection

Use **fields=<field_name>** to specify the exact field you want. When requesting multiple fields, the values should be separated using commas without spaces.



As a best practice, you should always identify the specific fields you want. You should only retrieve the set of common fields or all fields when needed. Which fields are classified as common, and returned using *fields=**, is determined by NetApp based on internal performance analysis. The classification of a field might change in future releases.

Sorting objects in the output set

The records in a resource collection are returned in the default order defined by the object. You can change the order using the `order_by` query parameter with the field name and sort direction as follows:

```
order_by=<field name> asc|desc
```

For example, you can sort the `type` field in descending order followed by `id` in ascending order:

```
order_by=type desc, id asc
```

- If you specify a sort field but do not provide a direction, the values are sorted in ascending order.
- When including multiple parameters, you should separate the fields with a comma.

Pagination when retrieving objects in a collection

When issuing an API call using GET to access a collection of objects of the same type, ONTAP tools for VMware vSphere attempts to return as many objects as possible based on two constraints. You can control each of these constraints using additional query parameters on the request. The first constraint reached for a specific GET request terminates the request and therefore limits the number of records returned.



If a request ends before iterating over all the objects, the response contains the link needed to retrieve the next batch of records.

Limiting the number of objects

By default, ONTAP tools for VMware vSphere returns a maximum of 10,000 objects for a GET request. You can change this limit using the `max_records` query parameter. For example:

```
max_records=20
```

The number of objects returned can be less than the maximum in effect, based on the related time constraint as well as the total number of objects in the system.

Limiting the time used to retrieve the objects

By default, ONTAP tools for VMware vSphere returns as many objects as possible within the time allowed for the GET request. The default timeout is 15 seconds. You can change this limit using the `return_timeout` query parameter. For example:

```
return_timeout=5
```

The number of objects returned can be less than the maximum in effect, based on the related constraint on the number of objects as well as the total number of objects in the system.

Narrowing the result set

If needed, you can combine these two parameters with additional query parameters to narrow the result set. For example, the following returns up to 10 EMS events generated after the specified time:

```
time⇒ 2018-04-04T15:41:29.140265Z&max_records=10
```

You can issue multiple requests to page through the objects. Each subsequent API call should use a new time value based on the latest event in the last result set.

Size properties

The input values used with some API calls as well as certain query parameters are numeric. Rather than provide an integer in bytes, you can optionally use a suffix as shown in the following table.

Suffix	Description
KB	KB Kilobytes (1024 bytes) or kibibytes
MB	MB Megabytes (KB x 1024 bytes) or mebibytes
GB	GB Gigabytes (MB x 1024 bytes) or gibibytes
TB	TB Terabytes (GB x 1024 bytes) or tebibytes
PB	PB Petabytes (TB x 1024 bytes) or pebibytes

Access ONTAP tools for VMware vSphere API reference documentation through the Swagger UI

You can access the ONTAP REST API documentation through the Swagger UI at your local ONTAP system.

Before you begin

You should have the following:

- IP address or host name of the ONTAP cluster management LIF
- Username and password for an account with authority to access the ONTAP REST API

Steps

1. Type the URL in your browser and press **Enter**: `https://<ip_address>/docs/api`
2. Sign in using the ONTAP account

The ONTAP API documentation page is displayed with the API calls organized in major resource categories at the bottom.

3. As an example of an individual API call, scroll down to the **cluster** category and click **GET /cluster**.

Get started with the REST API

You can quickly get started using ONTAP tools for VMware vSphere REST API. Accessing the API provides some perspective before you begin using it with the more complex workflow processes on a live setup.

Hello World

You can run a simple command on your system to get started using ONTAP tools for VMware vSphere REST API and confirm its availability.

Before you begin

- Ensure that the Curl utility is available on your system.
- IP address or host name of ONTAP tools for VMware vSphere server
- Username and password for an account with authority to access ONTAP tools for VMware vSphere REST API.



If your credentials include special characters, you need to format them in a way that is acceptable to Curl based on the shell you are using. For example, you can insert a backslash before each special character or wrap the entire `username:password` string in single quotes.

