IBM Spectrum NAS
Version 1.7.1.0

REST API Management Guide

SC27-9227-01

Introduction

This document describes the Management REST API for IBM Spectrum NAS, the types of operations it supports as well as how authentication and authorization are handled.

This edition applies to IBM Spectrum NAS, Version 1.7.1.0, and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Overview

The Management REST API makes it possible to programmatically manage an IBM Spectrum NAS file system. The API supports administration of file shares, users, groups and policies; it also supports manual snapshots. Requests are executed in the context of a single domain / file system.

Getting started

The storage cluster must have at least one file system configured. Before the REST API can be used, you will need to do the following in the Management tool (<domain> is the name of the domain / file system you want to configure):

1. Go to the **cluster** > **File System** tab > **<domain>** and enable **REST API**.
2. Go to the **cluster** > **File System** tab > **<domain>** > **Users** and create a user.
3. Go to a **node** > **Config** tab > **Gateway** and enable **REST API** (you will find this under *protocols*). Do this for any node that is going to be used for API communication.

Once the Web management service has been enabled and a user has been created, you are ready to start using the REST API. API requests are sent by using the HTTP protocol on port 81. Note that requests should be sent to a public IP address of the file system, or a DNS name that resolves to a public IP address. The user must also be defined in the same file system.

Example with Curl on a Linux machine (replace <username>, <password> and <node> with your own values):

```
$ CV_AUTH="$(echo -n <username>:<password> | base64)"
$ curl --head --header "Authorization: Basic $CV_AUTH" http://<node>:81/api/v1
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 09:52:42 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: ce1efac3-20ea-4a36-93f1-f0941ee2a650
```

In this example, a simple HEAD request is sent to the root URI. This is useful for checking that you are able to reach the REST API and that the user was successfully authenticated. No special permissions are needed for this request.

Authentication

IBM Spectrum NAS uses HTTP Basic authentication. To authenticate as a user, you must first generate a base64 encoding of the string <username>:<password> (replace <username> and <password> with one of your users). Note that ':' must be included in the string. Once you have the base64 encoded string, you set the HTTP Authorization header in the following way:

Authorization: Basic <base64_encoded_string>

Access control

There are two types of permissions to control access: management permissions on users and groups, and permissions on directories and files. User and group permissions determine the operations a user is allowed to perform – they are always checked before directory and file permissions. Directory and file permissions, or Access Control List (ACL), determine the operations a user is allowed to perform on a particular file system object.

User and group permissions

The following screenshot shows the user and group permissions. They can be found in the Management Tool: **cluster** > **File System** tab > **<domain>** > **Users/Groups** > **<User/Group name>.**

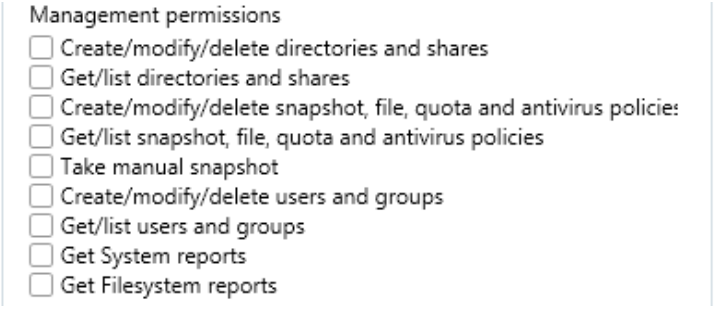

A user's effective management permissions is a combination of its own permissions and the permissions of all the groups it's a member of. If a user does not have the correct management permissions, then it will be denied access.

ACL permissions

For all requests that operate on the file system tree, ACL is checked before granting access. ACL can be modified using any of the supported protocols, or by specifying a new mode with the Modify Directory operation.

The following table lists the ACL permissions required by the operations. The Linux Mode column shows what the permissions correspond to in Linux. The Parent column means that the check is performed using the parent directory's ACL.

