



Cluster

ONTAP 9.6 REST API reference

NetApp
May 23, 2024

This PDF was generated from https://docs.netapp.com/us-en/ontap-restapi-96/ontap/cluster_overview.html on May 23, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Cluster 1
 - Cluster overview..... 1
 - Manage clusters..... 2
 - Retrieve cluster chassis 88
 - View and manage cluster jobs 101
 - Manage cluster licensing 113
 - Manage a cluster license package..... 133
 - Retrieve historical performance metrics for the cluster 142
 - Manage cluster nodes 150
 - Manage cluster peers..... 205
 - Manage cluster schedules 248
 - Manage cluster software 276

Cluster

Cluster overview

Overview

These APIs enable you to perform a number of independent workflows, including:

- Creating the cluster
- Adding nodes to the cluster
- Managing cluster configuration data (including name, version, NTP servers, name servers, and DNS domains)
- Managing node configuration data (including node names, models, serial numbers, and HA group information)
- Discovering the nodes on the cluster network that can be added to the cluster
- Viewing and updating current and recent jobs
- Updating the cluster software

Pre-Cluster APIs

A few of the cluster APIs (namely, POST/OPTIONS on /api/cluster, GET/HEAD/OPTIONS on /api/cluster/nodes, and calls on /api/cluster/jobs) are allowed before the cluster is created. These APIs support creation of the cluster and monitoring of its progress. Any other cluster API used before the cluster is created will fail.

APIs

cluster

The cluster APIs cover basic management of the cluster, including viewing and modifying settings such as the name, UUID, version, NTP servers, DNS domains, and the nodes in the cluster. POST /api/cluster allows creation of the cluster, including adding all of the nodes to the cluster available during setup, initial configuration of cluster and node management interfaces, NTP servers, name servers, licenses, and node names.

nodes

The node APIs allow you to gather information about the nodes in a cluster, including model number, serial number, HA Group information, names, management interfaces, and UUIDs. By setting a query parameter, the administrator can also discover nodes on the cluster network that have not been added to the cluster and use that information in the POST operation to add them. Nodes can be added to the cluster through the POST API.

jobs

The job APIs are used to monitor the progress of running and recent jobs using GET. Some jobs are capable of being paused and cancelled using the PATCH operation.

Manage clusters

Cluster endpoint overview

Overview

This API is used to create a cluster, update cluster-wide configurations, and retrieve the current configuration details.

Creating a cluster

You can create a new cluster by issuing a POST request to /cluster. Parameters are provided in the body of the POST request to configure cluster-wide settings and add nodes during the cluster setup.

Fields used for creating a cluster

The fields used for the cluster APIs fall into the following categories:

Required cluster-wide configuration

The following fields are always required for any POST /cluster request:

- name
- password

Optional cluster-wide configuration

The following fields are used to setup additional cluster-wide configuration:

- location
- contact
- dns_domains
- name_servers
- ntp_servers
- license
- configuration_backup
- management_interface
- nodes

Nodes field

The nodes field specifies the nodes to join to the cluster. All nodes must be at the same version to use this API. If no nodes are specified, the cluster is configured with one node added. The node added is the node to which

the REST request is issued. If one node is specified, the "node.cluster_interface.ip.address" field must not be used. If multiple nodes are specified, the node to which the REST request is issued must be provided in addition to the remote nodes, and the "node.cluster_interface.ip.address" field is required for each node to identify them. All other node fields are optional in all cases. If a field is provided for one node, it must be provided for all nodes.

Node networking fields

The cluster management interface and each node management interface use the cluster management interface netmask and gateway. For advanced configurations where the cluster and node management interfaces are on different subnets, the /network/ip/interface APIs must be used to configure network interfaces after setup is complete. The management interfaces are used to communicate with the name servers and NTP servers. The address family of the name servers and NTP servers must match the management interfaces address family.

Single node cluster field

When the "single_node_cluster" field is set to true, the cluster is created in single node cluster mode. A node field for this node can be provided for node-specific configuration and the "node.cluster_interface.ip.address" field must not be used. Storage failover is configured to non-HA mode, and ports used for cluster ports are moved to the default IPspace. This might cause the node to reboot during setup. While a node reboots, the RESTful interface might not be available. See 'Connection failures during cluster create' for more information.

Performance monitoring

Performance of the cluster can be monitored by the `metric.*` and `statistics.*` fields. These show the performance of the cluster in terms of IOPS, latency and throughput. The `metric.*` fields denote an average whereas `statistics.*` fields denote a real-time monotonically increasing value aggregated across all nodes.

Monitoring cluster create status

Errors before the job starts

Configuration in the POST /cluster request is validated before the cluster create job starts. If an invalid configuration is found, an HTTP error code in the 4xx range is returned. No cluster create job is started.

Polling on the job

After a successful POST /cluster has been issued, an HTTP error code of 202 is returned along with a job UUID and link in the body of the response. The cluster create job continues asynchronously and can be monitored with the job UUID using the /cluster/jobs API. The "message" field in the response of GET /cluster/jobs/{uuid} shows the current step in the job and the "state" field shows the overall state of the job.

Errors during the job

If a failure occurs during the cluster create job, the job body provides details of the error along with error code fields. See the error table in the 'Responses' of the POST /cluster documentation for common error codes and descriptions.

Re-running POST /cluster

The POST /cluster command can be re-run if errors occur. When re-running the request, the same body and query parameters must be used. The value of any field in the original body or query can be changed, but the fields that were provided cannot be changed. For example, an initial request might have a body section as

follows:

```
body =
{
  "name": "clusCreateRerun",
  "password": "openSesame",
  "nodes": [
    {
      "cluster_interface": {
        "ip": {
          "address": "1.1.1.1"
        }
      }
    },
    {
      "cluster_interface": {
        "ip": {
          "address": "2.2.2.2"
        }
      }
    }
  ]
}
```

A re-run request updates the body details to:

```
body =
{
  "name": "clusCreateRerun",
  "password": "openSesame",
  "nodes": [
    {
      "cluster_interface": {
        "ip": {
          "address": "3.3.3.3"
        }
      }
    },
    {
      "cluster_interface": {
        "ip": {
          "address": "4.4.4.4"
        }
      }
    }
  ]
}
```

A re-run request with the following body details is invalid:

```
body =
{
  "name": "clusCreateRerun",
  "password": "openSesame",
  "nodes": [
    {
      "cluster_interface": {
        "ip": {
          "address": "3.3.3.3"
        }
      }
    }
  ]
}
```

Also, note that the password might already be configured. If a password is already configured and a new password is provided, this request overwrites the existing password. If a password is already configured either by another interface or by a previous POST to /cluster, any future REST requests must be authenticated with that password. If POST to /cluster with the default return_timeout of 0 returns an error, then the password was not changed.

Connection failures during cluster create

There are two cases where a request to poll the job status might fail during the cluster create job. In these cases, programmatic use of the RESTful interface should be resilient to these connection failures.

1. When the "single_node_cluster" flag is set to true, the node might reboot. During this time, the RESTful interface might refuse connections, return errors on GET, or connection timeouts might occur. Any programmatic use of the RESTful interface during reboots must consider these effects while polling a cluster create job.
2. The "mgmt_auto" LIF is removed during the cluster create job. A POST /cluster request might be issued on the "mgmt_auto" LIF. However, requests to poll the job status might fail during cluster create when the "mgmt_auto" LIF is removed. The "mgmt_auto" LIF is only removed if a cluster management interface is provided as an argument to POST /cluster, and only after the cluster management interface is created. Programmatic use of the POST /cluster API on the "mgmt_auto" LIF should be configured to dynamically switch to polling the job on the cluster management LIF.

Modifying cluster configurations

The following fields can be used to modify a cluster-wide configuration:

- name
- location
- contact
- dns_domains
- name_servers

Examples

A minimal configuration of a 2-node setup

```
# Body
body =
{
"name": "clusCreateExample1",
"password": "openSesame",
"nodes": [
  {
    "cluster_interface": {
      "ip": {
        "address": "1.1.1.1"
      }
    }
  },
  {
    "cluster_interface": {
      "ip": {
        "address": "2.2.2.2"
      }
    }
  }
]
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster" -d body
```

A single node setup with additional node configuration

```
# Body
body =
{
  "name": "clusCreateExample2",
  "password": "openSesame",
  "nodes": [
    {
      "name": "singleNode",
      "location": "Sunnyvale"
    }
  ]
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster?single_node_cluster=true" -d
body
```

Modifying a cluster-wide configuration

```
# Body
body =
{
  "contact": "it@company.com"
}

# Request
curl -X PATCH "https://<mgmt-ip>/api/cluster" -d body
```

A detailed example of a cluster "create" operation

The following is an example of how a cluster can be created using the cluster APIs. This example shows the creation of a two node cluster and uses information from the nodes themselves combined with user supplied information to configure the cluster.

1) Preparing for setup

Before the REST APIs can be issued to create the cluster, the cluster must be wired up and powered on. The network connections between the nodes for the cluster network, as well as the connections to the management network, must be completed. Once the nodes are powered up, the nodes automatically configure interfaces on the platform's default cluster ports to allow the nodes to discover each other during setup and expansion workflows. You must configure a management interface on one node or use the `mgmt_auto` LIF, which is

assigned an IP address using DHCP, to start using the REST APIs. By making a console connection to a node, the cluster setup wizard guides you through the configuration of the initial node management interface to which the REST calls can be sent. Once this step is completed, exit the wizard by typing "exit". You can then issue REST API requests.

1. Wire and power up the nodes.
 2. Make a console connection to one node to access the cluster setup wizard.
 3. Enter node management interface information to enable REST API requests to be sent to the node.
-

```
Welcome to the cluster setup wizard.
You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
This system will send event messages and periodic reports to NetApp
Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.
Enabling AutoSupport can significantly speed problem determination and
resolution should a problem occur on your system.
For further information on AutoSupport, see:
  http://support.netapp.com/autosupport/
Type yes to confirm and continue {yes}: yes
Enter the node management interface port [e0c]:
  Enter the node management interface IP address: 10.224.82.249
  Enter the node management interface netmask: 255.255.192.0
  Enter the node management interface default gateway: 10.224.64.1
  A node management interface on port e0c with IP address 10.224.82.249
has been created.
  Use your web browser to complete cluster setup by accessing
  https://10.224.82.249
  Otherwise, press Enter to complete cluster setup using the command
line
  interface: exit
  Exiting the cluster setup wizard. Any changes you made have been
saved.
  The cluster administrator's account (username "admin") password is set
to the system default.
  Warning: You have exited the cluster setup wizard before completing
all
  of the tasks. The cluster is not configured. You can complete cluster
setup by typing
  "cluster setup" in the command line interface.
```

2) Discovering the nodes

Issuing a GET /api/cluster/nodes request when the nodes are not in a cluster, the API returns a list of nodes that were discovered on the cluster network. Information returned include the node's serial number, model, software version, UUID, and cluster interface address. The number of nodes returned should be the same as

the number of nodes expected to be in the cluster. If too many nodes are discovered, remove those nodes that should not be part of the cluster. If not enough nodes are discovered, ensure all the nodes are powered up, that the connections to the cluster network are complete, and retry the command.

```
# The API:
/api/cluster/nodes

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/nodes?fields=*" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "uuid": "60277d87-19e4-11e9-ba25-005056bb6eee",
      "name": "Computer.local",
      "serial_number": "4136233-26-3",
      "model": "FAS9000",
      "version": {
        "full": "NetApp Release 9.6.0: Wed Jan 16 18:20:57 UTC 2019",
        "generation": 9,
        "major": 6,
        "minor": 0
      },
      "membership": "available",
      "cluster_interfaces": [
        {
          "ip": {
            "address": "169.254.245.113"
          }
        }
      ],
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/60277d87-19e4-11e9-ba25-005056bb6eee"
        }
      }
    },
    {
      "uuid": "8071ba1b-19e3-11e9-b003-005056bb096a",
      "name": "Computer-6.local",
      "serial_number": "4136233-26-2",
      "model": "FAS9000",
      "version": {
```

```

    "full": "NetApp Release 9.6.0: Wed Jan 16 18:20:57 UTC 2019",
    "generation": 9,
    "major": 6,
    "minor": 0
  },
  "membership": "available",
  "cluster_interfaces": [
    {
      "ip": {
        "address": "169.254.217.95"
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/8071ba1b-19e3-11e9-b003-005056bb096a"
    }
  }
},
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/cluster/nodes?fields=*"
  }
}
}

```

3) Creating the cluster

Once the node information is available, including each node's cluster interface address, the information for creating the cluster can be assembled. You must provide the cluster name and the password for the admin account. The rest of the information is optional and can be configured later using other APIs. Each node to be included in the cluster must have the cluster interface address provided so that it can be connected to while adding it to the cluster. In addition to the cluster interface address, the optional node name, location, and management interface information can be supplied. If node names are not provided, nodes are named based on the cluster name. The nodes' management interface netmask and gateway values are omitted and must be the same as the cluster management interface's netmask and gateway.

```
# The API:
/api/cluster

# The call:
curl -X POST "https://<mgmt-ip>/api/cluster" -H "accept:
application/hal+json" -H "accept: application/hal+json" -d
'{"name":"cluster1","location":"datacenter1","contact":"me","dns_domains":
["example.com"],"name_servers":["10.224.223.130","10.224.223.131","10.224.
223.132"],"ntp_servers":["time.nist.gov"],"management_interface":{"ip":{"a
ddress":"10.224.82.25","netmask":"255.255.192.0","gateway":"10.224.64.1"}}
,"password":"mypassword","license":{"keys":["AMEPOS0IKLKGEEDGNDEKSJDE"]}}
,"nodes":[{"cluster_interface":{"ip":{"address":"169.254.245.113"}}, {"name"
:"node1","management_interface":{"ip":{"address":"10.224.82.29"}}}, {"clust
er_interface":{"ip":{"address":"169.254.217.95"}}, {"name":"node2","manageme
nt_interface":{"ip":{"address":"10.224.82.31"}}}]}'

# The response:
{
  "job": {
    "uuid": "b5bc07e2-19e9-11e9-a751-005056bbd95f",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/b5bc07e2-19e9-11e9-a751-005056bbd95f"
      }
    }
  }
}
```

4) Monitoring the progress of cluster creation

To monitor the progress of the cluster create operation, the job link returned should be polled until the state value is no longer "running" or "queued".

```
# The API:
/api/cluster/jobs/b5bc07e2-19e9-11e9-a751-005056bbd95f

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/b5bc07e2-1e9-11e9-a751-005056bbd95f" -H "accept: application/hal+json"

# The response:
{
  "uuid": "b5bc07e2-19e9-11e9-a751-005056bbd95f",
  "description": "POST /api/cluster",
  "state": "success",
  "message": "success",
  "code": 0,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/b5bc07e2-19e9-11e9-a751-005056bbd95f"
    }
  }
}
```

5) Verifying the cluster information

Once the cluster is created, the information applied can be verified using a number of APIs. Most of the information provided can be retrieved using the `/api/cluster` and `/api/cluster/nodes` APIs. In addition, the network interface and route information can be viewed using the `/api/network` APIs. The following example details how to retrieve the cluster information:

```
# The API:
/api/cluster

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster" -H "accept: application/hal+json"

# The response:
{
  "name": "cluster1",
  "uuid": "93d05f83-7d80-482d-b59c-a6661d272a47",
  "location": "datacenter1",
  "contact": "me",
  "version": {
```



```
"full": "NetApp Release 9.6.0: Wed Jan 16 18:20:57 UTC 2019",
"generation": 9,
"major": 6,
"minor": 0
},
"dns_domains": [
  "example.com"
],
"name_servers": [
  "10.224.223.130",
  "10.224.223.131",
  "10.224.223.132"
],
"ntp_servers": [
  "time.nist.gov"
],
"management_interfaces": [
  {
    "uuid": "c661725a-19e9-11e9-a751-005056bbd95f",
    "name": "cluster_mgmt",
    "ip": {
      "address": "10.224.82.25"
    }
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/c661725a-19e9-11e9-a751-005056bbd95f"
      }
    }
  }
],
"metric": {
  "timestamp": "2019-04-09T06:33:30Z",
  "duration": "PT15S",
  "status": "ok",
  "latency": {
    "other": 0,
    "total": 525,
    "read": 525,
    "write": 0
  },
  "iops": {
    "read": 200,
    "write": 0,
    "other": 0,
    "total": 200
  }
}
```

```
{
  "throughput": {
    "read": 820838,
    "write": 0,
    "other": 0,
    "total": 820838
  },
  "statistics": {
    "timestamp": "2019-04-09T06:33:50Z",
    "status": "ok",
    "latency_raw": {
      "other": 38928,
      "total": 3331918704,
      "read": 3331879776,
      "write": 0
    },
    "iops_raw": {
      "read": 6188132,
      "write": 0,
      "other": 5,
      "total": 6188137
    },
    "throughput_raw": {
      "read": 25346587876,
      "write": 0,
      "other": 0,
      "total": 25346587876
    }
  },
  "_links": {
    "self": {
      "href": "/api/cluster"
    }
  }
}
```

Retrieve the cluster configuration

GET /cluster

Retrieves the cluster configuration.

Learn more

- [DOC /cluster](#)

Parameters

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
configuration_backup	configuration_backup	
contact	string	
dns_domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none">• The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-" or "_" .• The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9.• The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9.• The top level domain must contain only the following characters: A through Z, a through z.• The system reserves the following names: "all", "local", and "localhost".
license	license	License keys or NLF contents.

Name	Type	Description
location	string	
management_interface	management_interface	The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.
management_interfaces	array[management_interfaces]	
metric	metric	Performance numbers, such as IOPS latency and throughput.
name	string	
name_servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
nodes	array[nodes]	
ntp_servers	array[string]	Host name, IPv4 address, or IPv6 address for the external NTP time servers.
password	string	Initial admin password used to create the cluster.
statistics	statistics	These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "configuration_backup": {
    "password": "yourpassword",
    "url": "http://10.224.65.198/backups",
    "username": "me"
  },
  "contact": "<a href="
mailto:support@company.com">support@company.com</a>",
  "dns_domains": [
    "example.com",
    "example2.example3.com"
  ],
  "license": {
    "keys": {
    }
  },
  "location": "building 1",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    }
  },
  "management_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "metric": {
    "_links": {
      "self": {
```

```

        "href": "/api/resourcelink"
    },
    },
    "duration": "PT15S",
    "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
    },
    "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
    },
    "status": "ok",
    "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
},
"name": "cluster1",
"name_servers": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
],
"nodes": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "cluster_interface": {
        "ip": {
            "address": "10.10.10.7"
        }
    },
    "cluster_interfaces": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "ip": {
            "address": "10.10.10.7"
        }
    }
}

```

```

    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "controller": {
    "flash_cache": {
      "capacity": 102400000000,
      "firmware_version": "NA05",
      "hardware_revision": "A1",
      "model": "X1970A",
      "part_number": "119-00207",
      "serial_number": "A22P5061550000187",
      "slot": "6-1",
      "state": "ok"
    },
    "frus": {
      "id": 0,
      "state": "ok",
      "type": "fan"
    },
    "over_temperature": "over"
  },
  "date": "2017-01-25 11:20:13 +0400",
  "ha": {
    "partners": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "location": "rack 2 row 5",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "management_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },

```

```

    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "membership": "available",
  "model": "FAS3070",
  "name": "node-01",
  "serial_number": "4048820-60-9",
  "service_processor": {
    "firmware_version": "string",
    "ipv4_interface": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    },
    "ipv6_interface": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    },
    "link_status": "up",
    "mac_address": "string",
    "state": "online"
  },
  "uptime": 300536,
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
  "version": {
    "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
    "generation": 9,
    "major": 4,
    "minor": 0
  }
},
"ntp_servers": [
  "time.nist.gov",
  "10.98.19.20",
  "2610:20:6F15:15::27"
],
"password": "mypassword",
"statistics": {
  "iops_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  }
}

```



```

    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "version": {
    "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
    "generation": 9,
    "major": 4,
    "minor": 0
  }
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

configuration_backup

Name	Type	Description
password	string	
url	string	An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster.
username	string	

license

License keys or NLF contents.