Step

At the command line interface, run the following to retrieve the plug-in information:

```
curl -X GET -u username:password -k  
"https://<ip_address>/api/hosts?fields=IncludePluginInfo"
```

Example:

```
curl -X GET -u admin:password -k  
"'https://10.225.87.97/api/hosts?fields=IncludePluginInfo'"
```

Workflows

Storage discovery

Discovery interval can be configured as part of the configuration map. Scheduled discovery runs for every 60 mins. The API given here is to run the discovery on demand for a given storage backend which is added in the local scope.

Use the following API to run discovery:

```
POST  
/virtualization/api/v1/vcenters/{vcguid}/storage-backends/{id}/discovery-  
jobs
```




See Onboard storage backend (SVM or Cluster) workflow and get ID from post storage backend API response.

Discovery from this API endpoint is supported only for local scoped storage backends and not for the global scoped storage backends. If the storage backend type is cluster, discovery implicitly runs for the child SVMs. If the storage backend type is SVM, discovery only runs for the selected SVM.

Example:

To run discovery on a storage backend specified by ID

```
POST
/api/v1/vcenters/3fa85f64-5717-4562-b3fc-2c963f66afa6/storage-
backends/74e85f64-5717-4562-b3fc-2c963f669dde/discovery-jobs
```

You need to pass x-auth for the API. You can generate this x-auth from the new API added under Auth in Swagger.

```
/virtualization/api/v1/auth/vcenter-login
```

SVM aggregate mapping requirements

To use SVM user credentials for provisioning datastores, ONTAP tools for VMware vSphere creates volumes on the aggregate specified in the datastores POST API. ONTAP does not allow the creation of volumes on unmapped aggregates on an SVM using SVM user credentials. To resolve this, map the SVMs with the aggregates using the ONTAP REST API or CLI as described here.

ONTAP REST API:

```
PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
 '{"aggregates":{"name":["aggr1","aggr2","aggr3"]}}'
```

ONTAP CLI:

```
still15_vsim_ucs630f_aggr1 vserver show-aggregates
AvailableVserver Aggregate State Size Type SnapLock
Type
svm_test still15_vsim_ucs630f_aggr1
online 10.11GB vmdisk non-snaplock
```

Onboard storage backend (SVM or Cluster) with a vCenter Server instance

Use the following API to onboard the storage backends and map the SVM to vCenter

locally. See [Configure ONTAP user roles and privileges](#) section for the ONTAP SVM user privileges.

```
POST /virtualization/api/v1/vcenters/<vcguid>/storage-backends

{
  "hostname_or_ip": "172.21.103.107",
  "username": "svm11",
  "password": "xxxxxx"
}
```



The ID from the above API response is used in discovery.

You need to pass x-auth for the API. You can generate this x-auth from the new API added under Auth in Swagger.

```
/virtualization/api/v1/auth/vcenter-login
```

Create vVols datastore

You can create vVols datastore with new volumes or with existing volumes. You can also create vVols datastore with the combination of existing volumes and new volumes.



Check to ensure root aggregates are not mapped to SVM.

Generate a JWT token before creating datastores or increase the SAML token expiry by setting 'Maximum Bearer Token Lifetime' to 60m on vCenter.

You need to pass x-auth for the API. You can generate this x-auth from the new API added under Auth in Swagger.

```
/virtualization/api/v1/auth/vcenter-login
```

1. Create vVols datastore with new volume.

Get Aggregate id, storage_id(SVM uuid) using ONTAP REST API. POST
`/virtualization/api/v1/vcenters/cdded9ad-6bsd-4c9e-b44g-691250bfe2df/vvols/datastores`

Use the following URI to check the status:

+

```
`\https://xx.xx.xx.xxx:8443/virtualization/api/jobmanager/v2/jobs/<JobID>?includeSubJobsAndTasks=true`
```