Operation	ACL Permissions	Linux Mode	Parent
Create Directory	Add subdirectory	2 (w)	Yes
Modify Directory	Only owner is able to change mode Read extended attributes on source parent directory Add subdirectory on target parent directory Source parent directory and target parent directory are checked when renaming the directory.	4 (r) on source and 2 (w) on target	-
Delete Directory	Delete subdirectory	2 (w)	Yes
Head Inode	Read extended attributes	2 (r)	Yes
Get Inode	Read extended attributes	2 (r)	Yes
List Directories and Files	List directory Read extended attributes	2 (r)	No
Create Policy	Write extended attributes	2 (w)	No
Modify Policy	Write extended attributes	2 (w)	No
Delete Policy	Write extended attributes	2 (w)	No
Get Policy	Read extended attributes	4 (r)	No
List Policies	Read extended attributes on file system root	4 (r)	-
Create Share	Write extended attributes	2 (w)	No
Modify Share	Write extended attributes	2 (w)	No
Delete Share	Write extended attributes	2 (w)	No
Get Share	Read extended attributes	2 (r)	No
List Shares	Read extended attributes on file system root	2 (r)	-
Take Snapshot	Write extended attributes	2 (w)	No
Get System reports	-	-	-
Get Filesystem reports	-	-	-

Note: Administrator, and members of the Administrators group, are exempted from the ACL check.

Current Working Directory

Every operation that takes a path as an argument, either in the path part of the URI, or in the request body, is executed relative to a current working directory (cwd). You specify a cwd by using a directory's unique ID in the x-cv-cwd-id header. The x-cv-cwd-id header is optional and the request is executed relative to the file system root if it's omitted. You can retrieve the unique ID of a directory by using the Head Inode operation or the Get Inode operation. The ID is also included in the result for an entry in the List Directories and Files operation.

Note that when the path is part of the request body, it's relative to the cwd only if it doesn't begin with a '/'. If it begins with a '/', the path is relative to the file system root. Paths that are included in the URI are always relative to the cwd.

File system Management

Operations for file system management. Create, rename and delete a directory; get an inode, and list files and directories.

Directories and Files

The Inode object


```
{
  "id": "8e82ba51-0000-0000-1956-56df2b153a43",
  "type": "directory",
  "name": "mydirectory",
  "createDate": "2016-10-14 12:44:21",
  "modifyDate": "2016-10-14 12:44:21",
  "accessDate": "2016-10-14 12:44:21",
  "metadataDate": "2016-10-14 12:44:21",
  "scanDate": "2016-10-14 12:44:21",
  "backupDate": "2000-01-01 00:00:00",
  "snapshotDate": "2016-10-14 12:44:21",
  "mode": "777",
  "size": 0,
  "usedSize": 0,
  "diskSize": 0,
  "directoryCount": 0,
  "fileCount": 0,
  "readBytes": 0,
  "writeBytes": 0,
  "readOperations": 0,
  "writeOperations": 0
}
```


Attribute Definitions

Name	Type	Description
id	String	The unique ID of the file / directory.
type	String	The type of inode. Valid values: directory, file, symlink, link
name	String	The name of the file / directory.
createDate	String	The date and time the file / directory was created.
modifyDate	String	The date and time the file / directory was last modified.
accessDate	String	The date and time the file / directory was last accessed.
metadataDate	String	The date and time the file / directory metadata was last modified.
scanDate	String	The date and time the file / directory was last scanned by the antivirus. Note: The value "2000-01-01 00:00:00" indicates that the item has never been scanned.
backupDate	String	The date and time the file / directory last had a backup taken. Note: The value "2000-01-01 00:00:00" indicates that the item has never been backed up.
snapshotDate	String	The date and time the file / directory last had a snapshot taken. Note: The value "2000-01-01 00:00:00" indicates that a snapshot has never been taken on the item.