Name	Type	Description
keys	array[string]	

ip

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.

Name	Type	Description
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

management_interface

The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.

Name	Type	Description
ip	ip	Object to setup an interface along with its default router.

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

management_interfaces

A network interface. Either UUID or name may be supplied on input.

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.

Name	Type	Description
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes does not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

node_setup_ip

The IP configuration for cluster setup.

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interface

The cluster network IP address of the node to be added.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

cluster_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

flash_cache

Name	Type	Description
capacity	integer	Size in bytes
firmware_version	string	
hardware_revision	string	
model	string	
part_number	string	
serial_number	string	
slot	string	
state	string	

frus

Name	Type	Description
id	integer	
state	string	
type	string	

controller

Controller information

Name	Type	Description
flash_cache	array[flash_cache]	A list of Flash-Cache devices. Only returned when requested by name.

Name	Type	Description
frus	array[frus]	A list of frus in the node. Only returned when requested by name.
over_temperature	string	Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds.

partners

Name	Type	Description
_links	_links	
name	string	
uuid	string	

ha

Name	Type	Description
auto_giveback	boolean	Specifies whether giveback is automatically initiated when the node that owns the storage is ready.
enabled	boolean	Specifies whether or not storage failover is enabled.
partners	array[partners]	The nodes in this node's High Availability (HA) group.

management_interface

The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

management_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

ipv4_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

ipv6_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

service_processor

Name	Type	Description
dhcp_enabled	boolean	Set to true to use DHCP to configure an IPv4 interface.
firmware_version	string	The version of firmware installed.
ipv4_interface	ipv4_interface	Object to setup an interface along with its default router.
ipv6_interface	ipv6_interface	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	
state	string	

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

nodes

Complete node information

Name	Type	Description
_links	_links	
cluster_interface	cluster_interface	The cluster network IP address of the node to be added.
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information

Name	Type	Description
date	string	Specifies the ISO-8601 format date and time on the node.
ha	ha	
location	string	
management_interface	management_interface	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - If a node is available, this means it is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
model	string	
name	string	

Name	Type	Description
serial_number	string	
service_processor	service_processor	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes does not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the cluster configuration

PATCH `/cluster`

Updates the cluster configuration once the cluster has been created.

Learn more

- [DOC /cluster](#)

Request Body

Name	Type	Description
<code>_links</code>	_links	
<code>configuration_backup</code>	configuration_backup	
<code>contact</code>	string	

Name	Type	Description
dns_domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> • The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-" or "_". • The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9. • The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9. • The top level domain must contain only the following characters: A through Z, a through z. • The system reserves the following names: "all", "local", and "localhost".
license	license	License keys or NLF contents.
location	string	
management_interface	management_interface	The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.
management_interfaces	array[management_interfaces]	
metric	metric	Performance numbers, such as IOPS latency and throughput.
name	string	
name_servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
nodes	array[nodes]	

Name	Type	Description
ntp_servers	array[string]	Host name, IPv4 address, or IPv6 address for the external NTP time servers.
password	string	Initial admin password used to create the cluster.
statistics	statistics	These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "configuration_backup": {
    "password": "yourpassword",
    "url": "http://10.224.65.198/backups",
    "username": "me"
  },
  "contact": "<a href="
mailto:support@company.com">support@company.com</a>",
  "dns_domains": [
    "example.com",
    "example2.example3.com"
  ],
  "license": {
    "keys": {
    }
  },
  "location": "building 1",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    }
  },
  "management_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "metric": {
    "_links": {
      "self": {
```

```

        "href": "/api/resourcelink"
    },
    },
    "duration": "PT15S",
    "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
    },
    "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
    },
    "status": "ok",
    "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
},
"name": "cluster1",
"name_servers": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
],
"nodes": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "cluster_interface": {
        "ip": {
            "address": "10.10.10.7"
        }
    },
    "cluster_interfaces": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "ip": {
            "address": "10.10.10.7"
        }
    }
}

```

```

    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "controller": {
    "flash_cache": {
      "capacity": 102400000000,
      "firmware_version": "NA05",
      "hardware_revision": "A1",
      "model": "X1970A",
      "part_number": "119-00207",
      "serial_number": "A22P5061550000187",
      "slot": "6-1",
      "state": "ok"
    },
    "frus": {
      "id": 0,
      "state": "ok",
      "type": "fan"
    },
    "over_temperature": "over"
  },
  "date": "2017-01-25 11:20:13 +0400",
  "ha": {
    "partners": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "location": "rack 2 row 5",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "management_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },

```

```

    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "membership": "available",
  "model": "FAS3070",
  "name": "node-01",
  "serial_number": "4048820-60-9",
  "service_processor": {
    "firmware_version": "string",
    "ipv4_interface": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    },
    "ipv6_interface": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    },
    "link_status": "up",
    "mac_address": "string",
    "state": "online"
  },
  "uptime": 300536,
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
  "version": {
    "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
    "generation": 9,
    "major": 4,
    "minor": 0
  }
},
"ntp_servers": [
  "time.nist.gov",
  "10.98.19.20",
  "2610:20:6F15:15::27"
],
"password": "mypassword",
"statistics": {
  "iops_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  }
}

```

```

    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "version": {
    "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
    "generation": 9,
    "major": 4,
    "minor": 0
  }
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
8847362	Too many name servers provided.
8847361	Too many DNS domains provided.
9240587	A name must be provided.
131727388	Hostnames for NTP servers cannot be used without DNS configured.
2097165	An NTP server could not be reached.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

configuration_backup

Name	Type	Description
password	string	
url	string	An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster.
username	string	

license

License keys or NLF contents.

Name	Type	Description
keys	array[string]	

ip

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.

Name	Type	Description
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

management_interface

The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.

Name	Type	Description
ip	ip	Object to setup an interface along with its default router.

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

management_interfaces

A network interface. Either UUID or name may be supplied on input.

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.

Name	Type	Description
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes does not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

node_setup_ip

The IP configuration for cluster setup.

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interface

The cluster network IP address of the node to be added.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

cluster_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

flash_cache

Name	Type	Description
capacity	integer	Size in bytes
firmware_version	string	
hardware_revision	string	
model	string	
part_number	string	
serial_number	string	
slot	string	
state	string	

frus

Name	Type	Description
id	integer	
state	string	
type	string	

controller

Controller information

Name	Type	Description
flash_cache	array[flash_cache]	A list of Flash-Cache devices. Only returned when requested by name.

Name	Type	Description
frus	array[frus]	A list of frus in the node. Only returned when requested by name.
over_temperature	string	Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds.

partners

Name	Type	Description
_links	_links	
name	string	
uuid	string	

ha

Name	Type	Description
auto_giveback	boolean	Specifies whether giveback is automatically initiated when the node that owns the storage is ready.
enabled	boolean	Specifies whether or not storage failover is enabled.
partners	array[partners]	The nodes in this node's High Availability (HA) group.

management_interface

The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

management_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

ipv4_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

ipv6_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

service_processor

Name	Type	Description
dhcp_enabled	boolean	Set to true to use DHCP to configure an IPv4 interface.
firmware_version	string	The version of firmware installed.
ipv4_interface	ipv4_interface	Object to setup an interface along with its default router.
ipv6_interface	ipv6_interface	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	
state	string	

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

nodes

Complete node information

Name	Type	Description
_links	_links	
cluster_interface	cluster_interface	The cluster network IP address of the node to be added.
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information

Name	Type	Description
date	string	Specifies the ISO-8601 format date and time on the node.
ha	ha	
location	string	
management_interface	management_interface	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - If a node is available, this means it is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
model	string	
name	string	

Name	Type	Description
serial_number	string	
service_processor	service_processor	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes does not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

cluster

Complete cluster information

Name	Type	Description
_links	_links	
configuration_backup	configuration_backup	
contact	string	
dns_domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> • The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_". • The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9. • The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9. • The top level domain must contain only the following characters: A through Z, a through z. • The system reserves the following names: "all", "local", and "localhost".
license	license	License keys or NLF contents.
location	string	
management_interface	management_interface	The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.
management_interfaces	array[management_interfaces]	
metric	metric	Performance numbers, such as IOPS latency and throughput.
name	string	

Name	Type	Description
name_servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
nodes	array[nodes]	
ntp_servers	array[string]	Host name, IPv4 address, or IPv6 address for the external NTP time servers.
password	string	Initial admin password used to create the cluster.
statistics	statistics	These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a cluster

POST /cluster

Sets up a cluster.

Required properties

- name
- password

Recommended optional properties

- location
- contact
- dns_domains
- name_servers
- ntp_servers
- license
- configuration_backup
- management_interface
- nodes

Learn more

- [DOC /cluster](#)

Parameters

Name	Type	In	Required	Description
single_node_cluster	boolean	query	False	Configures a single node cluster. All cluster ports are reassigned to the default network. The storage failover settings are configured to non-HA. The node reboots during this operation.

Request Body

Name	Type	Description
_links	_links	
configuration_backup	configuration_backup	
contact	string	
dns_domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> • The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_". • The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9. • The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9. • The top level domain must contain only the following characters: A through Z, a through z. • The system reserves the following names: "all", "local", and "localhost".
license	license	License keys or NLF contents.

Name	Type	Description
location	string	
management_interface	management_interface	The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.
management_interfaces	array[management_interfaces]	
metric	metric	Performance numbers, such as IOPS latency and throughput.
name	string	
name_servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
nodes	array[nodes]	
ntp_servers	array[string]	Host name, IPv4 address, or IPv6 address for the external NTP time servers.
password	string	Initial admin password used to create the cluster.
statistics	statistics	These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "configuration_backup": {
    "password": "yourpassword",
    "url": "http://10.224.65.198/backups",
    "username": "me"
  },
  "contact": "<a href="
mailto:support@company.com">support@company.com</a>",
  "dns_domains": [
    "example.com",
    "example2.example3.com"
  ],
  "license": {
    "keys": {
    }
  },
  "location": "building 1",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    }
  },
  "management_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "metric": {
    "_links": {
      "self": {
```

```

        "href": "/api/resourcelink"
    },
    },
    "duration": "PT15S",
    "iops": {
        "read": 200,
        "total": 1000,
        "write": 100
    },
    "latency": {
        "read": 200,
        "total": 1000,
        "write": 100
    },
    "status": "ok",
    "throughput": {
        "read": 200,
        "total": 1000,
        "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
},
"name": "cluster1",
"name_servers": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
],
"nodes": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "cluster_interface": {
        "ip": {
            "address": "10.10.10.7"
        }
    },
    "cluster_interfaces": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "ip": {
            "address": "10.10.10.7"
        }
    }
}

```

```

    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "controller": {
    "flash_cache": {
      "capacity": 102400000000,
      "firmware_version": "NA05",
      "hardware_revision": "A1",
      "model": "X1970A",
      "part_number": "119-00207",
      "serial_number": "A22P5061550000187",
      "slot": "6-1",
      "state": "ok"
    },
    "frus": {
      "id": 0,
      "state": "ok",
      "type": "fan"
    },
    "over_temperature": "over"
  },
  "date": "2017-01-25 11:20:13 +0400",
  "ha": {
    "partners": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "location": "rack 2 row 5",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "management_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },

```

```

    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "membership": "available",
  "model": "FAS3070",
  "name": "node-01",
  "serial_number": "4048820-60-9",
  "service_processor": {
    "firmware_version": "string",
    "ipv4_interface": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    },
    "ipv6_interface": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    },
    "link_status": "up",
    "mac_address": "string",
    "state": "online"
  },
  "uptime": 300536,
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
  "version": {
    "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
    "generation": 9,
    "major": 4,
    "minor": 0
  }
},
"ntp_servers": [
  "time.nist.gov",
  "10.98.19.20",
  "2610:20:6F15:15::27"
],
"password": "mypassword",
"statistics": {
  "iops_raw": {
    "read": 200,
    "total": 1000,
    "write": 100
  }
}

```

```

    },
    "latency_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput_raw": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "version": {
    "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
    "generation": 9,
    "major": 4,
    "minor": 0
  }
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
131727360	A node could not be added to the cluster. This is a generic code, see response message for details.
8978433	An invalid license key was provided.
8847362	Too many name servers provided.
8847361	Too many DNS domains provided.
9240587	A name must be provided.
39387137	The URL provided is invalid.
131727389	URL and username are required for configuration backup.
262245	The value provided is invalid.
1179817	The IP address, netmask, and gateway must all be provided for cluster management interface.
1179813	Fields set for one node must be set for all nodes.
1179818	The IP address and gateway must be of the same family.
1179821	An IP address and netmask conflicts with an existing entry.
131727388	Hostnames for NTP servers cannot be used without DNS configured.

Error Code	Description
2097165	An NTP server could not be reached.
1179825	All management and cluster config IP addresses must belong to the same address family.
8847394	An invalid DNS domain was provided.
9240594	An invalid name was provided.
1179824	An invalid gateway was provided.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

configuration_backup

Name	Type	Description
password	string	
url	string	An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster.
username	string	

license

License keys or NLF contents.

Name	Type	Description
keys	array[string]	

ip

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.

Name	Type	Description
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

management_interface

The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.

Name	Type	Description
ip	ip	Object to setup an interface along with its default router.

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

management_interfaces

A network interface. Either UUID or name may be supplied on input.

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.

Name	Type	Description
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes does not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

node_setup_ip

The IP configuration for cluster setup.

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interface

The cluster network IP address of the node to be added.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

cluster_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

flash_cache

Name	Type	Description
capacity	integer	Size in bytes
firmware_version	string	
hardware_revision	string	
model	string	
part_number	string	
serial_number	string	
slot	string	
state	string	

frus

Name	Type	Description
id	integer	
state	string	
type	string	

controller

Controller information

Name	Type	Description
flash_cache	array[flash_cache]	A list of Flash-Cache devices. Only returned when requested by name.

Name	Type	Description
frus	array[frus]	A list of frus in the node. Only returned when requested by name.
over_temperature	string	Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds.

partners

Name	Type	Description
_links	_links	
name	string	
uuid	string	

ha

Name	Type	Description
auto_giveback	boolean	Specifies whether giveback is automatically initiated when the node that owns the storage is ready.
enabled	boolean	Specifies whether or not storage failover is enabled.
partners	array[partners]	The nodes in this node's High Availability (HA) group.

management_interface

The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

management_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

ipv4_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

ipv6_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

service_processor

Name	Type	Description
dhcp_enabled	boolean	Set to true to use DHCP to configure an IPv4 interface.
firmware_version	string	The version of firmware installed.
ipv4_interface	ipv4_interface	Object to setup an interface along with its default router.
ipv6_interface	ipv6_interface	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	
state	string	

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

nodes

Complete node information

Name	Type	Description
_links	_links	
cluster_interface	cluster_interface	The cluster network IP address of the node to be added.
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information

Name	Type	Description
date	string	Specifies the ISO-8601 format date and time on the node.
ha	ha	
location	string	
management_interface	management_interface	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - If a node is available, this means it is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
model	string	
name	string	

Name	Type	Description
serial_number	string	
service_processor	service_processor	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes does not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

cluster

Complete cluster information

Name	Type	Description
_links	_links	
configuration_backup	configuration_backup	
contact	string	
dns_domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> • The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_". • The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9. • The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9. • The top level domain must contain only the following characters: A through Z, a through z. • The system reserves the following names: "all", "local", and "localhost".
license	license	License keys or NLF contents.
location	string	
management_interface	management_interface	The management interface of the cluster. The netmask and gateway for this interface are used for the node management interfaces provided in the node configuration.
management_interfaces	array[management_interfaces]	
metric	metric	Performance numbers, such as IOPS latency and throughput.
name	string	

Name	Type	Description
name_servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
nodes	array[nodes]	
ntp_servers	array[string]	Host name, IPv4 address, or IPv6 address for the external NTP time servers.
password	string	Initial admin password used to create the cluster.
statistics	statistics	These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve cluster chassis

Cluster chassis endpoint overview

Retrieving chassis information

The chassis GET API retrieves all of the chassis information in the cluster.

Examples

1) Retrieve a list of chassis from the cluster

The following example shows the response with a list of chassis in the cluster:

```
# The API:
/api/cluster/chassis

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/chassis" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "id": "021352005981",
      "_links": {
        "self": {
          "href": "/api/cluster/chassis/021352005981"
        }
      }
    },
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/cluster/chassis"
    }
  }
}
```

2) Retrieve a specific chassis from the cluster

The following example shows the response of the requested chassis. If there is no chassis with the requested id, an error is returned.

```
# The API:
/api/cluster/chassis/{id}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/chassis/021352005981" -H
"accept: application/hal+json"

# The response:
{
  "id": "021352005981",
```

```
"state": "ok",
"nodes": [
  {
    "name": "node-1",
    "uuid": "6ede364b-c3d0-11e8-a86a-00a098567f31",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/6ede364b-c3d0-11e8-a86a-00a098567f31"
      }
    }
  }
],
"frus": [
  {
    "id": "PSU2",
    "type": "psu",
    "state": "ok"
  },
  {
    "id": "PSU1",
    "type": "psu",
    "state": "ok"
  },
  {
    "id": "Fan2",
    "type": "fan",
    "state": "ok"
  },
  {
    "id": "Fan3",
    "type": "fan",
    "state": "ok"
  },
  {
    "id": "Fan1",
    "type": "fan",
    "state": "ok"
  }
],
"_links": {
  "self": {
    "href": "/api/cluster/chassis/021352005981"
  }
}
}
```

Retrieve a collection of chassis

GET /cluster/chassis

Retrieves a collection of chassis.

Related ONTAP commands

- `system chassis show`
- `system chassis fru show`

Learn more

- [DOC /cluster/chassis](#)

Parameters

Name	Type	In	Required	Description
frus.type	string	query	False	Filter by frus.type
frus.state	string	query	False	Filter by frus.state
frus.id	string	query	False	Filter by frus.id
nodes.name	string	query	False	Filter by nodes.name
nodes.uuid	string	query	False	Filter by nodes.uuid
state	string	query	False	Filter by state
shelves.uid	string	query	False	Filter by shelves.uid
id	string	query	False	Filter by id
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.