+ Request Body for NFS datastore

```
{ "name": "nfsds1", "protocol": "nfs", "platform_type": "aff", "moref": "domain-c8", "volumes": [ { "is_existing": false, "name": "vol_nfs_pvt", "size_in_mb": 2048000, "space_efficiency": "thin", "aggregate": { "id": "d7078b3c-3827-4ac9-9273-0a32909455c2" }, "qos": { "min_iops": 200, "max_iops": 5000 } } ], "storage_backend": { "storage_id": "654c67bc-0f75-11ee-8a8c-00a09860a3ff" } }
```

Request body for iSCSI datastore: { "name": "iscsi_custom", "protocol": "iscsi", "platform_type": "aff", "moref": "domain-c8", "volumes": [{ "is_existing": false, "name": "iscsi_custom", "size_in_mb": 8034, "space_efficiency": "thin", "aggregate": { "id": "54fe5dd4-e461-49c8-bb2d-6d62c5d75af2" } }], "custom_igroup_name": "igroup1", "storage_backend": { "storage_id": "eb9d33ab-1960-11ee-9506-00a0985c6d9b" } } . Create vVols datastore with existing volumes.

Get aggregate_id and volume_id using ONTAP REST API.

```
POST /virtualization/api/v1/vcenters/cdded9ad-6bsd-4c9e-b44g-691250bfe2df/vvols/datastores
```

Request Body

```
{
  "name" : "nfsds2",
  "protocol" : "nfs",
  "platform_type": "aff",
  "moref" : "domain-c8",
  "volumes" : [
    {
      "is_existing": true,
      "id": "e632a632-1412-11ee-8a8c-00a09860a3ff"
    }
  ],
  "storage_backend": {
    "storage_id": "33a8b6b3-10cd-11ee-8a8c-00a09860a3ff"
  }
}
```

```
}
```

Mount and unmount vVols datastore

You can mount a VMware Virtual Volumes (vVols) datastore to one or more additional hosts to provide storage access to additional hosts. You can unmount vVols datastore using APIs.

Use the following API to mount or unmount a vVols datastore. You need to pass x-auth for the API. You can generate this x-auth from the new API added under Auth in Swagger.

```
/virtualization/api/v1/auth/vcenter-login
```

PATCH

```
/virtualization/api/v1/vcenters/{vcguid}/vvols/datastores/{moref}/hosts
```

Get vVol datastore moref from vCenter.

Request Body

```
{
  "operation": "mount",
  "morefs": [
    "host-7044"
  ],
}
```

Examples: * Mount on additional host

Use the following API to mount on additional host:

```
/api/v1/vcenters/cdded9ad-6bsd-4c9e-b44g-
691250bfe2df/vvols/datastores/datastore-24/hosts
```

Request Body

```
{
  "operation": "mount",
  "morefs": ["host-13"],
}
```

- Unmount on additional host

Use the following API to unmount on additional host:

```
/api/v1/vcenters/cdded9ad-6bsd-4c9e-b44g-
691250bfe2df/vvols/datastores/datastore-24/hosts
```

Request Body

```
{
  "operation": "unmount",
  "morefs": ["host-13"],
}
```

Expand or shrink Storage of vVol Datastore

There are APIs to increase or decrease the available storage.

Steps

Use the following API to expand or shrink the vVols datastore:

```
PATCH
/virtualization/api/v1/vcenters/{vcguid}/vvols/datastores/{moref}/volumes
```

Examples

- Modify vVols datastore for add new volume

```
PATCH virtualization/api/v1/vcenters/cdded9ad-6bsd-4c9e-b44g-
691250bfe2df/vvols/datastores/datastore-24/volumes
```

Request Body

```
{
  "operation": "grow",
  "volumes": [{
    "is_existing": false,
    "name": "exp3",
    "size_in_mb": 51200,
    "space_efficiency": "thin",
    "aggregate": {
      "id": "1466e4bf-c6d6-411a-91d5-c4f56210e1ab"
    },
    "storage_backend": {
      "storage_id": "13d86e4f-1fb1-11ee-9509-005056a75778"
    },
    "qos": {
      "max_iops": 5000
    }
  }]
}
```