mode	String	The Linux permissions of the file / directory.
size	Integer	For a file, this attribute indicates the total size, in bytes, of the file. For a directory, this attribute indicates the sum of the “size” values for each file in the directory and all its subdirectories. If the directory has a .snapshot subfolder, the “size” of the .snapshot subfolder is not included in the value of this attribute.
usedSize	Integer	For a file, this attribute indicates the actual size, in bytes, of the data in the file (not- written areas in the file excluded). For a directory, this attribute indicates the sum of the “usedSize” values for each file in the directory and all its subdirectories. Note: due to thin-provisioning, “usedSize” is usually less than “size”. If the directory has a .snapshot subfolder, the “usedSize” of the .snapshot subfolder is not included in the value of this attribute.
diskSize	Integer	For a file, this attribute indicates the actual size, in bytes, of the file (not-written areas in the file excluded, erasure coding data included) and of the metadata associated with the file, as stored on the storage disk. For a directory, this attribute indicates the sum of the “diskSize” values for each file in the directory and all its subdirectories. If the directory has a .snapshot subfolder, the “diskSize” of the .snapshot subfolder is not included in the value of this attribute.
directoryCount	Integer	The number of subdirectory entries in a directory, including the directory itself. A directory without any subdirectories has the value 1 for this attribute. If the inode is a file, the value of the attribute is 0.
fileCount	Integer	The number of file entries in the directory. If the inode is a file, the value of the attribute is 0.
readBytes	Integer	The number of bytes that have been read from the file / directory (including all its contents, subdirectories and files), since its creation.
writeBytes	Integer	The number of bytes that have been written to the file / directory (including all its contents, subdirectories and files), since its creation.
readOperations	Integer	The number of read operations on the file / directory (including all its contents, subdirectories and files), since its creation.
writeOperations	Integer	The number of write operations on the file / directory (including all its contents, subdirectories and files), since its creation.

Create Directory

POST */api/v1/fs?type=directory*

Create a directory at the path specified in the request body.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Response Headers

x-cv-inode-id	The ID of the newly created directory.
x-cv-parent-id	The ID of the parent directory.

Request Body

Name	Type	Description	Status
createParents	Bool	If true, parent directories in the path are created as needed. If false, only the last part of the path is created.	Optional, default false.
path	String	The path to the new directory, relative to the current working directory.	Required
mode	String	The Linux permissions. This can be used to restrict access to directories. If createParents is true, all directories will be created with the same mode.	Optional, default “775”.

Sample Request

```
POST /api/v1/fs?type=directory HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "createParents": true,
  "path": "/path/to/directory",
  "mode": "744"
}
```

Sample Response

```
HTTP/1.1 201 Created
Date: Fri, 14 Oct 2016 10:51:57 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
Location: http://<node>:81/api/v1/fs/path/to/directory
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
x-cv-parent-id: 3b100000-0000-0000-1956-56df2b153a43
x-cv-inode-id: 3e4a8ed1-0000-0000-1956-56df2b153a43
```

Modify Directory

PUT */api/v1/fs/<path>?type=directory*

Change name or mode of a directory. Only the owner of the directory is allowed to change mode.

Path Parameters

path	The path of the directory, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Response Headers

x-cv-inode-id	The ID of the directory.
x-cv-parent-id	The ID of the parent directory.

Request Body

Same as for Create Directory, except createParents is not used.

Sample Request

```
PUT /api/v1/fs/path/to/directory?type=directory HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "path": "new_name",
  "mode": "700"
}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 11:50:56 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
x-cv-parent-id: 3b100000-0000-0000-1956-56df2b153a43
x-cv-inode-id: 3e4a8ed1-0000-0000-1956-56df2b153a43
```

Delete Directory

DELETE */api/v1/fs/<path>?type=directory*

Delete a directory. The directory must be empty for this operation to succeed.

Path Parameters

path	The path of the directory, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Response Headers

x-cv-parent-id	The ID of the parent directory.

Sample Request

```
DELETE /api/v1/fs/path/to/directory?type=directory HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 204 No Content
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
x-cv-parent-id: 3b100000-0000-0000-1956-56df2b153a43
```

Head Inode

HEAD /api/v1/fs/<path>

Check the existence of a file or directory.