Name	Type	In	Required	Description
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[chassis]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "frus": {
      "state": "ok",
      "type": "fan"
    },
    "id": "021352005981",
    "nodes": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "shelves": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "uid": 7777841915827391056
    },
    "state": "ok"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

frus

Name	Type	Description
id	string	
state	string	
type	string	

_links

Name	Type	Description
self	href	

nodes

Name	Type	Description
_links	_links	
name	string	
uuid	string	

shelf_reference

Shelf

Name	Type	Description
_links	_links	
uid	string	

chassis

Name	Type	Description
frus	array[frus]	List of frus in chassis
id	string	
nodes	array[nodes]	List of nodes in chassis
shelves	array[shelf_reference]	List of shelves in chassis
state	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a chassis

GET `/cluster/chassis/{id}`

Retrieves a specific chassis.

Related ONTAP commands

- `system chassis show`
- `system chassis fru show`

Learn more

- [DOC /cluster/chassis](#)

Parameters

Name	Type	In	Required	Description
id	string	path	True	Chassis ID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
frus	array[frus]	List of frus in chassis
id	string	
nodes	array[nodes]	List of nodes in chassis
shelves	array[shelf_reference]	List of shelves in chassis
state	string	

Example response

```
{
  "frus": {
    "state": "ok",
    "type": "fan"
  },
  "id": "021352005981",
  "nodes": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "shelves": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uid": 7777841915827391056
  },
  "state": "ok"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

frus

Name	Type	Description
id	string	
state	string	
type	string	

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

nodes

Name	Type	Description
_links	_links	
name	string	
uuid	string	

shelf_reference

Shelf

Name	Type	Description
_links	_links	
uid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

View and manage cluster jobs

Cluster jobs endpoint overview

Summary

This API is used to view and manipulate jobs. Jobs provide information about asynchronous operations. Some long-running jobs are paused or cancelled by calling PATCH. Individual operations will mention if they support PATCH on the job. Once a job transitions to a terminal state, it is deleted after a default time of 300 seconds. Attempts to GET or PATCH the job will return a 404 error code once the job has been deleted.

Example

The following examples show how to retrieve and update a job state

1) Retrieve job information

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/b5145e1d-b53b-11e8-8252-005056bbd8f5" -H "accept: application/json"

# The response:
{
  "uuid": "b5145e1d-b53b-11e8-8252-005056bbd8f5",
  "code": 0,
  "description": "Cluster Backup Job",
  "state": "running",
  "message": "creating_node_backups",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/b5145e1d-b53b-11e8-8252-005056bbd8f5"
    }
  }
}
```

2) Update a job that supports the new state

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/cluster/jobs/b5145e1d-b53b-11e8-8252-005056bbd8f5?action=cancel" -H "accept: application/json"
```

Retrieve recent asynchronous jobs

GET /cluster/jobs

Retrieves a list of recently running asynchronous jobs. Once a job transitions to a failure or success state, it is deleted after a default time of 300 seconds.

Learn more

- [DOC /cluster/jobs](#)

Parameters

Name	Type	In	Required	Description
state	string	query	False	Filter by state
start_time	string	query	False	Filter by start_time
end_time	string	query	False	Filter by end_time
code	integer	query	False	Filter by code
uuid	string	query	False	Filter by uuid
description	string	query	False	Filter by description
message	string	query	False	Filter by message
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	
records	array[job]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "code": 0,
    "description": "App Snapshot Job",
    "end_time": "string",
    "message": "Complete: Successful",
    "start_time": "string",
    "state": "queued",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

job

Name	Type	Description
_links	_links	
code	integer	If the state indicates "failure", this is the final error code.
description	string	The description of the job to help identify it independent of the UUID.
end_time	string	The time the job ended.
message	string	A message corresponding to the state of the job providing additional details about the current state.
start_time	string	The time the job started.
state	string	The state of the job.
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve details of an asynchronous job

GET `/cluster/jobs/{uuid}`

Retrieves the details of a specific asynchronous job. Once a job transitions to a failure or success state, it is deleted after a default time of 300 seconds.

Learn more

- [DOC /cluster/jobs](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Job UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	

Name	Type	Description
code	integer	If the state indicates "failure", this is the final error code.
description	string	The description of the job to help identify it independent of the UUID.
end_time	string	The time the job ended.
message	string	A message corresponding to the state of the job providing additional details about the current state.
start_time	string	The time the job started.
state	string	The state of the job.
uuid	string	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "code": 0,
  "description": "App Snapshot Job",
  "end_time": "string",
  "message": "Complete: Successful",
  "start_time": "string",
  "state": "queued",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the state of an asynchronous job


PATCH `/cluster/jobs/{uuid}`

Updates the state of a specific asynchronous job.

Learn more

- [DOC /cluster/jobs](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Job UUID
action	string	query	False	<p>Request a job to pause, resume, or cancel.</p> <div>  <p>Not all jobs support these actions. A job can only be resumed if it is in paused state. Upon successfully requesting a job to be cancelled, the job state changes to either success or failure.</p> <ul style="list-style-type: none"> enum: ["pause", "resume", "cancel"] </div>

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
459753	Command execution failed with custom error from the program.
458762	Job is already in a terminal state.
458773	The Job Manager is not initialized.
458771	The specified job is running.
458776	The specified job is not currently running.
458783	This job does not support pause.
458784	This job does not support cancel.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage cluster licensing

Cluster licensing licenses endpoint overview

Overview

Licensing allows you to tailor a system to meet an organization's specific needs. New features can be enabled by purchasing a license from a NetApp sales associate. After installation of the license, the new feature is available immediately.

This interface manages licenses according to their supported feature. By default, the interface displays packages with installed licenses, but you can also return unlicensed packages.

Each feature has a compliance state which is indicated at the package level. Individual licenses also contain a compliance state indicated in the 'licenses' array. The state of the package is determined by analyzing the underlying licenses according to the following criteria:

- Licensing terms
- Cluster state

Licensing terms

The licensing terms define the conditions under which a package is considered 'compliant'. Individual licenses are evaluated based on the following:

- Scope

- Time period
- Usage

Scope

A package can be licensed under the following scopes:

- site
- cluster
- node

A package licensed under 'site' or 'cluster' permits the feature to be used by any node that is a member of the cluster.

A package licensed under 'node' scope permits the authorized node to use the feature. Within a cluster, if you haven't supplied every node with a valid license, the package state will indicate 'noncompliant'. A license must be purchased for each node in a cluster for the package to be considered 'compliant'.

Time period

Some package licenses are only valid for a limited period of time. After a license has expired, the package state changes to 'noncompliant'. A new license will need to be purchased for the package to return to a 'compliant' state.

Usage

Some package licenses have additional terms that need to be maintained to keep a license in compliance. These conditions are defined by the individual license. For example, a license might define the maximum amount of storage that a node can allocate for the license to be 'compliant'.

Cluster state

A cluster's state consists of the following:

- Node online status
- Node cluster membership

Some features require that a node be online to display a valid compliance state. If a node can not be reached, or is not known to the cluster, the individual license may indicate an 'unknown' state.

Licensing keys

A license is issued in one of the following two formats:

- 26-character key
- NetApp License File (NLF)

The following is an example of a 26-character key:

```
AMEPOSIOIKLKGEEDGNDEKSJDE
```

The following is an example of a NLF key:

```
{
  "statusResp": {
    "version": "1",
    "serialNumber": "123456789",
    "message": "Success",
    "licenses": {
      "capacity": "1",
      "type": "capacity",
      "licenseProtocol": "FABRICPOOL-TB",
      "package": "FabricPool",
      "licenseScope": "cluster"
    },
    "snStatus": "Active",
    "product": "fabricpool",
    "statusCode": "S007"
  },
  "Signature": "signatureABC"
}
```

Either format can be submitted, via this API, to enable features.

Examples

Retrieving a collection of licenses organized by package

This example retrieves a collection that contains one entry for each package (filtered to only the 'fabricpool' package).

```

# API
GET /cluster/licensing/licenses/?fields=*&name=fabricpool"

# Response
200 OK

# JSON Body
{
  "records": [
    {
      "name": "fabricpool",
      "scope": "cluster",
      "state": "compliant",
      "licenses": [
        {
          "owner": "testcluster-1",
          "serial_number": "4149027342",
          "state": "compliant",
          "capacity": {
            "maximum_size": 1099511627776,
            "used_size": 0
          }
        }
      ],
      "_links": {
        "self": {
          "href": "/api/cluster/licensing/licenses/fabricpool"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/cluster/licensing/licenses/?fields=*&name=fabricpool"
    }
  }
}

```

Retrieving a collection of installed licenses

This example retrieves a collection containing all packages (except base) that have installed licenses.

```

# API
GET /cluster/licensing/licenses/?fields=*&name=!base

```

```
# Response
200 OK

# JSON Body
{
  "records": [
    {
      "name": "nfs",
      "scope": "node",
      "state": "compliant",
      "licenses": [
        {
          "owner": "testcluster-1",
          "serial_number": "1-81-0000000000000004149027492",
          "state": "compliant"
        }
      ],
      "_links": {
        "self": {
          "href": "/api/cluster/licensing/licenses/nfs"
        }
      }
    },
    {
      "name": "cifs",
      "scope": "node",
      "state": "compliant",
      "licenses": [
        {
          "owner": "testcluster-1",
          "serial_number": "1-81-0000000000000004149027492",
          "state": "compliant"
        }
      ],
      "_links": {
        "self": {
          "href": "/api/cluster/licensing/licenses/cifs"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/cluster/licensing/licenses/?fields=*&name=!base"
    }
  }
}
```

```
}  
}  
}
```

Installing a NLF license

This example installs a single license in the NLF format.



You must escape all double quotes and backslash characters, of the JSON license, before it can be placed in the POST request.

```
# API  
POST /cluster/licensing/licenses/  
  
# JSON Body  
{  
  "keys" : [ "{\\"statusResp\\":{\\"snStatus\\": \\"Active\\", \\"licenses\\":  
    {\\"package\\": \\"FabricPool\\", \\"capacity\\": \\"1\\", \\"licenseProtocol\\":  
    \\"FABRICPOOL-TB\\", \\"type\\": \\"capacity\\", \\"licenseScope\\": \\"cluster\\"},  
    \\"message\\": \\"Success\\", \\"statusCode\\": \\"S007\\", \\"version\\": \\"1\\",  
    \\"product\\": \\"fabricpool\\", \\"serialNumber\\": \\"4149027342\\"},  
    \\"Signature\\":\\"SignatureABC\\"}" ]  
}  
  
# Response  
201 Created
```

Installing a 26-character key

This example installs a single 26-character key formatted license.

```
# API  
POST /cluster/licensing/licenses/  
  
# JSON Body  
{  
  "keys" : [ "AAAAAAAAAAAAAAAAAAAAAAAAAAAA" ]  
}  
  
# Response  
201 Created
```

Installing multiple licenses with one API call

This example demonstrates how multiple keys can be provided to install multiple features in a single API call.


```
# API
POST /cluster/licensing/licenses/

# JSON Body
{
  "keys" : [ "AAAAAAAAAAAAAAAAAAAAAAAAAAAA",
             "BBBBBBBBBBBBBBBBBBBBBBBBBBBB" ]
}

# Response
201 Created
```

Retrieve license packages

GET /cluster/licensing/licenses

Retrieves a collection of license packages.

Related ONTAP commands

- `system license show-status`
- `system license show`

Learn more

- [DOC /cluster/licensing/licenses](#)

Parameters

Name	Type	In	Required	Description
scope	string	query	False	Filter by scope
licenses.expiry_time	string	query	False	Filter by licenses.expiry_time
licenses.compliance.state	string	query	False	Filter by licenses.compliance.state
licenses.capacity.used_size	integer	query	False	Filter by licenses.capacity.used_size
licenses.capacity.maximum_size	integer	query	False	Filter by licenses.capacity.maximum_size

Name	Type	In	Required	Description
licenses.serial_number	string	query	False	Filter by licenses.serial_number
licenses.active	boolean	query	False	Filter by licenses.active
licenses.owner	string	query	False	Filter by licenses.owner
licenses.evaluation	boolean	query	False	Filter by licenses.evaluation
licenses.start_time	string	query	False	Filter by licenses.start_time
state	string	query	False	Filter by state
name	string	query	False	Filter by name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

Name	Type	In	Required	Description
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    },
    "keys": {
    },
    "licenses": {
      "capacity": {
        "maximum_size": 0,
        "used_size": 0
      },
      "compliance": {
        "state": "compliant"
      },
      "expiry_time": "2019-03-02 19:00:00 UTC",
      "owner": "cluster1",
      "serial_number": "123456789",
      "start_time": "2019-02-02 19:00:00 UTC"
    },
    "name": "NFS",
    "scope": "not_available",
    "state": "compliant"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

capacity

Name	Type	Description
maximum_size	integer	Licensed capacity size (in bytes) that can be used.
used_size	integer	Capacity that is currently used (in bytes).

compliance

Name	Type	Description
state	string	Compliance state of the license.

licenses

Name	Type	Description
active	boolean	A flag indicating whether the license is currently being enforced.
capacity	capacity	
compliance	compliance	
evaluation	boolean	A flag indicating whether the license is in evaluation mode.

Name	Type	Description
expiry_time	string	Date and time when the license expires.
owner	string	Cluster, node or license manager that owns the license.
serial_number	string	Serial number of the license.
start_time	string	Date and time when the license starts.

records

Name	Type	Description
_links	_links	
keys	array[string]	
licenses	array[licenses]	Installed licenses of the package.
name	string	Name of the license.
scope	string	Scope of the license.
state	string	Summary state of package based on all installed licenses.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Install one or more feature licenses

POST `/cluster/licensing/licenses`

Installs one or more feature licenses.

Required properties

- `keys` - Array containing a list of NLF or 26-character license keys.

Related ONTAP commands

- `system license add`

Learn more

- [DOC /cluster/licensing/licenses](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned.

Request Body

Name	Type	Description
<code>_links</code>	_links	
<code>keys</code>	array[string]	
<code>licenses</code>	array[licenses]	Installed licenses of the package.
<code>name</code>	string	Name of the license.
<code>scope</code>	string	Scope of the license.
<code>state</code>	string	Summary state of package based on all installed licenses.

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "keys": {
  },
  "licenses": {
    "capacity": {
      "maximum_size": 0,
      "used_size": 0
    },
    "compliance": {
      "state": "compliant"
    },
    "expiry_time": "2019-03-02 19:00:00 UTC",
    "owner": "cluster1",
    "serial_number": "123456789",
    "start_time": "2019-02-02 19:00:00 UTC"
  },
  "name": "NFS",
  "scope": "not_available",
  "state": "compliant"
}
```

Response

Status: 201, Created

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "keys": {
    },
    "licenses": {
      "capacity": {
        "maximum_size": 0,
        "used_size": 0
      },
      "compliance": {
        "state": "compliant"
      },
      "expiry_time": "2019-03-02 19:00:00 UTC",
      "owner": "cluster1",
      "serial_number": "123456789",
      "start_time": "2019-02-02 19:00:00 UTC"
    },
    "name": "NFS",
    "scope": "not_available",
    "state": "compliant"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
1115117	Generic licensing error
1115122	No cluster serial number found
1115124	No node serial number found
1115130	No license code was provided
1115131	Installation of the license failed
1115132	License already exists on system
1115134	Serial number does not belong to node
1115141	License data is invalid
1115142	License signature is invalid
1115143	Internal error applying the requested license
1115152	License does not apply to the platform
1115154	Unable to retrieve cluster ID
1115155	Invalid cluster ID found
1115159	License is not in an acceptable format
1115164	Minimum ONTAP version requirements not met
1115179	FlexCache is not supported on this system
1115180	FlexCache is not supported on cloud systems
1115407	Capacity pool licenses cannot be installed directly
66846818	Failed to interpret FlexCache license information
66846821	FlexCache is not supported on cloud systems
66846822	Invalid FlexCache capacity information provided
655294464	Failed to extract license contents
655294465	License key is invalid
655294466	Serial number is invalid
655294467	Version number is invalid
655294468	Expired license
655294469	License does not apply to the platform
655294470	License does not apply to the product

Name	Type	Description
errors	array[error]	

Example error

```
{
  "errors": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

capacity

Name	Type	Description
maximum_size	integer	Licensed capacity size (in bytes) that can be used.
used_size	integer	Capacity that is currently used (in bytes).

compliance

Name	Type	Description
state	string	Compliance state of the license.

licenses

Name	Type	Description
active	boolean	A flag indicating whether the license is currently being enforced.
capacity	capacity	
compliance	compliance	
evaluation	boolean	A flag indicating whether the license is in evaluation mode.
expiry_time	string	Date and time when the license expires.
owner	string	Cluster, node or license manager that owns the license.

Name	Type	Description
serial_number	string	Serial number of the license.
start_time	string	Date and time when the license starts.

license_package

Name	Type	Description
_links	_links	
keys	array[string]	
licenses	array[licenses]	Installed licenses of the package.
name	string	Name of the license.
scope	string	Scope of the license.
state	string	Summary state of package based on all installed licenses.

_links

Name	Type	Description
next	href	
self	href	

records

Name	Type	Description
_links	_links	
keys	array[string]	
licenses	array[licenses]	Installed licenses of the package.
name	string	Name of the license.
scope	string	Scope of the license.
state	string	Summary state of package based on all installed licenses.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage a cluster license package

Cluster licensing licenses name endpoint overview

Overview

Manages a specific instance of a license package.

Examples

Retrieving information for a specific license package

This example shows how to retrieve information about the specific feature package `fabricpool`.

```
# API
GET /cluster/licensing/licenses/fabricpool/

# Response
200 OK

# JSON Body
{
  "name": "fabricpool",
  "scope": "cluster",
  "state": "compliant",
  "licenses": [
    {
      "owner": "testcluster-1",
      "serial_number": "123456789",
      "state": "compliant",
      "capacity": {
        "maximum_size": 109951162777600,
        "used_size": 0
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/cluster/licensing/licenses/fabricpool/"
    }
  }
}
```

Deleting a specific license

This example show how to delete a CIFS site license.

```
# API
DELETE /cluster/licensing/licenses/cifs/?serial_number=1-80-000011"

# JSON Body
{}

# Response
200 OK
```

Deleting with a query

The following example shows how to delete all NFS licenses specified with the '*' query.