- Modify vVols datastore for add existing volume

```
PATCH virtualization/api/v1/vcenters/cdded9ad-6bsd-4c9e-b44g-691250bfe2df/vvols/datastores/datastore-24/volumes
```

Request Body

```
{
  "operation": "grow",
  "volumes": [{
    "is_existing": true,
    "id": "vfded9ad-6bsd-4c9e-b44g-691250bfe2sd"
  }]
}
```

- Modify vVols datastore for remove volume and delete volume from storage

```
PATCH virtualization/api/v1/vcenters/cdded9ad-6bsd-4c9e-b44g-691250bfe2df/vvols/datastores/datastore-24/volumes?delete_volumes=true
```

Request Body

```
{
  "operation": "shrink",
  "volumes": [{
    "is_existing": true,
    "id": "vfded9ad-6bsd-4c9e-b44g-691250bfe2sd"
  }]
}
```

- Modify vVols datastore for remove volume and do not delete volume from storage

```
PATCH virtualization/api/v1/vcenters/cdded9ad-6bsd-4c9e-b44g-691250bfe2df/vvols/datastores/datastore-24/volumes?delete_volumes=false
```

Request Body

```
{
  "operation": "shrink",
  "volumes": [{
    "is_existing": true,
    "id": "vfded9ad-6bsd-4c9e-b44g-691250bfe2sd"
  }]
}
```

Delete vVols datastore

A vVols datastore exists as long as at least one FlexVol volume is available on the

datastore. If you want to delete a vVols datastore in a HA cluster, you should first unmount the datastore from all hosts within the HA cluster, and then delete the residing `.vsphere-HA` folder manually using the vCenter server user interface.

Steps

Use the following API to delete vVols datastore.

```
DELETE
/virtualization/api/v1/vcenters/{vcguid}/vvols/datastores/{moref}
```

Examples

- Delete vVols datastore and delete volumes from storage

```
DELETE /api/v1/vcenters/cdded9ad-6bsd-4c9e-b44g-
691250bfe2df/vvols/datastores/datastore-28?delete_volumes=true
```



Delete vVols Datastore workflow deletes datastore-volumes if you have passed the `delete_volume` flag as true irrespective of if the datastore-volume is managed or not managed.

- Delete vVols datastore and do not delete volumes from storage

```
DELETE /api/v1/vcenters/cdded9ad-6bsd-4c9e-b44g-
691250bfe2df/vvols/datastores/datastore-28?delete_volumes=false
```

Response:

```
{
  "id": "1889"
}
```

Manage Storage threshold

Use the following Get threshold API to retrieve the configured storage threshold limits for volume and aggregate.

```
GET/virtualization/api/v1/vcenters/{vcguid}/storage-thresholds
```

Examples: Get the Storage thresholds per vCenter Server instance by vCenter guid

```
GET "/api/v1/vcenters/beded9ad-6bbb-4c9e-b4c6-691250bfe2da/storage-
thresholds"
```

Use the following PATCH configure alarm for volume and aggregate to generate notification when configured threshold limits are reached.

```
PATCH/virtualization/api/v1/vcenters/{vcguid}/storage-thresholds
```

Examples: Update the Storage thresholds per vCenter by vCenter guid. Default limits are 80% for nearly-full and 90% for full. Modifying all threshold settings

```
{{}}PATCH "/api/v1/vcenters/beded9ad-6bbb-4c9e-b4c6-691250bfe2da/storage-
thresholds"
Request Body
{
"volume":

{ "nearly_full_percent": 80, "full_percent": 90 }
,
"aggregate": {
"nearly_full_percent": 80,
"full_percent": 90
}
}
}}}}
```

Manage network access

Use the following API to add IP addresses for whitelisting:

```
patch /api/v1/vcenters/{vcguid}/settings/ip-whitelist

{
  value: string
}

GET /api/v1/vcenters/{vcguid}/settings/ip-whitelist

{
  value: string
}
```


Copyright information

Copyright © 2024 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.