Path Parameters

path	The path of the directory / file, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Response Headers

x-cv-inode-id	The ID of the inode.
x-cv-inode-type	The type of the inode.
x-cv-parent-id	The ID of the parent directory.

Sample Request

```
HEAD /api/v1/fs/path/to/directory HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:05:32 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: c69a0b0c-236c-4372-8fc6-2caaf85551b8
x-cv-parent-id: 3b100000-0000-0000-1956-56df2b153a43
x-cv-inode-id: 3e4a8ed1-0000-0000-1956-56df2b153a43
x-cv-inode-type: directory
```

Get Inode / Statistics

GET */api/v1/fs/<path>*

Retrieve inode attributes.

Path Parameters

path	The path of the directory / file, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Response Headers

x-cv-parent-id	The ID of the parent directory.

Response Body

The Inode object.

Sample Request

```
GET /api/v1/fs/path/to/directory HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:05:32 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: c69a0b0c-236c-4372-8fc6-2caaf85551b8
x-cv-parent-id: 3b100000-0000-0000-1956-56df2b153a43

{
  "id": "8e82ba51-0000-0000-1956-56df2b153a43",
  "type": "directory",
  "name": "mydirectory",
  "createDate": "2016-10-14 12:44:21",
  "modifyDate": "2016-10-14 12:44:21",
  "accessDate": "2016-10-14 12:44:21",
  "metadataDate": "2016-10-14 12:44:21",
  "scanDate": "2016-10-14 12:44:21",
  "backupDate": "2016-10-14 12:44:21",
  "snapshotDate": "2016-10-14 12:44:21",
  "mode": "777",
  "size": 0,
  "usedSize": 0,
  "diskSize": 0,
  "directoryCount": 0,
  "fileCount": 0,
  "readBytes": 0,
  "writeBytes": 0,
  "readOperations": 0,
  "writeOperations": 0
}
```

List Directories and Files

GET */api/v1/fs/<path>?list*

List the content of a directory.

Path Parameters

path	The path of the directory, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Response Headers

x-cv-inode-id	The ID of the directory.
x-cv-parent-id	The ID of the parent directory.
x-cv-directory-count	The number of entries in the directory.

Response Body

A JSON array of Inode objects. Optional attributes are not included.

Sample Request

```
GET /api/v1/fs/path/to/directory?list HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:05:32 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: c69a0b0c-236c-4372-8fc6-2caaf85551b8
x-cv-parent-id: 3b100000-0000-0000-1956-56df2b153a43
x-cv-inode-id: 3e4a8ed1-0000-0000-1956-56df2b153a43
x-cv-directory-count: 5

[
  {
    "id": "8e82ba51-0000-0000-1956-56df2b153a43",
    "type": "directory",
    "name": "mydirectory",
    "createDate": "2016-10-14 12:44:21",
    "modifyDate": "2016-10-14 12:44:21",
    "accessDate": "2016-10-14 12:44:21",
    "metadataDate": "2016-10-14 12:44:21",
    "scanDate": "2016-10-14 12:44:21",
    "backupDate": "2016-10-14 12:44:21",
    "snapshotDate": "2016-10-14 12:44:21",
    "mode": "777",
    "size": 0,
    "usedSize": 0,
    "diskSize": 0
  },
  ...
]
```

Policy Management

Operations for policy management. Create, modify, delete, and get a file/quota/snapshot policy; list file/quota/snapshot policies. A policy is created by configuring an existing directory.

