



# **Additional considerations when using the REST API**

## **ONTAP Automation**

NetApp  
September 24, 2021

# Table of Contents

- Additional considerations when using the REST API ..... 1
  - Object references and access ..... 1
  - Users of the ONTAPI API and ONTAP CLI ..... 2
  - Accessing the ONTAP CLI through the REST API ..... 3
  - Security using RBAC ..... 3
  - Performance metrics for storage resources ..... 4

# Additional considerations when using the REST API

There are several characteristics of the ONTAP REST API affecting its operation and use. You should be aware of these considerations before issuing an API call.

You should review the resources exposed through the API.

## Related links

[Summary of the REST resource categories](#)

## Object references and access

The resource instances or objects exposed through the ONTAP REST API can be referenced and accessed in several different ways.

### Object access paths

At a high level, there are two path types when accessing an object:

- Primary

The object is the primary or direct target of the API call.

- Foreign

The object is not the primary reference of the API call, but rather is linked to from the primary object. It is therefore a foreign or downstream object and referenced through a field in the primary object.

## Accessing an object using the UUID

Every object is assigned a unique identifier when it is created, which in most cases is a 128-bit UUID. The assigned UUID values are immutable and are used internally within ONTAP to access and manage the resources. Because of this, the UUID generally provides the fastest and most stable way to access objects.

For many of the resource types, a UUID value can be provided as part of the path key in the URL to access a specific object. For example, you can use the following to access a node instance:

```
`/cluster/nodes/{uuid}
```

### Accessing an object using an object property

In addition to a UUID, you can also access an object using an object property. In most cases, it is convenient to use the name property. For example, you can use the following query parameter in the URL string to access a node instance by its name: `/cluster/nodes?name=node_one`. In addition to a query parameter, a foreign object can be accessed through a property in the primary object.

While you can use the name or other property to access an object instead of the UUID, there are several possible disadvantages:

- The name field is not immutable and can be changed. If the name of an object is changed before accessing an object, the wrong object will be returned or an object access error will fail.



This issue can occur with a POST or PATCH method on a foreign object or with a GET method on a primary object.

- ONTAP must translate the name field into the corresponding UUID. This is a type of indirect access which can become a performance issue.

In particular, a performance degradation is possible when one or more of the following is true:

- GET method is used
- A large collection of objects is accessed
- A complex or elaborate query is used

### Cluster versus SVM context

There are several REST endpoints that support both a cluster and SVM. When using one of these endpoints, you can indicate the context of the API call through the `scope=[svm|cluster]` value. Examples of endpoints supporting a dual context include IP interfaces and security roles.



The scope value has a default value base on the properties provided for each API call.

### Using PATCH and DELETE on a collection of objects

Every REST endpoint supporting PATCH or DELETE on a resource instance also supports the same method on a collection of objects. The only requirement is that at least one field must be provided through a query parameter in the URL string. When issuing a PATCH or DELETE over a collection, this is equivalent to doing the following internally:

- Query-based GET to retrieve the collection
- Serial sequence of PATCH or DELETE calls on each object in the collection

The time out for the operation can be set by `return_timeout` with a default of 15 seconds. If not completed before the timeout, the response includes a link to the next object. You must reissue the same HTTP method using the next link to continue the operation.

## Users of the ONTAPI API and ONTAP CLI

ONTAPI is a set of API calls provided through the Network Manageability SDK that can be used to administer ONTAP storage systems. There are differences between the REST API and the ONTAPI API calls, as well as between the REST API and ONTAP CLI. You should understand these differences before using the REST API in a production environment.

### General design differences

The ONTAP REST API and command line interface have a fundamentally different designs. The CLI commands and parameters do not map directly to the REST API calls. And even where there might be a similarity, the details of the input parameters can be different. For example, numeric units might be specified in bytes or using a suffix (such as KB). You should review [Input variables controlling an API request](#) as well as the [API reference](#) for more information.

## Data SVMs exposed through the REST API

ONTAP supports several types of storage virtual machines (SVMs). However, only the data SVMs are directly exposed through the ONTAP REST API. The configuration information describing the cluster and nodes is available through the REST API, however the cluster and nodes are not treated as separate SVMs.