```
# API
DELETE /cluster/licensing/licenses/nfs/?serial_number=*

# JSON Body
{}

# Response
200 OK
```

Delete a license

DELETE /cluster/licensing/licenses/{name}

Deletes a license.

Related ONTAP commands

- `system license delete`

Learn more

- [DOC /cluster/licensing/licenses/{name}](#)

Parameters

Name	Type	In	Required	Description
name	string	path	True	
serial_number	string	query	True	

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
1115144	Cloud licenses cannot be deleted
1115178	A tier license, that is still in use, cannot be deleted
1115406	Capacity pool licenses cannot be deleted

Error Code	Description
66846823	A flexcache licenses, that is still in use, cannot be deleted
1115213	License is still in use and cannot be removed
1115137	Cluster license requires a base license to be installed
525028	Error during volume limit check, cannot remove license
525029	Current volume use will exceed limits if license is removed

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a license package

GET `/cluster/licensing/licenses/{name}`

Retrieves a specific license package.

Related ONTAP commands

- `system license show`
- `system license show-status`

Learn more

- [DOC /cluster/licensing/licenses/{name}](#)

Parameters

Name	Type	In	Required	Description
name	string	path	True	Name of the license package.
scope	string	query	False	Filter by scope

Name	Type	In	Required	Description
licenses.expiry_time	string	query	False	Filter by licenses.expiry_time
licenses.compliance.state	string	query	False	Filter by licenses.compliance.state
licenses.capacity.used_size	integer	query	False	Filter by licenses.capacity.used_size
licenses.capacity.maximum_size	integer	query	False	Filter by licenses.capacity.maximum_size
licenses.serial_number	string	query	False	Filter by licenses.serial_number
licenses.active	boolean	query	False	Filter by licenses.active
licenses.owner	string	query	False	Filter by licenses.owner
licenses.evaluation	boolean	query	False	Filter by licenses.evaluation
licenses.start_time	string	query	False	Filter by licenses.start_time
state	string	query	False	Filter by state
name	string	query	False	Filter by name
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	

Name	Type	Description
keys	array[string]	
licenses	array[licenses]	Installed licenses of the package.
name	string	Name of the license.
scope	string	Scope of the license.
state	string	Summary state of package based on all installed licenses.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "keys": {
  },
  "licenses": {
    "capacity": {
      "maximum_size": 0,
      "used_size": 0
    },
    "compliance": {
      "state": "compliant"
    },
    "expiry_time": "2019-03-02 19:00:00 UTC",
    "owner": "cluster1",
    "serial_number": "123456789",
    "start_time": "2019-02-02 19:00:00 UTC"
  },
  "name": "NFS",
  "scope": "not_available",
  "state": "compliant"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

capacity

Name	Type	Description
maximum_size	integer	Licensed capacity size (in bytes) that can be used.
used_size	integer	Capacity that is currently used (in bytes).

compliance

Name	Type	Description
state	string	Compliance state of the license.

licenses

Name	Type	Description
active	boolean	A flag indicating whether the license is currently being enforced.
capacity	capacity	
compliance	compliance	
evaluation	boolean	A flag indicating whether the license is in evaluation mode.
expiry_time	string	Date and time when the license expires.
owner	string	Cluster, node or license manager that owns the license.

Name	Type	Description
serial_number	string	Serial number of the license.
start_time	string	Date and time when the license starts.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve historical performance metrics for the cluster

GET `/cluster/metrics`

Retrieves historical performance metrics for the cluster.

Parameters

Name	Type	In	Required	Description
status	string	query	False	Filter by status
timestamp	string	query	False	Filter by timestamp
duration	string	query	False	Filter by duration
iops.total	integer	query	False	Filter by iops.total

Name	Type	In	Required	Description
iops.other	integer	query	False	Filter by iops.other
iops.read	integer	query	False	Filter by iops.read
iops.write	integer	query	False	Filter by iops.write
latency.total	integer	query	False	Filter by latency.total
latency.other	integer	query	False	Filter by latency.other
latency.read	integer	query	False	Filter by latency.read
latency.write	integer	query	False	Filter by latency.write
throughput.total	integer	query	False	Filter by throughput.total
throughput.other	integer	query	False	Filter by throughput.other
throughput.read	integer	query	False	Filter by throughput.read
throughput.write	integer	query	False	Filter by throughput.write
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1
fields	array[string]	query	False	Specify the fields to return.

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	return_records	boolean	query	False
The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1	interval	string	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "latency": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "status": "ok",
    "throughput": {
      "read": 200,
      "total": 1000,
      "write": 100
    },
    "timestamp": "2017-01-25 11:20:13 UTC"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

records

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes does not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage cluster nodes

Cluster nodes endpoint overview

Overview

This API is used to add nodes to a cluster, update node-specific configurations, and retrieve the current node configuration details.

Adding a node to a cluster

A node can be added to a cluster by issuing a POST `/cluster/nodes` request to a node currently in the cluster. All nodes must be at the same version to use this API. Mixed version joins are not supported in this release. Properties can be provided as fields in the body of the POST request to configure node-specific settings. On a successful request, POST `/cluster/nodes` returns a status code of 202 and job information in the body. The `/cluster/jobs` APIs can be used to track the status of the node add job.

Fields used for adding a node

Fields used for the `/cluster/nodes` APIs fall into the following categories

Required node fields

The following field is required for any POST `/cluster/nodes` request:

- `cluster_interface.ip.address`

Optional fields

All of the following fields are used to setup additional cluster-wide configuration:

- name
- location
- records

Network interface fields

Each node can have a node-specific configuration set in POST /cluster/nodes. If a field is provided in the body of a node, it must be provided for all nodes in the POST body. The node management interface can be provided for each node if all node management interfaces in the cluster use the same netmask. If the node management interfaces use different netmasks, then configuration of the node management interfaces should be done using the /network/ip/interfaces API.

The records field

Multiple nodes can be added to the cluster in one request by providing an array named "records" with multiple node entries. Each node entry in records must follow the required and optional fields listed previously. When only adding a single node, no records field is needed. See 'Example usecases' for an example of how to use the records field.

Modifying node configurations

The following fields can be used to modify a node configuration:

- name
- location

Examples

The following examples show how to shutdown/reboot a node and how to update a node configuration.

Adding a single node with a minimal configuration

```
# Body
body =
{
  "cluster_interface": {
    "ip": {
      "address": "1.1.1.1"
    }
  }
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster/nodes" -d body
```

Adding multiple nodes in the same request

```
# Body
body =
{
  "records": [
    {
      "name": "node1",
      "cluster_interface": {
        "ip": {
          "address": "1.1.1.1"
        }
      }
    },
    {
      "name": "node2",
      "cluster_interface": {
        "ip": {
          "address": "2.2.2.2"
        }
      }
    }
  ],
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster/nodes" -d body
```

Modifying a cluster-wide configuration

```
# Body
body =
{
  "name": "renamedNode",
  "location": "newLocation"
}

# Request
curl -X PATCH "https://<mgmt-ip>/api/cluster/nodes" -d body
```

Shutting down a node

```
curl -X PATCH "https://<mgmt-ip>/api/cluster/nodes/{uuid}?action=shutdown"
```

Retrieve nodes in a cluster

GET /cluster/nodes

Retrieves the nodes in the cluster.

Learn more

- [DOC /cluster/nodes](#)

Parameters

Name	Type	In	Required	Description
ha.auto_giveback	boolean	query	False	Filter by ha.auto_giveback
ha.enabled	boolean	query	False	Filter by ha.enabled
ha.partners.name	string	query	False	Filter by ha.partners.name
ha.partners.uuid	string	query	False	Filter by ha.partners.uuid
uptime	integer	query	False	Filter by uptime
date	string	query	False	Filter by date
membership	string	query	False	Filter by membership
serial_number	string	query	False	Filter by serial_number
controller.over_temperature	string	query	False	Filter by controller.over_temperature

Name	Type	In	Required	Description
controller.flash_cache.serial_number	string	query	False	Filter by controller.flash_cache.serial_number
controller.flash_cache.hardware_revision	string	query	False	Filter by controller.flash_cache.hardware_revision
controller.flash_cache.capacity	integer	query	False	Filter by controller.flash_cache.capacity
controller.flash_cache.model	string	query	False	Filter by controller.flash_cache.model
controller.flash_cache.slot	string	query	False	Filter by controller.flash_cache.slot
controller.flash_cache.state	string	query	False	Filter by controller.flash_cache.state
controller.flash_cache.firmware_version	string	query	False	Filter by controller.flash_cache.firmware_version
controller.flash_cache.part_number	string	query	False	Filter by controller.flash_cache.part_number
controller.frus.id	integer	query	False	Filter by controller.frus.id
controller.frus.type	string	query	False	Filter by controller.frus.type
controller.frus.state	string	query	False	Filter by controller.frus.state
location	string	query	False	Filter by location
model	string	query	False	Filter by model

Name	Type	In	Required	Description
management_interfaces.ip.address	string	query	False	Filter by management_interfaces.ip.address
management_interfaces.uuid	string	query	False	Filter by management_interfaces.uuid
management_interfaces.name	string	query	False	Filter by management_interfaces.name
service_processor.link_status	string	query	False	Filter by service_processor.link_status
service_processor.state	string	query	False	Filter by service_processor.state
service_processor.firmware_version	string	query	False	Filter by service_processor.firmware_version
service_processor.dhcp_enabled	boolean	query	False	Filter by service_processor.dhcp_enabled
service_processor.ipv4_interface.address	string	query	False	Filter by service_processor.ipv4_interface.address
service_processor.ipv4_interface.netmask	string	query	False	Filter by service_processor.ipv4_interface.netmask
service_processor.ipv4_interface.gateway	string	query	False	Filter by service_processor.ipv4_interface.gateway
service_processor.mac_address	string	query	False	Filter by service_processor.mac_address

Name	Type	In	Required	Description
service_processor.ip v6_interface.address	string	query	False	Filter by service_processor.ip v6_interface.address
service_processor.ip v6_interface.netmask	string	query	False	Filter by service_processor.ip v6_interface.netmask
service_processor.ip v6_interface.gateway	string	query	False	Filter by service_processor.ip v6_interface.gateway
name	string	query	False	Filter by name
version.minor	integer	query	False	Filter by version.minor
version.full	string	query	False	Filter by version.full
version.major	integer	query	False	Filter by version.major
version.generation	integer	query	False	Filter by version.generation
cluster_interfaces.ip. address	string	query	False	Filter by cluster_interfaces.ip. address
cluster_interfaces.uuid	string	query	False	Filter by cluster_interfaces.uuid
cluster_interfaces.name	string	query	False	Filter by cluster_interfaces.name
uuid	string	query	False	Filter by uuid
fields	array[string]	query	False	Specify the fields to return.

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "cluster_interface": {
      "ip": {
        "address": "10.10.10.7"
      }
    },
    "cluster_interfaces": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ip": {
        "address": "10.10.10.7"
      },
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "controller": {
      "flash_cache": {
        "capacity": 1024000000000,
        "firmware_version": "NA05",
        "hardware_revision": "A1",
        "model": "X1970A",
        "part_number": "119-00207",
        "serial_number": "A22P5061550000187",
        "slot": "6-1",
        "state": "ok"
      },
      "frus": {
```



```

        "id": 0,
        "state": "ok",
        "type": "fan"
    },
    "over_temperature": "over"
},
"date": "2017-01-25 11:20:13 +0400",
"ha": {
    "partners": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
},
"location": "rack 2 row 5",
"management_interface": {
    "ip": {
        "address": "10.10.10.7"
    }
},
"management_interfaces": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "ip": {
        "address": "10.10.10.7"
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"membership": "available",
"model": "FAS3070",
"name": "node-01",
"serial_number": "4048820-60-9",
"service_processor": {
    "firmware_version": "string",
    "ipv4_interface": {
        "address": "10.10.10.7",
        "gateway": "10.1.1.1",
        "netmask": "24"
    }
}

```

```

    },
    "ipv6_interface": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    },
    "link_status": "up",
    "mac_address": "string",
    "state": "online"
  },
  "uptime": 300536,
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
  "version": {
    "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
    "generation": 9,
    "major": 4,
    "minor": 0
  }
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```


See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

node_setup_ip

The IP configuration for cluster setup.

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interface

The cluster network IP address of the node to be added.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

flash_cache

Name	Type	Description
capacity	integer	Size in bytes
firmware_version	string	
hardware_revision	string	
model	string	
part_number	string	
serial_number	string	
slot	string	
state	string	

frus

Name	Type	Description
id	integer	
state	string	
type	string	

controller

Controller information

Name	Type	Description
flash_cache	array[flash_cache]	A list of Flash-Cache devices. Only returned when requested by name.
frus	array[frus]	A list of frus in the node. Only returned when requested by name.

Name	Type	Description
over_temperature	string	Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds.

partners

Name	Type	Description
_links	_links	
name	string	
uuid	string	

ha

Name	Type	Description
auto_giveback	boolean	Specifies whether giveback is automatically initiated when the node that owns the storage is ready.
enabled	boolean	Specifies whether or not storage failover is enabled.
partners	array[partners]	The nodes in this node's High Availability (HA) group.

management_interface

The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

management_interfaces

Network interface

Name	Type	Description
_links	_links	

Name	Type	Description
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

ipv4_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

ipv6_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

service_processor

Name	Type	Description
dhcp_enabled	boolean	Set to true to use DHCP to configure an IPv4 interface.

Name	Type	Description
firmware_version	string	The version of firmware installed.
ipv4_interface	ipv4_interface	Object to setup an interface along with its default router.
ipv6_interface	ipv6_interface	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	
state	string	

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

records

Complete node information

Name	Type	Description
_links	_links	
cluster_interface	cluster_interface	The cluster network IP address of the node to be added.
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information
date	string	Specifies the ISO-8601 format date and time on the node.
ha	ha	

Name	Type	Description
location	string	
management_interface	management_interface	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - If a node is available, this means it is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
model	string	
name	string	
serial_number	string	
service_processor	service_processor	

Name	Type	Description
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Add a node or nodes to a cluster

POST `/cluster/nodes`

Adds a node or nodes to the cluster

Required properties

- `cluster_interface.ip.address`

Learn more

- [DOC /cluster/nodes](#)

Request Body

Name	Type	Description
_links	_links	
cluster_interface	cluster_interface	The cluster network IP address of the node to be added.
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information
date	string	Specifies the ISO-8601 format date and time on the node.
ha	ha	
location	string	
management_interface	management_interface	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - If a node is available, this means it is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
model	string	
name	string	
serial_number	string	
service_processor	service_processor	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "cluster_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "controller": {
    "flash_cache": {
      "capacity": 1024000000000,
      "firmware_version": "NA05",
      "hardware_revision": "A1",
      "model": "X1970A",
      "part_number": "119-00207",
      "serial_number": "A22P5061550000187",
      "slot": "6-1",
      "state": "ok"
    },
    "frus": {
      "id": 0,
      "state": "ok",
      "type": "fan"
    },
    "over_temperature": "over"
  },
  "date": "2017-01-25 11:20:13 +0400",
  "ha": {
    "partners": {
```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"location": "rack 2 row 5",
"management_interface": {
  "ip": {
    "address": "10.10.10.7"
  }
},
"management_interfaces": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ip": {
    "address": "10.10.10.7"
  },
  "name": "lif1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"membership": "available",
"model": "FAS3070",
"name": "node-01",
"serial_number": "4048820-60-9",
"service_processor": {
  "firmware_version": "string",
  "ipv4_interface": {
    "address": "10.10.10.7",
    "gateway": "10.1.1.1",
    "netmask": "24"
  },
  "ipv6_interface": {
    "address": "10.10.10.7",
    "gateway": "10.1.1.1",
    "netmask": "24"
  },
  "link_status": "up",
  "mac_address": "string",
  "state": "online"
}

```

```

    },
    "uptime": 300536,
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
    "version": {
      "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
      "generation": 9,
      "major": 4,
      "minor": 0
    }
  }
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```

{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
131727360	A node cannot be added to the cluster. This is a generic code, see response message for details.

Error Code	Description
262245	The value provided was invalid.
1179817	The IP address, netmask, and gateway must all be provided for cluster management interface.
1179813	Fields set for one node must be set for all nodes.
1179818	The IP address and gateway must be of the same family.
1179821	An IP address and netmask conflicts with an existing entry.
1179795	A node being added is already in the cluster.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node_setup_ip

The IP configuration for cluster setup.

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interface

The cluster network IP address of the node to be added.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.

Name	Type	Description
uuid	string	The UUID that uniquely identifies the interface.

flash_cache

Name	Type	Description
capacity	integer	Size in bytes
firmware_version	string	
hardware_revision	string	
model	string	
part_number	string	
serial_number	string	
slot	string	
state	string	

frus

Name	Type	Description
id	integer	
state	string	
type	string	

controller

Controller information

Name	Type	Description
flash_cache	array[flash_cache]	A list of Flash-Cache devices. Only returned when requested by name.
frus	array[frus]	A list of frus in the node. Only returned when requested by name.
over_temperature	string	Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds.

partners

Name	Type	Description
_links	_links	
name	string	
uuid	string	

ha

Name	Type	Description
auto_giveback	boolean	Specifies whether giveback is automatically initiated when the node that owns the storage is ready.
enabled	boolean	Specifies whether or not storage failover is enabled.
partners	array[partners]	The nodes in this node's High Availability (HA) group.

management_interface

The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

management_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

ipv4_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

ipv6_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

service_processor

Name	Type	Description
dhcp_enabled	boolean	Set to true to use DHCP to configure an IPv4 interface.
firmware_version	string	The version of firmware installed.
ipv4_interface	ipv4_interface	Object to setup an interface along with its default router.
ipv6_interface	ipv6_interface	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	

Name	Type	Description
state	string	

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

node

Complete node information

Name	Type	Description
_links	_links	
cluster_interface	cluster_interface	The cluster network IP address of the node to be added.
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information
date	string	Specifies the ISO-8601 format date and time on the node.
ha	ha	
location	string	
management_interface	management_interface	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - If a node is available, this means it is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
model	string	
name	string	
serial_number	string	
service_processor	service_processor	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve node information

GET /cluster/nodes/{uuid}

Retrieves information for the node.