File Policies

The File Policy object

```
{
  "filters": {
    "pattern": {
      "enable": true,
      "value": "*.doc"
    },
    "age": {
      "enable": true,
      "count": 2,
      "type": "months"
    }
  },
  "actions": {
    "fileCoding": {
      "enable": true,
      "value": "ERASURE_2_1"
    },
    "tier": {
      "enable": true,
      "value": "TIER_0"
    },
    "encryption": {
      "enable": true
    },
    "retention": {
      "enable": true,
      "modifiedAgo": {
        "count": 2 ,
        "type": "days"
      },
      "accessedAgo": {
        "count": 1,
        "type": "days"
      }
    },
    "worm": {
      "enable": true,
      "count": 2,
      "type": "hours"
    }
  }
}
```

Attribute Definitions

Name	Type	Description
filters	Object	A set of criteria that determine which files will be affected by the actions.
filters.pattern	Object	A filter for one or several file name patterns.
filters.pattern.enable	Bool	Flag that signals if the pattern filter is enabled (true) or disabled (false).
filters.pattern.value	String	The file name patterns, separated by semicolon, which determine which files will be affected by the actions (Ex: *.jpg; *.gif).
filters.age	Object	A filter for the files that are older than the specified age.
filters.age.enable	Bool	Flag that signals if the age filter is enabled (true) or disabled (false).
filters.age.count	Integer	The unit count for the age filter. Valid values: 1-30.
filters.age.type	String	The unit type for the age filter. Valid values: “days”, “weeks”, “months, “years”
actions	Object	A set of events triggered on files that match all enabled filters.
actions.fileCoding	Object	An action that will apply a certain file encoding to the files that match all enabled filters.
actions.fileCoding .enable	Bool	Flag that signals if the file encoding action is enabled (true) or disabled (false).
actions.fileCoding .value	String	The file encoding argument (Ex: COPIES_3, ERASURE_2_1 etc.). See the appendix “File encodings” for all valid values.
actions.tier	Object	An action that will move files that match all enabled filters to nodes belonging to a certain tier (tiers define different importance levels of the data).
actions.tier.enable	Bool	Flag that signals if the tier action is enabled (true) or disabled (false).
actions.tier.value	String	The tier argument (Ex: TIER_0, TIER_1 etc.). See the appendix “Tiers” for all valid file encodings.
actions.encryption	Object	An action that will encrypt files that match all enabled filters.
actions.encryption .enable	Bool	Flag that signals if the encryption action is enabled (true) or disabled (false).
actions.retention	Object	An action that will delete files that have not been accessed / modified in the retention period. Files that have been modified / accessed at least once since the specified time (for example, in the last 2 months) are kept; the remaining files are deleted. At least one of modifiedAgo and accessedAgo must be specified (have a positive value). If both are specified, only files that do not match any of the two retention periods will be deleted.
actions.retention .enable	Bool	Flag that signals if the retention action is enabled (true) or disabled (false).

actions.retention .modifiedAgo	Object	An attribute for the retention action. Once given a value, only files that have been modified since the specified time are retained, and the rest are deleted (unless they are retained by a different retention attribute, i.e. “accessedAgo”).
actions.retention .modifiedAgo.count	Integer	The unit count for the modifiedAgo attribute. Valid values: 0-30. A value of 0 means the modifiedAgo attribute is disabled.
actions.retention .modifiedAgo.type	String	The unit type for the modifiedAgo attribute. Valid values: “days”, “weeks”, “months”, “years”
actions.retention .accessedAgo	Object	An attribute for the retention action. Once given a value, only files that have been accessed since the specified time are retained, and the rest are deleted (unless they are retained by a different retention attribute, i.e. “modifiedAgo”).
actions.retention .accessedAgo.count	Integer	The unit count for the accessedAgo attribute. Valid values: 0-30. A value of 0 means the accessedAgo attribute is disabled.
actions.retention .accessedAgo.type	String	The unit type for the accesedAgo attribute. Valid values: “days”, “weeks”, “months”, “years”
actions.worm	Object	An action (name stands for Write Once Read Many) that will permanently mark as read-only the files that are older than the specified period.
actions.worm.enable	Bool	Flag that signals if the worm action is enabled (true) or disabled (false).
actions.worm.count	Integer	The unit count for the worm period. Valid values: > 0
actions.worm.type	String	The unit type for the worm period. Valid values: “seconds”, “minutes”, “hours”, “days”, “weeks”, “months”, “years”

Create File Policy

POST */api/v1/fs/<path>?policy=file*

Add a file policy to a directory.