## Migrating from ONTAPI to REST

The REST API is the primary and strategic choice for automating ONTAP system administration. If you are currently using the ONTAPI API, you should consider migrating to the REST API. NetApp provides mapping documentation to assist with the migration from ONTAPI to REST.

### Related links

[Input variables controlling an API request](#)

[ONTAPI to REST API mapping information](#)

## Accessing the ONTAP CLI through the REST API

To assist CLI and ONTAPI API users in their transition to the ONTAP REST API, ONTAP provides a REST endpoint to access the CLI. You can use this passthrough feature to execute any CLI command. Usage of the REST endpoint is returned in AutoSupport data so NetApp can identify gaps in the REST API and make improvements in future releases.

To issue a CLI command, you must make a REST API call that is properly formed based on rules regarding:

- Resource paths
- Field names
- HTTP methods

The base resource path for CLI access is `/private/cli`. Refer to the ONTAP API online documentation page for details about accessing the CLI through the REST API.



NetApp maintains a GitHub repository containing code samples and other helpful information. You can access the repository for examples of how to use the CLI passthrough facility.

### Related links

[ONTAP REST Python GitHub repository - CLI passthrough samples](#)

## Security using RBAC

The REST API expands the role-based access control (RBAC) capabilities when using ONTAP. You can create user accounts with custom roles to restrict access to the REST endpoints.

### Creating roles for the REST endpoints

A REST role is defined through a set of one or more privileges. Each privilege consists of a path to a REST endpoint and the associated access level. Access to each endpoint is provided in one of three levels and determines the HTTP methods that can be used against the resource. The access levels include:

- All

All HTTP methods can be used

- Read-only

Only GET can be used

- None

No access is allowed

Here is an example of two privileges that can be assigned to a REST role:

- `access="readonly", path="/api/storage/volumes"`
- `access="none", path="/api/snapmirror/policies"`

### Creating a user account with a custom role

At a high level, you can create an account with a custom REST role as follows:

1. Create a user account with access to the HTTP management protocol.
2. Create a REST role with the desired privileges.
3. Associate the user account with the role.

## Performance metrics for storage resources

ONTAP collects performance metrics about selected SVM storage objects and protocols, and reports this information through the REST API. You can use this data to monitor the performance of an ONTAP system.

For a given storage object or protocol, the performance data falls into three categories:

- IOPS
- Latency
- Throughput

Within each category, one or more of the following types of data is available:

- Read (R)
- Write (W)
- Other (O)
- Total (T)

The following table summarizes the performance data available through the ONTAP REST API, including the release when it was added. Refer to the REST API online documentation page at your ONTAP system for more information.

<b>Storage object or protocol</b>	<b>IOPS</b>	<b>Latency</b>	<b>Throughput</b>	<b>ONTAP release</b>
Ethernet port	Not applicable	Not applicable	RWT	9.8
FC port	RWOT	RWOT	RWT	9.8
IP interface	Not applicable	Not applicable	RWT	9.8
FC interface	RWOT	RWOT	RWT	9.8
NVMe namespace	RWOT	RWOT	RWOT	9.8
Qtree statistics	Raw RWOT	Not applicable	Raw RWOT	9.8
Volume Flexcache	RWOT	RWOT	RWT	9.8
Node – process utilization	Process utilization as a numerical value	Process utilization as a numerical value	Process utilization as a numerical value	9.8
Cloud volume	RWOT	RWOT	Not applicable	9.7
LUN	RWOT	RWOT	RWOT	9.7
Aggregate	RWOT	RWOT	RWOT	9.7
SVM NFS protocol	RWOT	RWOT	RWT	9.7
SVM CIFS protocol	RWOT	RWOT	RWT	9.7
SVM FCP protocol	RWOT	RWOT	RWT	9.7
SVM iSCSI protocol	RWOT	RWOT	RWT	9.7
SVM NVMe protocol	RWOT	RWOT	RWT	9.7
Cluster	RWOT	RWOT	RWOT	9.6
Volumes	RWOT	RWOT	RWOT	9.6

## Copyright Information

Copyright © 2021 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## 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.