Learn more

- [DOC /cluster/nodes](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	• format: uuid

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
cluster_interface	cluster_interface	The cluster network IP address of the node to be added.
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information
date	string	Specifies the ISO-8601 format date and time on the node.
ha	ha	
location	string	
management_interface	management_interface	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - If a node is available, this means it is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
model	string	
name	string	
serial_number	string	
service_processor	service_processor	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "cluster_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "controller": {
    "flash_cache": {
      "capacity": 1024000000000,
      "firmware_version": "NA05",
      "hardware_revision": "A1",
      "model": "X1970A",
      "part_number": "119-00207",
      "serial_number": "A22P5061550000187",
      "slot": "6-1",
      "state": "ok"
    },
    "frus": {
      "id": 0,
      "state": "ok",
      "type": "fan"
    },
    "over_temperature": "over"
  },
  "date": "2017-01-25 11:20:13 +0400",
  "ha": {
    "partners": {
```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"location": "rack 2 row 5",
"management_interface": {
  "ip": {
    "address": "10.10.10.7"
  }
},
"management_interfaces": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ip": {
    "address": "10.10.10.7"
  },
  "name": "lif1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"membership": "available",
"model": "FAS3070",
"name": "node-01",
"serial_number": "4048820-60-9",
"service_processor": {
  "firmware_version": "string",
  "ipv4_interface": {
    "address": "10.10.10.7",
    "gateway": "10.1.1.1",
    "netmask": "24"
  },
  "ipv6_interface": {
    "address": "10.10.10.7",
    "gateway": "10.1.1.1",
    "netmask": "24"
  },
  "link_status": "up",
  "mac_address": "string",
  "state": "online"
}

```

```

    },
    "uptime": 300536,
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
    "version": {
      "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
      "generation": 9,
      "major": 4,
      "minor": 0
    }
  }
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node_setup_ip

The IP configuration for cluster setup.

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interface

The cluster network IP address of the node to be added.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.

Name	Type	Description
uuid	string	The UUID that uniquely identifies the interface.

flash_cache

Name	Type	Description
capacity	integer	Size in bytes
firmware_version	string	
hardware_revision	string	
model	string	
part_number	string	
serial_number	string	
slot	string	
state	string	

frus

Name	Type	Description
id	integer	
state	string	
type	string	

controller

Controller information

Name	Type	Description
flash_cache	array[flash_cache]	A list of Flash-Cache devices. Only returned when requested by name.
frus	array[frus]	A list of frus in the node. Only returned when requested by name.
over_temperature	string	Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds.

partners

Name	Type	Description
_links	_links	
name	string	
uuid	string	

ha

Name	Type	Description
auto_giveback	boolean	Specifies whether giveback is automatically initiated when the node that owns the storage is ready.
enabled	boolean	Specifies whether or not storage failover is enabled.
partners	array[partners]	The nodes in this node's High Availability (HA) group.

management_interface

The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

management_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

ipv4_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

ipv6_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

service_processor

Name	Type	Description
dhcp_enabled	boolean	Set to true to use DHCP to configure an IPv4 interface.
firmware_version	string	The version of firmware installed.
ipv4_interface	ipv4_interface	Object to setup an interface along with its default router.
ipv6_interface	ipv6_interface	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	

Name	Type	Description
state	string	

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update node information

PATCH /cluster/nodes/{uuid}

Updates the node information or performs shutdown/reboot actions on a node.

Learn more

- [DOC /cluster/nodes](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	<ul style="list-style-type: none">• format: uuid
action	string	query	False	<p>The shutdown action shuts the node down and transfers storage control to its HA group if storage failover is enabled. The reboot action reboots the node and transfers storage control to its HA group if storage failover is enabled.</p> <ul style="list-style-type: none">• enum: ["shutdown", "reboot"]
shutdown_reboot_reason	string	query	False	<p>Indicates the reason for the reboot or shutdown. This only applies when an action of reboot or shutdown is provided.</p>
allow_data_outage	boolean	query	False	<p>This only applies when an action of reboot or shutdown is provided. It allows storage failover to be bypassed along with any failures related to maintaining quorum in the cluster.</p> <ul style="list-style-type: none">• Default value:

Request Body

Name	Type	Description
_links	_links	
cluster_interface	cluster_interface	The cluster network IP address of the node to be added.
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information
date	string	Specifies the ISO-8601 format date and time on the node.
ha	ha	
location	string	
management_interface	management_interface	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - If a node is available, this means it is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
model	string	
name	string	
serial_number	string	
service_processor	service_processor	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "cluster_interfaces": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "controller": {
    "flash_cache": {
      "capacity": 1024000000000,
      "firmware_version": "NA05",
      "hardware_revision": "A1",
      "model": "X1970A",
      "part_number": "119-00207",
      "serial_number": "A22P5061550000187",
      "slot": "6-1",
      "state": "ok"
    },
    "frus": {
      "id": 0,
      "state": "ok",
      "type": "fan"
    },
    "over_temperature": "over"
  },
  "date": "2017-01-25 11:20:13 +0400",
  "ha": {
    "partners": {
```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"location": "rack 2 row 5",
"management_interface": {
  "ip": {
    "address": "10.10.10.7"
  }
},
"management_interfaces": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ip": {
    "address": "10.10.10.7"
  },
  "name": "lif1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"membership": "available",
"model": "FAS3070",
"name": "node-01",
"serial_number": "4048820-60-9",
"service_processor": {
  "firmware_version": "string",
  "ipv4_interface": {
    "address": "10.10.10.7",
    "gateway": "10.1.1.1",
    "netmask": "24"
  },
  "ipv6_interface": {
    "address": "10.10.10.7",
    "gateway": "10.1.1.1",
    "netmask": "24"
  },
  "link_status": "up",
  "mac_address": "string",
  "state": "online"
}

```

```

    },
    "uptime": 300536,
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
    "version": {
      "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
      "generation": 9,
      "major": 4,
      "minor": 0
    }
  }
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```

{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
852046	HA partner node
65562	Internal RPC error

Error Code	Description
852115	The reboot/shutdown is prevented because LIFs cannot be moved away from the node
9240606	The reboot/shutdown is prevented due to quorum warnings.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node_setup_ip

The IP configuration for cluster setup.

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interface

The cluster network IP address of the node to be added.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.

Name	Type	Description
uuid	string	The UUID that uniquely identifies the interface.

flash_cache

Name	Type	Description
capacity	integer	Size in bytes
firmware_version	string	
hardware_revision	string	
model	string	
part_number	string	
serial_number	string	
slot	string	
state	string	

frus

Name	Type	Description
id	integer	
state	string	
type	string	

controller

Controller information

Name	Type	Description
flash_cache	array[flash_cache]	A list of Flash-Cache devices. Only returned when requested by name.
frus	array[frus]	A list of frus in the node. Only returned when requested by name.
over_temperature	string	Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds.

partners

Name	Type	Description
_links	_links	
name	string	
uuid	string	

ha

Name	Type	Description
auto_giveback	boolean	Specifies whether giveback is automatically initiated when the node that owns the storage is ready.
enabled	boolean	Specifies whether or not storage failover is enabled.
partners	array[partners]	The nodes in this node's High Availability (HA) group.

management_interface

The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

management_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

ipv4_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

ipv6_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

service_processor

Name	Type	Description
dhcp_enabled	boolean	Set to true to use DHCP to configure an IPv4 interface.
firmware_version	string	The version of firmware installed.
ipv4_interface	ipv4_interface	Object to setup an interface along with its default router.
ipv6_interface	ipv6_interface	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	

Name	Type	Description
state	string	

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

node

Complete node information

Name	Type	Description
_links	_links	
cluster_interface	cluster_interface	The cluster network IP address of the node to be added.
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information
date	string	Specifies the ISO-8601 format date and time on the node.
ha	ha	
location	string	
management_interface	management_interface	The management interface of the node to be added. The netmask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - If a node is available, this means it is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. A query on the "membership" property for <i>available</i> must be provided to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node may be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
model	string	
name	string	
serial_number	string	
service_processor	service_processor	
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage cluster peers

Cluster peers endpoint overview

Cluster peer operations

Cluster peering allows administrators of ONTAP systems to establish relationships between two or more independent clusters. Once a relationship exists between two clusters, they may then exchange user data, configuration information and coordinate operations. The `/cluster/peers` endpoint supports operations such as create, get, modify and delete using GET, PATCH and POST HTTP requests.

Creating a cluster peer

A new cluster peer relationship can be set up by issuing a POST request to `/cluster/peers`. Parameters in the POST body define the settings of the peering relationship. A successful POST request that succeeds in creating a peer returns a HTTP status code, code 201, along with the details of the created peer such as peer UUID, name, authentication information. A failed POST request returns an HTTP error code along with a message indicating the reason for the error. This can include malformed request and invalid operations.

Sample request

```
curl -X POST 'https://<mgmt-ip>/api/cluster/peers/' -d  
'{"authentication":{"expiry_time":"12/25/2018  
12:34:56","generate_passphrase":true}}'
```

Examples


```

# Create - no params
body = {}

# Create with a peer address and a passphrase
body =
{
  "remote":
    {
      "ip_addresses":["1.2.3.4"]
    }
}

# Create with a peer name and a generated passphrase that is true
body =
{
  "name":"cp_xyz123",
  "authentication":
    {
      "generate_passphrase":true
    }
}

# Create with a name, a peer address, and a passphrase
body =
{
  "name":"cp_xyz123",
  "remote":
    {
      "ip_addresses": ["1.2.3.4"]
    },
  "authentication":
    {
      "passphrase":"xyz12345"
    }
}

# Create with a proposed encryption protocol
body =
{
  "encryption":
    {
      "proposed":"tls-psk"
    }
}

```

Creating local intercluster LIFs

The local cluster must have an intercluster LIF on each node for the correct operation of cluster peering. If no local intercluster LIFs exist, you can optionally specify LIFs to be created for each node in the local cluster. These local interfaces, if specified, are created on each node before proceeding with the creation of the cluster peering relationship. Cluster peering relationship would be established if there is an error preventing the LIFs from being created. Local interfaces, once created, should not be specified for subsequent cluster peering relationships.

Local LIF creation fields

- `local_network.ip_addresses` - list of IP addresses to assign, one per node in the local cluster
- `local_network.netmask` - IPv4 mask or netmask length
- `local_network.broadcast_domain` - Broadcast domain that is in use within the IPspace.
- `local_network.gateway` - The IPv4 or IPv6 address of the default router.

Additional information on network routes

It might happen that when creating LIFs the network route discovery mechanism could take additional time (1-5 seconds) to become visible in the network outside of the cluster. This delay in publishing the routes might cause an initial cluster peer "create" request to fail. This error disappears with a retry of the same request.

Example

```
curl -X POST "https://<mgmt-ip>/api/cluster/peers" -d body
```

where "<mgmt-ip>" is replaced by the IP address of the cluster management LIF, and "body" is replaced by the JSON body of the POST, containing the fields for the new peering relationship and local LIFs.</mgmt-ip>

Example POST body

To create 4 intercluster LIFs on a 4-node cluster before creating a cluster peer relationship:

```
{
  "local_network":
  {
    "interfaces": [
      {"ip_address": "1.2.3.4"},
      {"ip_address": "1.2.3.5"},
      {"ip_address": "1.2.3.6"}
    ],
    "netmask": "255.255.0.0",
    "broadcast_domain": "Default",
    "gateway": "1.2.0.1"
  }
  "remote.ip_addresses": ["1.2.9.9"],
  "authentication.passphrase": "xyz12345"
}
```

Retrieve a cluster peer

Peers in a cluster can be retrieved by issuing a GET request to `/cluster/peers`. It is also possible to retrieve a specific peer when qualified by its UUID to `/cluster/peers/{uuid}`.

Overview of fields used for retrieving a cluster peer

A GET request might have no query parameters or a valid cluster UUID. The former retrieves all records while the latter retrieves the record for the cluster peer with that UUID.

Required fields

There are no required fields for GET requests.

Optional fields

The following fields are optional for GET requests

- `UUID` - uuid of the cluster peer

Examples

```
curl -X GET "https://<mgmt-ip>/api/cluster/peers/"
curl -X GET "https://<mgmt-ip>/api/cluster/peers/{uuid}"
curl -X GET "https://<mgmt-ip>/api/cluster/peers/{uuid}?fields=*"

```

Update a cluster peer

A cluster peer relationship can be updated by issuing a PATCH request to `/cluster/peers/{uuid}`. As in the CLI mode, you can toggle the proposed encryption protocol, update the passphrase, or specify a new set of stable addresses. All PATCH requests take the parameters that are to be updated in the request body. If the `generate_passphrase` is 'true', the passphrase is returned in the PATCH response.

Fields overview

This sections highlights the parameters that control the modification of an existing cluster peering relationship.

Required fields

A PATCH request with an empty body has no effect on the cluster peer instance. All other fields and the combinations in which they are valid are indicated below:

- `encryption_proposed` - Toggle the proposed encryption protocol (from 'none' to 'tls-psk' or otherwise). Authentication must be true and a passphrase must be present in body.
- `passphrase`
- `passphrase` or `generate passphrase`
- `remote.ip_addresses`

Optional fields

- `expiration time` - Set the expiration time of the passphrase

Examples

```

# Update with an empty body
body = {}

# Update the proposed encryption protocol from tls-psk to none
body =
{
  "authentication":
    {
      "passphrase": "xyz12345",
      "in_use": "ok"
    },
  "encryption":
    {
      "proposed": "none"
    }
}

# Update the passphrase
body =
{
  "authentication":
    {
      "passphrase": "xyz12345",
      "in_use": "ok"
    }
}

# Set an auto-generated passphrase
body =
{
  "authentication":
    {
      "generate_passphrase": true,
      "in_use": "ok"
    }
}

# Update remote IP addresses
body =
{
  "remote":
    {
      "ip_addresses": ["10.224.65.30"]
    }
}

```

Sample requests

```
# Set a passphrase
curl -X PATCH 'https://<mgmt-ip>/api/cluster/peers/73123071-d0b9-11e8-a686-005056a7179a' -d
'{"authentication":{"passphrase":"xyz12345","in_use":"ok"}}'

# Update a peer address
curl -X PATCH 'https://<mgmt-ip>/api/cluster/peers/73123071-d0b9-11e8-a686-005056a7179a' -d '{"remote":{"ip_addresses":["1.2.3.4"]}}'
```

Delete a cluster peer

This interface allows you to delete a cluster peer using the HTTP DELETE request.

Required fields

All delete operations must be performed on a valid peer UUID. Deleting an invalid peer returns 'HTTP 404' indicating an error.

Optional fields

The DELETE operation has no optional fields.

Request format

DELETE "https://<mgmt-ip>/api/cluster/peers/{uuid}"</mgmt-ip>

Examples

The request -

```
curl -X DELETE "https://<mgmt-ip>/api/cluster/peers/8becc0d4-c12c-11e8-9ceb-005056bbd143"
```

deletes a peer with peer UUID '8becc0d4-c12c-11e8-9ceb-005056bbd143'

Retrieve cluster peers

GET /cluster/peers

Retrieve the collection of cluster peers.

Learn more

- [DOC /cluster/peers](#)

Parameters

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[cluster_peer]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "interfaces": {
        "href": "/api/resourcelink"
      },
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "authentication": {
      "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
      "in_use": "ok",
      "state": "ok"
    },
    "encryption": {
      "proposed": "none",
      "state": "none"
    },
    "initial_allowed_svms": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "ipspace": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "exchange",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  }
}
```



```

    },
    "local_network": {
      "broadcast_domain": "bd1",
      "gateway": "10.1.1.1",
      "interfaces": {
        "ip_address": "10.10.10.7"
      },
      "netmask": "255.255.0.0"
    },
    "name": "cluster2",
    "remote": {
      "ip_addresses": {
      },
      "name": "cluster2",
      "serial_number": "4048820-60-9"
    },
    "status": {
      "state": "available",
      "update_time": "2017-01-25 11:20:13 UTC"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
interfaces	href	
self	href	

authentication

Name	Type	Description
expiry_time	string	The time when the passphrase will expire, in ISO 8601 duration format or date and time format. The default is 1 hour.
generate_passphrase	boolean	Auto generate a passphrase when true.
in_use	string	
passphrase	string	A password to authenticate the cluster peer relationship.
state	string	

encryption

Name	Type	Description
proposed	string	
state	string	

_links

Name	Type	Description
self	href	

initial_allowed_svms

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ipspace

The IPspace of the local intercluster LIFs

Name	Type	Description
_links	_links	
name	string	IPspace name
uuid	string	IPspace UUID

interfaces

Name	Type	Description
ip_address	string	IPv4 or IPv6 address

local_network

Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.

Name	Type	Description
broadcast_domain	string	Broadcast domain that is in use within the IPspace.
gateway	string	The IPv4 or IPv6 address of the default router.
interfaces	array[interfaces]	
netmask	string	IPv4 mask or netmask length.

remote

Name	Type	Description
ip_addresses	array[string]	The IPv4 addresses, IPv6 addresses, or hostnames of the peers.
name	string	The name of the remote cluster.
serial_number	string	The serial number of the remote cluster.

status

Name	Type	Description
state	string	
update_time	string	The last time the state was updated.

cluster_peer

Name	Type	Description
_links	_links	
authentication	authentication	
encryption	encryption	
initial_allowed_svms	array[initial_allowed_svms]	The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship.
ipspace	ipspace	The IPspace of the local intercluster LIFs
local_network	local_network	Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.
name	string	Optional name for the cluster peer relationship. By default it is the name of the remote cluster.

Name	Type	Description
remote	remote	
status	status	
uuid	string	UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a peering relationship

POST `/cluster/peers`

Creates a peering relationship and, optionally, the IP interfaces it will use. There are two ways to create a peering relationship.

Provide remote IP

Here the user provides the remote IP address. Creating a new cluster peer relationship with a specific remote cluster requires at least one remote intercluster IP address from that cluster.

Required properties

- `remote.ip_addresses` - Addresses of the remote peers. The local peer must be able to reach and connect to these addresses for the request to succeed in creating a peer.

- Either set `generate_passphrase` to true or provide a passphrase in the body of the request; only one of them is required.

Optional properties

The following fields are optional for a POST on `/cluster/peer/`. All fields must follow the structure in the cluster peer API definition.

- `name` - Name of the peering relationship.
- `passphrase` - User generated passphrase for use in authentication.
- `generate_passphrase` (true/false) - When this option is true, ONTAP automatically generates a passphrase to authenticate cluster peers.
- `ipspace` - IPspace of the local intercluster LIFs. Assumes Default IPspace if not provided.
- `initial_allowed_svms` - the local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship.
- `local_network` - fields to create a local intercluster LIF. See section on "Creating local intercluster lifs".
- `expiry_time` - Duration in ISO 8601 format for which the user-supplied or auto-generated passphrase is valid. Expiration time must not be greater than seven days into the future. ISO 8601 duration format is "PnDTnHnMnS" or "PnW" where n is a positive integer. The nD, nH, nM and nS fields can be dropped if zero. "P" should always be present and "T" should be present if there are any hours, minutes or seconds fields.
- `encryption_proposed` (none/tls-psk) - Encryption mechanism of the communication channel between the two peers.

Do not provide remote IP

This method is used when the remote IP address is not provided. This method is used when the filer is ready to accept peering request from foreign clusters.

Required properties

- `generate_passphrase` (true) - This option must be set to true. ONTAP automatically generates a passphrase to authenticate cluster peers. Either set `generate_passphrase` to true or provide a passphrase in the body of the request; only one of them is required.

Optional properties

The following fields are optional for a POST on `/cluster/peer/`. All fields must follow the structure in the cluster peer API definition.

- `name` - Name of the remote peer.
- `ipspace` - IPspace of the local intercluster LIFs. Assumes Default IPspace if not provided.
- `initial_allowed_svms` - Local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship.
- `local_network` - Fields to create a local intercluster LIF. See section on "Creating local intercluster lifs".
- `expiry_time` - Duration in ISO 8601 format for which the user-supplied or auto-generated passphrase is valid. Expiration time must not be greater than seven days into the future. ISO 8601 duration format is "PnDTnHnMnS" or "PnW" where n is a positive integer. The nD, nH, nM and nS fields can be dropped if

zero. "P" should always be present and "T" should be present if there are any hours, minutes or seconds fields.