Path Parameters

path	The path of the directory the policy should be added to, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Request Body

The File Policy object. At least one action must be specified.

Sample Request

```
POST /api/v1/fs/path/to/directory?policy=file HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "filters": {
    "pattern": {
      "enable": true,
      "value": "*.doc"
    },
    "age": {
      "enable": true,
      "count": 2,
      "type": "months"
    }
  },
  "actions": {
    "fileCoding": {
      "enable": true,
      "value": "ERASURE_2_1"
    },
    "tier": {
      "enable": true,
      "value": "TIER_0"
    },
    "encryption": {
      "enable": true
    },
    "retention": {
      "enable": true,
      "modifiedAgo": {
        "count": 2 ,
        "type": "days"
      },
      "accessedAgo": {
        "count": 1,
        "type": "days"
      }
    },
    "worm": {
      "enable": true,
      "count": 2,
      "type": "hours"
    }
  }
}
```

Sample Response

```
HTTP/1.1 201 Created
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Modify File Policy

PUT /api/v1/fs/<path>?policy=file

Modify a directory's file policy.

Path Parameters

path	The path of the directory that has the file policy, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Request Body

The File Policy object. Any parameters not provided remains unchanged.

Sample Request

```
PUT /api/v1/fs/path/to/directory?policy=file HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "filters": {
    "pattern": {
      "value": "*.jpg"
    }
  },
  "actions": {
    "retention": {
      "enable": false
    },
    "worm": {
      "enable": false
    }
  }
}
```

This sample request changes the pattern from *.doc to *.jpg, and disables the retention and worm action. The other file policy configuration remains unchanged.

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Delete File Policy

DELETE /api/v1/fs/<path>?policy=file

Delete a directory's file policy.

Path Parameters

path	The path of the directory that has the file policy, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Sample Request

```
DELETE /api/v1/fs/path/to/directory?policy=file HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 204 No Content
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Get File Policy

GET */api/v1/fs/<path>?policy=file*

Retrieve a directory's file policy configuration.

Path Parameters

path	The path of the directory that has the file policy, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Response Body

The File Policy object.

Sample Request

```
GET /api/v1/fs/path/to/directory?policy=file HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

{
  "filters": {
    "pattern": {
      "enable": true,
      "value": "*.doc"
    },
    "age": {
      "enable": true,
      "count": 2,
      "type": "months"
    }
  },
  "actions": {
    "fileCoding": {
      "enable": true,
      "value": "ERASURE_2_1"
    },
    "tier": {
      "enable": true,
      "value": "TIER_0"
    },
    "encryption": {
      "enable": true
    },
    "retention": {
      "enable": true,
      "modifiedAgo": {
        "count": 2 ,
        "type": "days"
      },
      "accessedAgo": {
        "count": 1,
        "type": "days"
      }
    },
    "worm": {
      "enable": true,
      "count": 2,
      "type": "hours"
    }
  }
}
```

List File Policies

GET */api/v1/fs?policy=file&list*

List all file policies in the file system.

Response Body

Name	Type	Description
path	String	The path of the directory that has the file policy.
filePolicy	Object	The File Policy object.

Sample Request

```
GET /api/v1/fs?policy=file&list HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

[
  {
    "path": "/path/to/directory",
    "filePolicy": {
      "filters": {
        "pattern": {
          "enable": true,
          "value": "*.doc"
        },
        "age": {
          "enable": true,
          "count": 2,
          "type": "months"
        }
      },
      "actions": {
        "fileCoding": {
          "enable": true,
          "value": "ERASURE_2_1"
        },
        "tier": {
          "enable": true,
          "value": "TIER_0"
        },
        "encryption": {
          "enable": true
        },
        "retention": {
          "enable": true,
          "modifiedAgo": {
            "count": 2 ,
            "type": "days"
          },
          "accessedAgo": {
            "count": 1,
            "type": "days"
          }
        },
        "worm": {
          "enable": true,
          "count": 2,
          "type": "hours"
        }
      }
    }
  },
  ...
]
```