- `encryption_proposed` (none/tls-psk) - Encryption mechanism of the communication channel between the two peers.

Additional information

As with creating a cluster peer through the CLI, the combinations of options must be valid in order for the create operation to succeed. The following list shows the combinations that will succeed and those that will fail:

- a passphrase only (fail)
- a peer IP address (fail)
- a passphrase with an expiration time > 7 days into the future (fail)
- peer IP address and a passphrase (OK)
- `generate_passphrase=true` (OK)
- any proposed encryption protocol (OK)
- an IPspace name or UUID (OK)
- a passphrase, peer IP address, and any proposed encryption protocol (OK)
- a non empty list initial allowed vservers peer names or UUIDs. (OK)

Learn more

- [DOC /cluster/peers](#)

Request Body

Name	Type	Description
<code>_links</code>	_links	
<code>authentication</code>	authentication	
<code>encryption</code>	encryption	
<code>initial_allowed_svms</code>	<code>array[initial_allowed_svms]</code>	The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship.
<code>ipspace</code>	ipspace	The IPspace of the local intercluster LIFs
<code>local_network</code>	local_network	Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.

Name	Type	Description
name	string	Optional name for the cluster peer relationship. By default it is the name of the remote cluster.
remote	remote	
status	status	
uuid	string	UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship.

Example request

```
{
  "_links": {
    "interfaces": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication": {
    "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
    "in_use": "ok",
    "state": "ok"
  },
  "encryption": {
    "proposed": "none",
    "state": "none"
  },
  "initial_allowed_svms": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "local_network": {
    "broadcast_domain": "bd1",
    "gateway": "10.1.1.1",
    "interfaces": {
      "ip_address": "10.10.10.7"
    },
    "netmask": "255.255.0.0"
  },
}
```

```

"name": "cluster2",
"remote": {
  "ip_addresses": {
  },
  "name": "cluster2",
  "serial_number": "4048820-60-9"
},
"status": {
  "state": "available",
  "update_time": "2017-01-25 11:20:13 UTC"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}

```

Response

Status: 201, Created

Name	Type	Description
_links	_links	
authentication	authentication	
ip_address	string	IPv4 or IPv6 address
name	string	Optional name for the cluster peer relationship. By default it is the name of the remote cluster, or a temporary name may be autogenerated for anonymous cluster peer offers.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication": {
    "expiry_time": "2017-01-25 11:20:13 UTC"
  },
  "ip_address": "10.10.10.7",
  "name": "cluster2"
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
4656069	Specifying a passphrase without remote IP addresses is not supported.
4656070	The encryption protocol is meaningful only with authenticated cluster peer relationships.
4656071	Cannot peer with a cluster bearing the same name as the local cluster.
4656072	The name must conform to the same rules as a cluster name.
4656074	Cannot check whether all nodes of this cluster support encryption.
4656077	Specify either remote IP addresses or generate_passphrase.
4656075	Cannot specify encryption: this operation requires an ECV of 9.6.0 or later.
4656079	No cluster nodes were found. Check your cluster configuration.
4656085	Cannot create an intercluster LIF with an empty list of local IP addresses.

Error Code	Description
4656087	The number of local intercluster IP addresses must be less than or equal to the number of available nodes.
4656086	Creating an intercluster LIF requires a broadcast domain that is in use within the IPspace.
4653365	IPspaces are unavailable with cluster peering: {ipspace}.
4656088	Found no ports matching the IPspace and the broadcast domain.
4656089	Found no matching entry for IPspace.
4656090	The given IPspace differs from the IPspace entry found.
4656091	Creating an intercluster LIF requires a network mask or a network mask length.
4656081	Creating an intercluster LIF requires a list of local IP addresses.
1966366	The System SVM of the cluster IPspace hosts Cluster LIFs only.
4656096	Creating an intercluster LIF requires an IPv4 or IPv6 address of the default router.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
interfaces	href	
self	href	

authentication

Name	Type	Description
expiry_time	string	The time when the passphrase will expire, in ISO 8601 duration format or date and time format. The default is 1 hour.
generate_passphrase	boolean	Auto generate a passphrase when true.
in_use	string	
passphrase	string	A password to authenticate the cluster peer relationship.
state	string	

encryption

Name	Type	Description
proposed	string	
state	string	

_links

Name	Type	Description
self	href	

initial_allowed_svms

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ipspace

The IPspace of the local intercluster LIFs

Name	Type	Description
_links	_links	
name	string	IPspace name
uuid	string	IPspace UUID

interfaces

Name	Type	Description
ip_address	string	IPv4 or IPv6 address

local_network

Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.

Name	Type	Description
broadcast_domain	string	Broadcast domain that is in use within the IPspace.
gateway	string	The IPv4 or IPv6 address of the default router.
interfaces	array[interfaces]	
netmask	string	IPv4 mask or netmask length.

remote

Name	Type	Description
ip_addresses	array[string]	The IPv4 addresses, IPv6 addresses, or hostnames of the peers.

Name	Type	Description
name	string	The name of the remote cluster.
serial_number	string	The serial number of the remote cluster.

status

Name	Type	Description
state	string	
update_time	string	The last time the state was updated.

cluster_peer

Name	Type	Description
_links	_links	
authentication	authentication	
encryption	encryption	
initial_allowed_svms	array[initial_allowed_svms]	The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship.
ipspace	ipspace	The IPspace of the local intercluster LIFs
local_network	local_network	Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.
name	string	Optional name for the cluster peer relationship. By default it is the name of the remote cluster.
remote	remote	
status	status	

Name	Type	Description
uuid	string	UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship.

authentication

Name	Type	Description
expiry_time	string	The date and time the passphrase will expire. The default expiry time is one hour.
passphrase	string	A password to authenticate the cluster peer relationship.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a cluster peer

DELETE /cluster/peers/{uuid}

Deletes a cluster peer.

Learn more

- [DOC /cluster/peers](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
4663070	Unable to delete cluster peer relationship due to an ongoing vserver migration.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a cluster peer instance

GET `/cluster/peers/{uuid}`

Retrieves a specific cluster peer instance.

Learn more

- [DOC /cluster/peers](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Cluster peer relationship UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
authentication	authentication	
encryption	encryption	
initial_allowed_svms	array[initial_allowed_svms]	The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship.
ipspace	ipspace	The IPspace of the local intercluster LIFs
local_network	local_network	Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.
name	string	Optional name for the cluster peer relationship. By default it is the name of the remote cluster.
remote	remote	
status	status	
uuid	string	UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship.

Example response

```
{
  "_links": {
    "interfaces": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication": {
    "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
    "in_use": "ok",
    "state": "ok"
  },
  "encryption": {
    "proposed": "none",
    "state": "none"
  },
  "initial_allowed_svms": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "local_network": {
    "broadcast_domain": "bd1",
    "gateway": "10.1.1.1",
    "interfaces": {
      "ip_address": "10.10.10.7"
    },
    "netmask": "255.255.0.0"
  },
}
```

```

"name": "cluster2",
"remote": {
  "ip_addresses": {
  },
  "name": "cluster2",
  "serial_number": "4048820-60-9"
},
"status": {
  "state": "available",
  "update_time": "2017-01-25 11:20:13 UTC"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
interfaces	href	
self	href	

authentication

Name	Type	Description
expiry_time	string	The time when the passphrase will expire, in ISO 8601 duration format or date and time format. The default is 1 hour.
generate_passphrase	boolean	Auto generate a passphrase when true.
in_use	string	
passphrase	string	A password to authenticate the cluster peer relationship.
state	string	

encryption

Name	Type	Description
proposed	string	
state	string	

_links

Name	Type	Description
self	href	

initial_allowed_svms

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ipspace

The IPspace of the local intercluster LIFs

Name	Type	Description
_links	_links	
name	string	IPspace name
uuid	string	IPspace UUID

interfaces

Name	Type	Description
ip_address	string	IPv4 or IPv6 address

local_network

Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.

Name	Type	Description
broadcast_domain	string	Broadcast domain that is in use within the IPspace.
gateway	string	The IPv4 or IPv6 address of the default router.
interfaces	array[interfaces]	
netmask	string	IPv4 mask or netmask length.

remote

Name	Type	Description
ip_addresses	array[string]	The IPv4 addresses, IPv6 addresses, or hostnames of the peers.

Name	Type	Description
name	string	The name of the remote cluster.
serial_number	string	The serial number of the remote cluster.

status

Name	Type	Description
state	string	
update_time	string	The last time the state was updated.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update a cluster peer instance

PATCH /cluster/peers/{uuid}

Updates a cluster peer instance.

Learn more

- [DOC /cluster/peers](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Cluster peer relationship UUID

Request Body

Name	Type	Description
_links	_links	
authentication	authentication	
encryption	encryption	
initial_allowed_svms	array[initial_allowed_svms]	The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship.
ipspace	ipspace	The IPspace of the local intercluster LIFs
local_network	local_network	Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.
name	string	Optional name for the cluster peer relationship. By default it is the name of the remote cluster.
remote	remote	
status	status	
uuid	string	UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship.

Example request

```
{
  "_links": {
    "interfaces": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication": {
    "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
    "in_use": "ok",
    "state": "ok"
  },
  "encryption": {
    "proposed": "none",
    "state": "none"
  },
  "initial_allowed_svms": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "local_network": {
    "broadcast_domain": "bd1",
    "gateway": "10.1.1.1",
    "interfaces": {
      "ip_address": "10.10.10.7"
    },
    "netmask": "255.255.0.0"
  },
}
```

```
"name": "cluster2",
"remote": {
  "ip_addresses": {
  },
  "name": "cluster2",
  "serial_number": "4048820-60-9"
},
"status": {
  "state": "available",
  "update_time": "2017-01-25 11:20:13 UTC"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
authentication	authentication	
ip_address	string	IPv4 or IPv6 address
name	string	Optional name for the cluster peer relationship. By default it is the name of the remote cluster, or a temporary name may be autogenerated for anonymous cluster peer offers.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication": {
    "expiry_time": "2017-01-25 11:20:13 UTC"
  },
  "ip_address": "10.10.10.7",
  "name": "cluster2"
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
4656070	The encryption protocol is meaningful only with authenticated cluster peer relationships.
4656072	The name must conform to the same rules as a cluster name.
4656073	Changing the encryption state requires the refreshing of the authentication passphrase.
4656076	Cluster peer modify was attempted mismatched IPv4 and IPv6 addresses.
4656075	Cannot specify encryption: this operation requires an ECV of Data ONTAP 9.6.0 or later.
4656084	Passphrase can only be modified with an authenticated cluster peer relationship.
4656082	Specify either a passphrase or "-generate-passphrase".
4656083	Cannot auto-generate a passphrase when "generate-passphrase" is false. Modifying a passphrase using an auto-generated passphrase requires "generate-passphrase" be true.
4656081	The remote IP address list is empty.

Error Code	Description
4656092	Cluster peer modify was attempted with a host name that did not resolve to an IPv4 or IPv6 address.
4655058	Expiration time cannot be more than 7 days in the future.
4653261	Error finding IPspace.
4656095	The address family of the specified peer addresses is not valid in this IPspace. Use /api/network/interfaces/ to verify that required LIFs are present and operational on each cluster node.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
interfaces	href	
self	href	

authentication

Name	Type	Description
expiry_time	string	The time when the passphrase will expire, in ISO 8601 duration format or date and time format. The default is 1 hour.
generate_passphrase	boolean	Auto generate a passphrase when true.
in_use	string	
passphrase	string	A password to authenticate the cluster peer relationship.
state	string	

encryption

Name	Type	Description
proposed	string	
state	string	

_links

Name	Type	Description
self	href	

initial_allowed_svms

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ipspace

The IPspace of the local intercluster LIFs

Name	Type	Description
_links	_links	
name	string	IPspace name
uuid	string	IPspace UUID

interfaces

Name	Type	Description
ip_address	string	IPv4 or IPv6 address

local_network

Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.

Name	Type	Description
broadcast_domain	string	Broadcast domain that is in use within the IPspace.
gateway	string	The IPv4 or IPv6 address of the default router.
interfaces	array[interfaces]	
netmask	string	IPv4 mask or netmask length.

remote

Name	Type	Description
ip_addresses	array[string]	The IPv4 addresses, IPv6 addresses, or hostnames of the peers.

Name	Type	Description
name	string	The name of the remote cluster.
serial_number	string	The serial number of the remote cluster.

status

Name	Type	Description
state	string	
update_time	string	The last time the state was updated.

cluster_peer

Name	Type	Description
_links	_links	
authentication	authentication	
encryption	encryption	
initial_allowed_svms	array[initial_allowed_svms]	The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship.
ipspace	ipspace	The IPspace of the local intercluster LIFs
local_network	local_network	Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.
name	string	Optional name for the cluster peer relationship. By default it is the name of the remote cluster.
remote	remote	
status	status	

Name	Type	Description
uuid	string	UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship.

authentication

Name	Type	Description
expiry_time	string	The date and time the passphrase will expire. The default expiry time is one hour.
passphrase	string	A password to authenticate the cluster peer relationship.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage cluster schedules

Cluster schedules endpoint overview

Overview

The `/cluster/schedules` API is used to view, create, and modify job schedules in a cluster.

Retrieving a job schedule

Job schedules can be retrieved by issuing a GET request to `/cluster/schedules`. It is also possible to retrieve a specific schedule when qualified by its UUID to `/cluster/schedules/{uuid}`. Queries on fields can be applied to retrieve all schedules that match the combined query.

Example

```
# The API:
/api/cluster/schedules/

# The call:
curl -X GET 'https://<mgmt-ip>/api/cluster/schedules?type=interval'

# The response:
{
  "records": [
    {
      "uuid": "08ceae53-0158-11e9-a82c-005056bb4301",
      "name": "RepositoryBalanceMonitorJobSchedule",
      "type": "interval",
      "interval": "PT10M",
      "_links": {
        "self": {
          "href": "/api/cluster/schedules/08ceae53-0158-11e9-a82c-005056bb4301"
        }
      }
    },
    {
      "uuid": "0941e980-0158-11e9-a82c-005056bb4301",
      "name": "Balanced Placement Model Cache Update",
      "type": "interval",
      "interval": "PT7M30S",
      "_links": {
        "self": {
          "href": "/api/cluster/schedules/0941e980-0158-11e9-a82c-005056bb4301"
        }
      }
    },
    {
      "uuid": "0944b975-0158-11e9-a82c-005056bb4301",
```

```

    "name": "Auto Balance Aggregate Scheduler",
    "type": "interval",
    "interval": "PT1H",
    "_links": {
      "self": {
        "href": "/api/cluster/schedules/0944b975-0158-11e9-a82c-005056bb4301"
      }
    }
  },
  {
    "uuid": "0c65f1fb-0158-11e9-a82c-005056bb4301",
    "name": "Application Templates ASUP Dump",
    "type": "interval",
    "interval": "P1D",
    "_links": {
      "self": {
        "href": "/api/cluster/schedules/0c65f1fb-0158-11e9-a82c-005056bb4301"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/cluster/schedules?type=interval"
  }
}
}

```

```
# The API:
/api/cluster/schedules/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/cluster/schedules/25312bd8-0158-11e9-a82c-005056bb4301'

# The response:
{
  "uuid": "25312bd8-0158-11e9-a82c-005056bb4301",
  "name": "monthly",
  "cluster": {
    "name": "rodan-tsundere",
    "uuid": "f3f9bbfa-0157-11e9-a82c-005056bb4301"
  },
  "type": "cron",
  "cron": {
    "minutes": [
      20
    ],
    "hours": [
      0
    ],
    "days": [
      1
    ]
  },
  "_links": {
    "self": {
      "href": "/api/cluster/schedules/25312bd8-0158-11e9-a82c-005056bb4301"
    }
  }
}
```

Creating a job schedule

A job schedule is created by issuing a POST request to `/cluster/schedules` to a node in the cluster. For a successful request, the POST request returns a status code of 201. Job schedules can be of either type "cron" or type "interval". A cron schedule is run at specific minutes within the hour, or hours of the day, days of the week, days of the month, or months of the year. An interval schedule runs repeatedly at fixed intervals.

Required fields

- **name** - Name of the job schedule You are required to provide a "minutes" field for a cron schedule. An "interval" field is required for an interval schedule. You must not provide both a "cron" field and an "interval" field. The schedule UUID is created by the system.

Cron schedule fields

- cron.minutes - Minutes within the hour (0 through 59)
- cron.hours - Hours of the day (0 through 23)
- cron.weekdays - Weekdays (0 through 6, where 0 is Sunday and 6 is Saturday.)
- cron.days - Days of the month (1 through 31)
- cron.months - Months of the year (1 through 12)

Interval schedule field

- interval - Length of time in ISO 8601 duration format

Example

```
# The API:
/api/cluster/schedules

# The call:
curl -X POST "https://<mgmt-ip>/api/cluster/schedules" -d body

# The response of a successful POST is empty.
Example body to create an interval schedule with a 1-week interval:
{
  "name": "test_interval_1",
  "interval": "P1W"
}
Example body to create a cron schedule that runs daily at 12:05 :
{
  "name": "test_cron_1",
  "cron":
  {
    "minutes": [ 5 ],
    "hours": [ 12 ]
  }
}
```

Optional fields

By default, the schedule is owned by the local cluster. In a MetroCluster configuration, the partner cluster can be specified if the local cluster is in the switchover state.

- cluster.name - Name of the cluster owning the schedule
- cluster.uuid - UUID of the cluster owning the schedule

Records field

Multiple schedules can be created in one request by providing an array of named records with schedule

entries. Each entry must follow the required and optional fields listed above.

Updating a job schedule

The following fields of an existing schedule can be modified:

- cron.minutes
- cron.hours
- cron.weekdays
- cron.days
- cron.months
- interval Note: The name, cluster, and type of schedule cannot be modified. Also, you cannot modify a cron field of an interval schedule, or the interval field of a cron schedule. Queries on fields can be applied to modify all schedules that match the combined query.

Example

```
# The API:
/api/cluster/schedules/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/cluster/schedules/{uuid}" -d body

# The response of a successful PATCH is empty.
Example body to modify an interval schedule with a 2-day and 5-minute
interval:
{
  "interval": "P2DT5M"
}
Example body to modify a cron schedule to run Mondays at 2:
{
  "cron":
  {
    "hours": [ 2 ],
    "weekdays": [ 1 ]
  }
}
```

Deleting a job schedule

Job schedules can be deleted based on their UUID. Queries on fields can be applied to delete all schedules that match the combined query.

Example

```
# The API:
/api/cluster/schedules/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/cluster/schedules/{uuid}"

# The response of a successful DELETE of one schedule is empty.
```

```
# The API:
/api/cluster/schedules/

# The call:
curl -X DELETE "https://<mgmt-ip>/api/cluster/schedules/?name=test*"

# The response of a successful DELETE indicates the number of schedules
affected:
{
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/cluster/schedules?name=test*"
    }
  }
}
```

MetroCluster configurations

In a MetroCluster configuration, user-created schedules owned by the local cluster are replicated to the partner cluster. Likewise, user-created schedules owned by the partner cluster are replicated to the local cluster. The owning cluster for a particular schedule is shown in the "cluster.name" and "cluster.uuid" fields. Normally, only schedules owned by the local cluster can be created, modified, and deleted on the local cluster. However, when a MetroCluster configuration is in switchover, the cluster in switchover state can create, modify, and delete schedules owned by the partner cluster.

Retrieve schedules

GET /cluster/schedules

Retrieves a schedule.

Learn more

- [DOC /cluster/schedules](#)

Parameters

Name	Type	In	Required	Description
cron.hours	integer	query	False	Filter by cron.hours
cron.months	integer	query	False	Filter by cron.months
cron.minutes	integer	query	False	Filter by cron.minutes
cron.weekdays	integer	query	False	Filter by cron.weekdays
cron.days	integer	query	False	Filter by cron.days
cluster.name	string	query	False	Filter by cluster.name
cluster.uuid	string	query	False	Filter by cluster.uuid
interval	string	query	False	Filter by interval
uuid	string	query	False	Filter by uuid
type	string	query	False	Filter by type
name	string	query	False	Filter by name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.

Name	Type	In	Required	Description
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[schedule]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": 1,
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "cluster": {
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "cron": {
      "days": {
      },
      "hours": {
      },
      "minutes": {
      },
      "months": {
      },
      "weekdays": {
      }
    },
    "interval": "P1DT2H3M4S",
    "type": "cron",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
459760	The schedule specified is not a valid schedule.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

cluster

The cluster that owns the schedule. Defaults to the local cluster.

Name	Type	Description
name	string	Cluster name
uuid	string	Cluster UUID

cron

Details for schedules of type cron.

Name	Type	Description
days	array[integer]	The days of the month the schedule runs. Leave empty for all.
hours	array[integer]	The hours of the day the schedule runs. Leave empty for all.
minutes	array[integer]	The minutes the schedule runs. Required on POST for a cron schedule.

Name	Type	Description
months	array[integer]	The months of the year the schedule runs. Leave empty for all.
weekdays	array[integer]	The weekdays the schedule runs. Leave empty for all.

schedule

Complete schedule information

Name	Type	Description
_links	_links	
cluster	cluster	The cluster that owns the schedule. Defaults to the local cluster.
cron	cron	Details for schedules of type cron.
interval	string	An ISO-8601 duration formatted string.
name	string	Schedule Name. Required in the URL or POST body.
type	string	Schedule type
uuid	string	Job schedule UUID

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a schedule

POST `/cluster/schedules`

Create a schedule.

Required Fields

- **name** - Name of the job schedule It is required to provide a minutes field for a cron schedule. An interval field is required for an interval schedule. You must not provide both a cron field and an interval field.

Learn more

- [DOC /cluster/schedules](#)

Request Body

Name	Type	Description
<code>_links</code>	_links	
<code>cluster</code>	cluster	The cluster that owns the schedule. Defaults to the local cluster.
<code>cron</code>	cron	Details for schedules of type cron.
<code>interval</code>	string	An ISO-8601 duration formatted string.
<code>name</code>	string	Schedule Name. Required in the URL or POST body.
<code>type</code>	string	Schedule type
<code>uuid</code>	string	Job schedule UUID

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cron": {
    "days": {
    },
    "hours": {
    },
    "minutes": {
    },
    "months": {
    },
    "weekdays": {
    }
  },
  "interval": "P1DT2H3M4S",
  "type": "cron",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
459760	The schedule specified is not a valid schedule.

Error Code	Description
458788	The schedule specified is not a valid schedule.
459764	Cannot create a schedule with the same name as an existing schedule from the MetroCluster partner cluster but of a different schedule type.
460783	As this is a MetroCluster configuration and the local cluster is waiting for switchback, changes to non-system schedules are not allowed.
459763	Schedule cannot be created locally using the remote cluster name as the owner.
460784	An error occurred creating the remote cluster version of this schedule.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

cluster

The cluster that owns the schedule. Defaults to the local cluster.

Name	Type	Description
name	string	Cluster name
uuid	string	Cluster UUID

cron

Details for schedules of type cron.

Name	Type	Description
days	array[integer]	The days of the month the schedule runs. Leave empty for all.
hours	array[integer]	The hours of the day the schedule runs. Leave empty for all.
minutes	array[integer]	The minutes the schedule runs. Required on POST for a cron schedule.
months	array[integer]	The months of the year the schedule runs. Leave empty for all.
weekdays	array[integer]	The weekdays the schedule runs. Leave empty for all.

schedule

Complete schedule information

Name	Type	Description
_links	_links	
cluster	cluster	The cluster that owns the schedule. Defaults to the local cluster.
cron	cron	Details for schedules of type cron.
interval	string	An ISO-8601 duration formatted string.
name	string	Schedule Name. Required in the URL or POST body.
type	string	Schedule type
uuid	string	Job schedule UUID

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a schedule

DELETE /cluster/schedules/{uuid}

Deletes a schedule.

Learn more

- [DOC /cluster/schedules](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

Response

```
Status: 200, Ok
```

Error

```
Status: Default
```

ONTAP Error Response Codes

Error Code	Description
459758	Cannot delete a job schedule that is in use. Remove all references to the schedule, and then try to delete again.
459761	Schedule cannot be deleted on this cluster because it is replicated from the remote cluster.
459762	The schedule cannot be deleted because it is a system-level schedule.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a schedule

GET /cluster/schedules/{uuid}

Retrieves a schedule.

Learn more

- [DOC /cluster/schedules](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Schedule UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
cluster	cluster	The cluster that owns the schedule. Defaults to the local cluster.
cron	cron	Details for schedules of type cron.
interval	string	An ISO-8601 duration formatted string.
name	string	Schedule Name. Required in the URL or POST body.
type	string	Schedule type
uuid	string	Job schedule UUID

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cron": {
    "days": {
    },
    "hours": {
    },
    "minutes": {
    },
    "months": {
    },
    "weekdays": {
    }
  },
  "interval": "P1DT2H3M4S",
  "type": "cron",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

cluster

The cluster that owns the schedule. Defaults to the local cluster.

Name	Type	Description
name	string	Cluster name
uuid	string	Cluster UUID

cron

Details for schedules of type cron.

Name	Type	Description
days	array[integer]	The days of the month the schedule runs. Leave empty for all.
hours	array[integer]	The hours of the day the schedule runs. Leave empty for all.
minutes	array[integer]	The minutes the schedule runs. Required on POST for a cron schedule.
months	array[integer]	The months of the year the schedule runs. Leave empty for all.
weekdays	array[integer]	The weekdays the schedule runs. Leave empty for all.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update a schedule

PATCH `/cluster/schedules/{uuid}`

Updates a schedule. Note that you cannot modify a cron field of an interval schedule, or the interval field of a cron schedule.

Learn more

- [DOC /cluster/schedules](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Schedule UUID

Request Body

Name	Type	Description
<code>_links</code>	_links	
cluster	cluster	The cluster that owns the schedule. Defaults to the local cluster.
cron	cron	Details for schedules of type cron.

Name	Type	Description
interval	string	An ISO-8601 duration formatted string.
name	string	Schedule Name. Required in the URL or POST body.
type	string	Schedule type
uuid	string	Job schedule UUID

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cron": {
    "days": {
    },
    "hours": {
    },
    "minutes": {
    },
    "months": {
    },
    "weekdays": {
    }
  },
  "interval": "P1DT2H3M4S",
  "type": "cron",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
459760	The schedule specified is not a valid schedule.
458788	The schedule specified is not a valid schedule.
460783	As this is a MetroCluster configuration and the local cluster is waiting for switchback, changes to non-system schedules are not allowed.
459761	Schedule cannot be modified on this cluster because it is replicated from the remote cluster.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

cluster

The cluster that owns the schedule. Defaults to the local cluster.

Name	Type	Description
name	string	Cluster name
uuid	string	Cluster UUID

cron

Details for schedules of type cron.

Name	Type	Description
days	array[integer]	The days of the month the schedule runs. Leave empty for all.
hours	array[integer]	The hours of the day the schedule runs. Leave empty for all.
minutes	array[integer]	The minutes the schedule runs. Required on POST for a cron schedule.
months	array[integer]	The months of the year the schedule runs. Leave empty for all.
weekdays	array[integer]	The weekdays the schedule runs. Leave empty for all.

schedule

Complete schedule information

Name	Type	Description
_links	_links	
cluster	cluster	The cluster that owns the schedule. Defaults to the local cluster.
cron	cron	Details for schedules of type cron.
interval	string	An ISO-8601 duration formatted string.
name	string	Schedule Name. Required in the URL or POST body.
type	string	Schedule type
uuid	string	Job schedule UUID

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage cluster software

Cluster software endpoint overview

Overview

ONTAP cluster software API retrieves and displays relevant information about the software profile, software packages collection, and software history collection. The API retrieves the information about all software packages present in the cluster, or specific software package.

The POST request provides the ability to download a software package from an HTTP or FTP server. The PATCH request provides the option to upgrade the cluster software version. The client can validate the package before triggering the update by selecting the `validate_only` field. Setting the `version` field triggers the installation of the package in the cluster. The client can pause, resume, or cancel any ongoing software upgrade by selecting `action`. The DELETE request can remove a specific software package present in the cluster.

Examples

Retrieving software profile information

The following example shows how to retrieve software profile information. The client can check the validation results after selecting `validate_only` field. Upgrade progress information is available after an upgrade has started.

```
# The API:
/api/cluster/software

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software?return_timeout=15" -H
"accept: application/hal+json"

# The response:
{
  "validation_results": [
    {
      "update_check": "NFS mounts",
      "status": "warning",
      "message": [
        {
          "code": 166,
          "message": "Use NFS hard mounts, if possible.",
          "arguments": [
            "string"
          ]
        }
      ],
      "action": [
```

```

    {
      "code": 166,
      "message": "Use NFS hard mounts, if possible.",
      "arguments": [
        "string"
      ]
    }
  ]
},
"version": "9.5.0",
"pending_version": "9.6.0",
"nodes": [
  {
    "node": "sti70-vsim-ucs165n",
    "version": "9.5.0"
  }
],
"metrocluster": {
  "progress_summary": "Update paused by user",
  "progress_details": "Installing Data ONTAP software image on cluster
\\\"sti70-vsim-ucs165n_siteA\\\".",
  "clusters": [
    {
      "name": "sti70-vsim-ucs165n_siteA",
      "uuid": "720f046c-4b13-11e9-9c34-005056ac5626",
      "estimated_duration": 3480,
      "elapsed_duration": 0,
      "state": "waiting"
    }
  ],
},
"state": "in_progress",
"start_time": "2018-05-21T09:53:04+05:30",
"end_time": "2018-05-21T11:53:04+05:30",
"estimated_time": 5220,
"elapsed_time": 2140,
"update_details": [
  {
    "phase": "Data ONTAP updates",
    "state": "in_progress",
    "estimated_duration": 4620,
    "elapsed_duration": 29,
    "node": {
      "name": "sti70-vsim-ucs165n"
    }
  }
]

```



```

    }
  ],
  "status_details": [
    {
      "name": "do-download-job",
      "state": "completed",
      "message": "Image update complete",
      "action": "",
      "start_time": "2018-05-21T09:53:04+05:30",
      "end_time": "2018-05-21T11:53:04+05:30",
      "node": {
        "name": "sti70-vsimg-ucs165n"
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/cluster/software/"
    }
  }
}

```

Upgrading the software version

The following example shows how to upgrade cluster software. Setting the `version` field triggers the installation of the package. The client can select the `validate_only` field to validate the package before the installation starts. Setting `skip_warning` as `true` ignores the validation warning before the installation starts. Setting the `action` field performs a pause, resume, or 'cancel' to an ongoing upgrade. An upgrade can only be resumed if it is in the paused state.

The client can start the upgrade process at the cluster-level. There are no options available to start the upgrade for a specific node or HA pairs.

1. Validating the package and verifying the validation results

The following example shows how to validate a cluster software package. The client has to validate the package before the software upgrade. The client must set the `validate_only` field to `true` to start the validation. The client can check for validation results in the GET `/cluster/software` endpoint.

```
# The API:
/api/cluster/software

# The call:
curl -X PATCH "https://<mgmt_ip>/api/cluster/software?validate_only=true"
-H "accept: application/json" -H "Content-Type: application/hal+json" -d
'{ "version": "9.5.0"}'

# The response:
{
  "job": {
    "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
      }
    }
  }
}
```

The call to validate the software cluster version returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a state and a message to indicate the progress of the job. When the job is complete and the application has been fully created, the message indicates success and the state field of the job is set to success.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc" -H "accept: application/hal+json"

# The response:
{
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "description": "PATCH /api/cluster/software",
  "state": "success",
  "message": "success",
  "code": 0,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
```

The client can check for validation results in the GET /cluster/software endpoint. The following example shows how to check the validation warnings and errors after setting the `validate_only` field to true.

```
# The API:
/api/cluster/software

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software" -H "accept: application/hal+json"

# The response:
{
  "version": "9.7.0",
  "validation_results": [
    {
      "update_check": "High Availability status",
      "status": "error",
      "message": "Cluster HA is not configured in the cluster. Storage failover is not enabled on node \"node1\", \"node2\".",
      "action": "Check cluster HA configuration. Check storage failover status."
    }
  ]
}
```

```

    },
    {
      "update_check": "Manual checks",
      "status": "warning",
      "message": "Manual validation checks need to be performed. Refer to
the Upgrade Advisor Plan or \"Performing manual checks before an automated
cluster upgrade\" section in the \"Clustered Data ONTAP Upgrade Express
Guide\" for the remaining validation checks that need to be performed
before update. Failing to do so can result in an update failure or an I/O
disruption.",
      "action": "Refer to the Upgrade Advisor Plan or \"Performing manual
checks before an automated cluster upgrade\" section in the \"Clustered
Data ONTAP Upgrade Express Guide\" for the remaining validation checks
that need to be performed before update."
    }
  ],
  "nodes": [
    {
      "node": "node1",
      "version": "9.7.0"
    },
    {
      "node": "node2",
      "version": "9.7.0"
    }
  ],
  "state": "failed",
  "elapsed_duration": 56,
  "estimated_duration": 600,
  "_links": {
    "self": {
      "href": "/api/cluster/software"
    }
  }
}

```

2. Updating the cluster

The following example shows how to initiate a cluster software upgrade. The client must validate the package before the software upgrade starts. The client must set the `skip_warnings` field to `true` in order to skip any validation warnings and start the software package upgrade.

```
# The API:
/api/cluster/software

# The call:
curl -X PATCH "https://<mgmt_ip>/api/cluster/software?skip_warnings=true"
-H "accept: application/json" -H "Content-Type: application/hal+json" -d
'{ "version": "9.5.0"}'

# The response:
{
  "job": {
    "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
      }
    }
  }
}
```

The call to update the software cluster version returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a state and a message to indicate the progress of the job. When the job is complete and the application has been fully created, the message indicates success and the state field of the job is set to success.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc" -H "accept: application/hal+json"

# The response:
{
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "description": "PATCH /api/cluster/software",
  "state": "success",
  "message": "success",
  "code": 0,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
```

The client can check the update progress information in the GET /cluster/software endpoint. The following example shows how to check the progress of an update after setting the `skip_warnings` field to `true`.

```
# The API:
/api/cluster/software

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software" -H "accept: application/hal+json"

# The response:
{
  "version": "9.7.0",
  "validation_results": [
    {
      "update_check": "Manual checks",
      "status": "warning",
      "message": "Manual validation checks need to be performed. Refer to the Upgrade Advisor Plan or \"Performing manual checks before an automated cluster upgrade\" section in the \"Clustered Data ONTAP Upgrade Express Guide\" for the remaining validation checks that need to be performed before update. Failing to do so can result in an update failure or an I/O disruption.",
      "action": "Refer to the Upgrade Advisor Plan or \"Performing manual
```

checks before an automated cluster upgrade\" section in the \"Clustered Data ONTAP Upgrade Express Guide\" for the remaining validation checks that need to be performed before update."

```
    }
  ],
  "nodes": [
    {
      "node": "node1",
      "version": "9.7.0"
    },
    {
      "node": "node2",
      "version": "9.7.0"
    }
  ],
  "pending_version": "9.7.0",
  "state": "in_progress",
  "elapsed_duration": 63,
  "estimated_duration": 5220,
  "status_details": [
    {
      "name": "do-download-job",
      "status": "running",
      "message": "",
      "action": "",
      "start_time": "2019-01-14T23:12:14+05:30",
      "end_time": "2019-01-14T23:12:14+05:30",
      "node": {
        "name": "node1"
      }
    },
    {
      "name": "do-download-job",
      "status": "running",
      "message": "",
      "action": "",
      "start_time": "2019-01-14T23:12:14+05:30",
      "end_time": "2019-01-14T23:12:14+05:30",
      "node": {
        "name": "node2"
      }
    }
  ],
  "update_details": [
    {
      "phase": "Data ONTAP updates",
```

```
"status": "in-progress",
"estimated_duration": 4620,
"elapsed_duration": 10,
"node": {
  "name": "node1"
}
},
{
  "phase": "Data ONTAP updates",
  "status": "in-progress",
  "estimated_duration": 4620,
  "elapsed_duration": 10,
  "node": {
    "name": "node2"
  }
}
],
"_links": {
  "self": {
    "href": "/api/cluster/software"
  }
}
}
```

3. Pausing/resuming/cancelling the upgrade

The following example shows how to pause an ongoing cluster software package upgrade. The client must set the `action` field to `pause`, `resume`, or `cancel` which pauses, resumes or cancels the upgrade respectively. Not all update operations support these actions. An update can only be resumed if it is in the paused state.

```
# The API:
/api/cluster/software

# The call:
curl -X PATCH "https://<mgmt_ip>/api/cluster/software?action=pause" -H
"accept: application/json" -H "Content-Type: application/hal+json" -d '{
"version": "9.5.0"}'

# The response:
{
  "job": {
    "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
      }
    }
  }
}
```

The call to update the software cluster version returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a state and a message to indicate the progress of the job. When the job is complete and the application has been fully created, the message indicates success and the state field of the job is set to success.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc" -H "accept: application/hal+json"

# The response:
{
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "description": "PATCH /api/cluster/software",
  "state": "success",
  "message": "success",
  "code": 0,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
```

The client can check the progress of the upgrade in the GET /cluster/software endpoint. The following example shows how to check the progress of the pause upgrade state after setting the action field to pause.

```
# The API:
/api/cluster/software

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software" -H "accept: application/hal+json"

# The response:
{
  "version": "9.7.0",
  "validation_results": [
    {
      "update_check": "Manual checks",
      "status": "warning",
      "message": "Manual validation checks need to be performed. Refer to the Upgrade Advisor Plan or \"Performing manual checks before an automated cluster upgrade\" section in the \"Clustered Data ONTAP Upgrade Express Guide\" for the remaining validation checks that need to be performed"
    }
  ]
}
```

```

before update. Failing to do so can result in an update failure or an I/O
disruption.",
    "action": "Refer to the Upgrade Advisor Plan or \"Performing manual
checks before an automated cluster upgrade\" section in the \"Clustered
Data ONTAP Upgrade Express Guide\" for the remaining validation checks
that need to be performed before update.",
    }
],
"nodes": [
    {
        "node": "node1",
        "version": "9.7.0"
    },
    {
        "node": "node2",
        "version": "9.7.0"
    }
],
"pending_version": "9.7.0",
"state": "pause_pending",
"elapsed_duration": 103,
"estimated_duration": 5220,
"status_details": [
    {
        "status": "in-progress",
        "message": "Installing Data ONTAP software image.",
        "action": "",
        "start_time": "2019-01-08T02:54:36+05:30",
        "node": {
            "name": "node1"
        }
    },
    {
        "status": "in-progress",
        "message": "Installing Data ONTAP software image.",
        "action": "",
        "start_time": "2019-01-08T02:54:36+05:30",
        "node": {
            "name": "node2"
        }
    }
],
"update_details": [
    {
        "phase": "Pre-update checks",
        "status": "completed",

```

```

    "estimated_duration": 600,
    "elapsed_duration": 54,
    "node": {
      "name": "node1"
    }
  },
  {
    "phase": "Data ONTAP updates",
    "status": "pause-pending",
    "estimated_duration": 4620,
    "elapsed_duration": 49,
    "node": {
      "name": "node2"
    }
  },
  {
    "phase": "Data ONTAP updates",
    "status": "pause-pending",
    "estimated_duration": 4620,
    "elapsed_duration": 49
  }
],
"_links": {
  "self": {
    "href": "/api/cluster/software"
  }
}
}

```

Downloading the software package

The following example shows how to download the software package from an HTTP or FTP server. The client provides the `url`, `username`, and `password` to start the download of the software package to the cluster.

```
# The API:
/api/cluster/software/download

# The call:
curl -X POST "https://<mgmt-
ip>/api/cluster/software/download?return_timeout=0" -H "accept:
application/json" -H "Content-Type: application/hal+json" -d '{ "url":
"http://nbsweb.eng.btc.netapp.in/~suvadipd/99/image1.tgz", "username":
"admin", "password": "*****"}'

# The response:
{
  "job": {
    "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
      }
    }
  }
}
```

The call to download the software package returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a state and a message to indicate the progress of the job. When the job is complete and the application has been fully created, the message indicates success and the job state field is set to success.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc" -H "accept: application/hal+json"

# The response:
{
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "description": "POST /api/cluster/software/download",
  "state": "success",
  "message": "success",
  "code": 0,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
```

Retrieving cluster software packages information

The following example shows how to retrieve the ONTAP software packages in a cluster.

```
# The API:
/api/cluster/software/packages

# The call:
curl -X GET "https://<mgmt-
ip>/api/cluster/software/packages?return_records=true&return_timeout=15"
-H "accept: application/hal+json"

# The response:
{
  "records": [
    {
      "version": "9.7.0",
      "_links": {
        "self": {
          "href": "/api/cluster/software/packages/9.7.0"
        }
      }
    },
    {
      "version": "9.5.0",
      "_links": {
        "self": {
          "href": "/api/cluster/software/packages/9.5.0"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/cluster/software/packages"
    }
  }
}
```

The following example shows how to retrieve the details of a given cluster software package.

```
# The API:
/api/cluster/software/packages/{version}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software/packages/9.7.0" -H
"accept: application/hal+json"

# The response:
{
  "version": "9.7.0",
  "create_time": "2018-05-21T10:06:59+05:30",
  "_links": {
    "self": {
      "href": "/api/cluster/software/packages/9.7.0"
    }
  }
}
```

Deleting a cluster software package

The following example shows how to delete a package from the cluster. The client needs to provide the package version that they want to delete. The software package delete creates a job to perform the delete operation.

```
# The API:
/api/cluster/software/packages/{version}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/cluster/software/packages/9.6.0" -H
"accept: application/hal+json"

# The response:
{
  "job": {
    "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
      }
    }
  }
}
```

The call to delete the package returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a state and a message to indicate the progress of the job. When the job is complete and the application has been fully created, the message indicates success and the job state field is set to success.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc" -H "accept: application/hal+json"

# The response:
{
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "description": "DELETE /api/cluster/software/packages/9.6.0",
  "state": "success",
  "message": "success",
  "code": 0,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
```

HTTPS error codes

The following is a list of possible error codes that can be returned during a package delete operation.

ONTAP Error Response codes

Error codes	Description
10551315	Package store is empty
10551322	Error in retrieving package cleanup status
10551323	Error in cleaning up package information on a node
10551324	Error in cleaning up package information on both nodes
10551325	Package does not exist on the system
10551326	Error in deleting older package cleanup tasks
10551346	Package delete failed since a validation is in progress
10551347	Package delete failed since an update is in progress
10551367	A package synchronization is in progress
10551388	Package delete operation timed out

Retrieving software installation history information

The following example shows how to

- retrieve the software package installation history information.
- display specific node level software installation history information.
- provide all the attributes by default in response when the self referential link is not present.

```
# The API:
/api/cluster/software/history

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software/history" -H "accept:
application/hal+json"

# The response:
{
  "node": {
    "uuid": "58cd3a2b-af63-11e8-8b0d-0050568e7279",
    "name": "sti70-vsimsim-ucs165n",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/58cd3a2b-af63-11e8-8b0d-0050568e7279"
      }
    }
  },
  "start_time": "2018-09-03T16:18:46+05:30",
  "state": "successful"
  "from_version": "9.4.0",
  "to_version": "9.5.0",
  "end_time": "2018-05-21T10:14:51+05:30"
}
```

Retrieve the cluster software profile

GET /cluster/software

Retrieves the software profile of a cluster.

Related ONTAP commands

- cluster image show
- cluster image show-update-progress

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
action	string	User triggered action to apply to the install operation
elapsed_duration	integer	Elapsed time during the upgrade or validation operation
estimated_duration	integer	Estimated time remaining until completion of the upgrade or validation operation.
metrocluster	metrocluster	
nodes	array[software_node_reference]	List of nodes and active versions.

Name	Type	Description
pending_version	string	Version being installed on the system. <ul style="list-style-type: none"> • example: ONTAP_X_1 • readOnly: 1
state	string	Operational state of the upgrade
status_details	array[software_status_details_reference]	Display status details.
update_details	array[software_update_details_reference]	Display update process details.
validation_results	array[software_validation_reference]	List of validation warnings, errors, and advice.
version	string	Version of ONTAP installed and currently active on the system. During PATCH, using the 'validate_only' parameter on the request executes pre-checks, but does not perform the full installation. <ul style="list-style-type: none"> • example: ONTAP_X

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "action": "pause",
  "elapsed_duration": 2140,
  "estimated_duration": 5220,
  "metrocluster": {
    "clusters": {
      "elapsed_duration": 2140,
      "estimated_duration": 3480,
      "name": "cluster_A",
      "state": "in_progress"
    },
    "progress_details": "Switchover in progress.",
    "progress_summary": "MetroCluster updated successfully."
  },
  "nodes": {
    "name": "node1",
    "version": "ONTAP_X"
  },
  "pending_version": "ONTAP_X_1",
  "state": "completed",
  "status_details": {
    "action": "string",
    "end_time": "2019-02-02 19:00:00 UTC",
    "message": "Post-update checks successful",
    "name": "initialize",
    "node": {
      "name": "node1"
    },
    "start_time": "2019-02-02 19:00:00 UTC",
    "state": "failed"
  },
  "update_details": {
    "elapsed_duration": 2100,
    "estimated_duration": 4620,
    "node": {
      "name": "node1"
    },
    "phase": "Pre-update checks",
    "state": "failed"
  }
}
```

```

    },
    "validation_results": {
      "action": "string",
      "message": "string",
      "status": "warning",
      "update_check": "nfs_mounts"
    },
    "version": "ONTAP_X"
  }
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

software_mcc_reference

Name	Type	Description
elapsed_duration	integer	Elapsed duration of update time (in seconds) in MetroCluster.
estimated_duration	integer	Estimated duration of update time (in seconds) in MetroCluster.
name	string	Name of the site in MetroCluster.
state		Upgrade state of MetroCluster.

metrocluster

Name	Type	Description
clusters	array[software_mcc_reference]	List of MetroCluster sites, statuses, and active versions.
progress_details	string	MetroCluster update progress details.
progress_summary	string	MetroCluster update progress summary.

software_node_reference

Name	Type	Description
name	string	Name of the node.

Name	Type	Description
version	string	ONTAP version of the node. <ul style="list-style-type: none"> • example: ONTAP_X • readOnly: 1

node

Name	Type	Description
name	string	Name of the node to be retrieved for status details.

software_status_details_reference

Name	Type	Description
action	string	Corrective action to be taken to resolve the status error.
end_time	string	End time for each status phase.
message	string	Detailed message of the phase details.
name	string	Name of the phase to be retrieved for status details.
node	node	
start_time	string	Start time for each status phase.
state	string	Status of the phase

node

Name	Type	Description
name	string	Name of the node to be retrieved for update details.

software_update_details_reference

Name	Type	Description
elapsed_duration	integer	Elapsed duration for each update phase

Name	Type	Description
estimated_duration	integer	Estimated duration for each update phase
node	node	
phase	string	Phase details
state	string	State of the update phase

software_validation_reference

Name	Type	Description
action	string	Corrective action to resolve errors or warnings for update checks.
message	string	Details of the error or warning encountered by the update check.
status	string	Status of this update check.
update_check	string	Name of the update check to be validated.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the cluster software version

PATCH /cluster/software

Upgrades the cluster software version. Setting `version` triggers the installation of the package to start. To validate the package for installation but not perform the installation, use the `validate_only` field on request. Important note:

- Setting 'version' triggers the package installation.
- To validate the package for installation but not perform the installation, use the `validate_only` field on the request.

Required properties

- `version` - Software version to be installed on the cluster

Recommended optional parameters

- `validate_only` - Required to validate a software package before an upgrade
- `skip_warnings` - Used to skip validation warnings when starting a software upgrade
- `action` - Used to pause, resume, or cancel an ongoing software upgrade

Related ONTAP commands

- `cluster image validate`
- `cluster image update`
- `cluster image pause-update`
- `cluster image resume-update`
- `cluster image cancel-update`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
<code>validate_only</code>	boolean	query	False	Validate the operation and its parameters, without actually performing the operation.
<code>skip_warnings</code>	boolean	query	False	Ignore warnings and proceed with the install.

Name	Type	In	Required	Description
action	string	query	False	<p>Requests an upgrade to pause, resume, or cancel. Note that not all upgrades support these actions. An upgrade can only be resumed if it is in the paused state. When a request to cancel an upgrade is successful, the upgrade state changes to either success or failure.</p> <ul style="list-style-type: none"> enum: ["pause", "resume", "cancel"]

Request Body

Name	Type	Description
_links	_links	
action	string	User triggered action to apply to the install operation
elapsed_duration	integer	Elapsed time during the upgrade or validation operation
estimated_duration	integer	Estimated time remaining until completion of the upgrade or validation operation.
metrocluster	metrocluster	
nodes	array[software_node_reference]	List of nodes and active versions.
pending_version	string	<p>Version being installed on the system.</p> <ul style="list-style-type: none"> example: ONTAP_X_1 readOnly: 1
state	string	Operational state of the upgrade

Name	Type	Description
status_details	array[software_status_details_reference]	Display status details.
update_details	array[software_update_details_reference]	Display update process details.
validation_results	array[software_validation_reference]	List of validation warnings, errors, and advice.
version	string	<p>Version of ONTAP installed and currently active on the system. During PATCH, using the 'validate_only' parameter on the request executes pre-checks, but does not perform the full installation.</p> <ul style="list-style-type: none"> • example: ONTAP_X

Example request

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "action": "pause",
  "elapsed_duration": 2140,
  "estimated_duration": 5220,
  "metrocluster": {
    "clusters": {
      "elapsed_duration": 2140,
      "estimated_duration": 3480,
      "name": "cluster_A",
      "state": "in_progress"
    },
    "progress_details": "Switchover in progress.",
    "progress_summary": "MetroCluster updated successfully."
  },
  "nodes": {
    "name": "node1",
    "version": "ONTAP_X"
  },
  "pending_version": "ONTAP_X_1",
  "state": "completed",
  "status_details": {
    "action": "string",
    "end_time": "2019-02-02 19:00:00 UTC",
    "message": "Post-update checks successful",
    "name": "initialize",
    "node": {
      "name": "node1"
    },
    "start_time": "2019-02-02 19:00:00 UTC",
    "state": "failed"
  },
  "update_details": {
    "elapsed_duration": 2100,
    "estimated_duration": 4620,
    "node": {
      "name": "node1"
    },
    "phase": "Pre-update checks",
    "state": "failed"
  }
}
```

```
},
"validation_results": {
  "action": "string",
  "message": "string",
  "status": "warning",
  "update_check": "nfs_mounts"
},
"version": "ONTAP_X"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

software_mcc_reference

Name	Type	Description
elapsed_duration	integer	Elapsed duration of update time (in seconds) in MetroCluster.
estimated_duration	integer	Estimated duration of update time (in seconds) in MetroCluster.
name	string	Name of the site in MetroCluster.
state		Upgrade state of MetroCluster.

metrocluster

Name	Type	Description
clusters	array[software_mcc_reference]	List of MetroCluster sites, statuses, and active versions.
progress_details	string	MetroCluster update progress details.
progress_summary	string	MetroCluster update progress summary.

software_node_reference

Name	Type	Description
name	string	Name of the node.

Name	Type	Description
version	string	ONTAP version of the node. <ul style="list-style-type: none"> • example: ONTAP_X • readOnly: 1

node

Name	Type	Description
name	string	Name of the node to be retrieved for status details.

software_status_details_reference

Name	Type	Description
action	string	Corrective action to be taken to resolve the status error.
end_time	string	End time for each status phase.
message	string	Detailed message of the phase details.
name	string	Name of the phase to be retrieved for status details.
node	node	
start_time	string	Start time for each status phase.
state	string	Status of the phase

node

Name	Type	Description
name	string	Name of the node to be retrieved for update details.

software_update_details_reference

Name	Type	Description
elapsed_duration	integer	Elapsed duration for each update phase

Name	Type	Description
estimated_duration	integer	Estimated duration for each update phase
node	node	
phase	string	Phase details
state	string	State of the update phase

software_validation_reference

Name	Type	Description
action	string	Corrective action to resolve errors or warnings for update checks.
message	string	Details of the error or warning encountered by the update check.
status	string	Status of this update check.
update_check	string	Name of the update check to be validated.

software_reference

Name	Type	Description
_links	_links	
action	string	User triggered action to apply to the install operation
elapsed_duration	integer	Elapsed time during the upgrade or validation operation
estimated_duration	integer	Estimated time remaining until completion of the upgrade or validation operation.
metrocluster	metrocluster	
nodes	array[software_node_reference]	List of nodes and active versions.

Name	Type	Description
pending_version	string	Version being installed on the system. <ul style="list-style-type: none"> • example: ONTAP_X_1 • readOnly: 1
state	string	Operational state of the upgrade
status_details	array[software_status_details_reference]	Display status details.
update_details	array[software_update_details_reference]	Display update proccess details.
validation_results	array[software_validation_reference]	List of validation warnings, errors, and advice.
version	string	Version of ONTAP installed and currently active on the system. During PATCH, using the 'validate_only' parameter on the request executes pre-checks, but does not perform the full installation. <ul style="list-style-type: none"> • example: ONTAP_X

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Download a software or firmware package

POST `/cluster/software/download`

Downloads a software package from the server.

Required properties

- `url` - URL location of the software package

Recommended optional parameters

- `username` - Username of HTTPS/FTP server
- `password` - Password of HTTPS/FTP server

Related ONTAP commands

- `cluster image package get`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

Request Body

Name	Type	Description
password	string	Password for download
url	string	HTTP or FTP URL of the package via a server
username	string	Username for download

Example request

```
{
  "password": "admin_password",
  "url": "http://server/package",
  "username": "admin"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

software_package_download

Name	Type	Description
password	string	Password for download
url	string	HTTP or FTP URL of the package via a server
username	string	Username for download

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve the software installation request history details

GET `/cluster/software/history`

Retrieves the history details for software installation requests.

Related ONTAP commands

- `cluster image show-update-history`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.

Name	Type	In	Required	Description
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	
records	array[software_history_reference]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "end_time": "2019-02-02 20:00:00 UTC",
    "from_version": "ONTAP_X1",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "start_time": "2019-02-02 19:00:00 UTC",
    "state": "successful",
    "to_version": "ONTAP_X2"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

software_history_reference

Name	Type	Description
end_time	string	Completion time of this installation request.
from_version	string	Previous version of node <ul style="list-style-type: none">• example: ONTAP_X1• readOnly: 1
node	node	
start_time	string	Start time of this installation request.
state	string	Status of this installation request.

Name	Type	Description
to_version	string	Updated version of node <ul style="list-style-type: none"> • example: ONTAP_X2 • readOnly: 1

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve cluster software packages

GET /cluster/software/packages

Retrieves the software packages for a cluster.

Related ONTAP commands

- `cluster image package show-repository`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned.
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	
records	array[software_package]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "create_time": "2019-02-04 19:00:00 UTC",
    "version": "ONTAP_X"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

software_package

Name	Type	Description
_links	_links	
create_time	string	Indicates when this package was loaded
version	string	Version of this package <ul style="list-style-type: none">• example: ONTAP_X• readOnly: 1

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code

Name	Type	Description
message	string	Error message
target	string	The target parameter that caused the error.

Delete a software package from the cluster

DELETE /cluster/software/packages/{version}

Deletes a software package from the cluster. The delete operation fails if the package is currently installed.

Related ONTAP commands

- `cluster image package delete`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
version	string	path	True	

Response

Status: 202, Accepted

Error

Status: Default

ONTAP Error Response codes

Error codes	Description
10551315	Package store is empty
10551322	Error in retrieving package cleanup status
10551323	Error in cleaning up package information on a node
10551324	Error in cleaning up package information on multiple nodes
10551325	Package does not exist on the system

Error codes	Description
10551326	Error in deleting older package cleanup tasks. Clean up images from the store and retry
10551346	Package delete failed since a validation is in progress
10551347	Package delete failed since an update is in progress
10551367	A package synchronization is in progress
10551388	Package delete operation timed out

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve the software package information

GET /cluster/software/packages/{version}

Retrieves the software package information.

Related ONTAP commands

- `cluster image package show-repository`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
version	string	path	True	
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
create_time	string	Indicates when this package was loaded
version	string	Version of this package <ul style="list-style-type: none">• example: ONTAP_X• readOnly: 1

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "create_time": "2019-02-04 19:00:00 UTC",
  "version": "ONTAP_X"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": {
      "code": "string",
      "message": "string"
    },
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

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.