Quota Policies

The Quota Policy object

```
{
  "limit": {
    "count": 2,
    "type": "GB"
  }
}
```

Attribute Definitions

Name	Type	Description
limit	Object	The quota limit for the directory (including all subdirectories and files). The total size of the data contained in the directory cannot normally surpass this limit; however, in case of multithreaded parallel write operations, there is a theoretical possibility that the limit is slightly exceeded due to the delay required by the system to compute the total size of the parallel write operations. As soon as the system has realized that the limit has been reached, no more write operations will be allowed on the directory (until some data is deleted and the directory size goes below the quota limit). Note: It is a valid operation to apply a quota limit on a directory that already contains more data than the applied limit. Any new write attempt to the directory will however be denied, until the total size of the data in the directory goes below the quota limit (by deleting some of the already existing data).
limit.count	String	The unit count for the limit. Valid values: 1-999.
limit.type	String	The unit type for the limit. Valid values: “GB”, “TB”, “PB”

Create Quota Policy

POST */api/v1/fs/<path>?policy=quota*

Add a quota policy to a directory.

Path Parameters

path	The path of the directory the policy should be added to, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Request Body

The Quota Policy object.

Sample Request

```
POST /api/v1/fs/path/to/directory?policy=quota HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "limit": {
    "count": 2,
    "type": "GB"
  }
}
```

Sample Response

```
HTTP/1.1 201 Created
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Modify Quota Policy

PUT */api/v1/fs/<path>?policy=quota*

Modify a directory's quota policy.

Path Parameters

path	The path of the directory that has the quota policy, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Request Body

The Quota Policy object. Any parameters not provided remains unchanged.

Sample Request

```
PUT /api/v1/fs/path/to/directory?policy=quota HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "limit": {
    "count": 5,
    "type": "TB"
  }
}
```

This sample request changes the quota limit from 2 GB to 5 TB. The other quota policy configuration remains unchanged.

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Delete Quota Policy

DELETE */api/v1/fs/<path>?policy=quota*

Delete a directory's quota policy.

Path Parameters

path	The path of the directory that has the quota policy, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Sample Request

```
DELETE /api/v1/fs/path/to/directory?policy=quota HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 204 No Content
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Get Quota Policy

GET */api/v1/fs/<path>?policy=quota*

Retrieve a directory's quota policy configuration.

Path Parameters

path	The path of the directory that has the quota policy, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Response Body

The Quota Policy object.

Sample Request

```
GET /api/v1/fs/path/to/directory?policy=quota HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

{
  "limit": {
    "count": 2,
    "type": "GB"
  }
}
```

List Quota Policies

GET */api/v1/fs?policy=quota&list*

List all quota policies in the file system.

Response Body

Name	Type	Description
path	String	The path of the directory that has the quota policy.
quotaPolicy	Object	The Quota Policy object.

Sample Request

```
GET /api/v1/fs?policy=quota&list HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

[
  {
    "path": "/path/to/directory",
    "quotaPolicy": {
      "limit": {
        "count": 2,
        "type": "GB"
      }
    }
  },
  ...
]
```

Snapshot Policies

The Snapshot Schedule object

```
{
  "type": "weekly",
  "numberOfCopies": 5,
  "timeOfDay": "13:56",
  "dayOfWeek": "monday",
  "hoursOfDay": [1, 12, 23]
}
```

Attribute Definitions

Name	Type	Description
type	String	The type of snapshot schedule. The type determines how the numberOfCopies, timeOfDay, dayOfWeek and hoursOfDay attributes are used. Note that the schedule type “manually” is not required for snapshots to be triggered manually. See actions: Take Snapshot. Valid values: “manually”, “hourly”, “daily”, “weekly”. The type “manually” uses numberOfCopies. The type “hourly” uses numberOfCopies and hoursOfDay. The type “daily” uses numberOfCopies and timeOfDay. The type “weekly” uses numberOfCopies , timeOfDay and dayOfWeek.
numberOfCopies	Integer	The maximum number of snapshot copies that can be retained at any point in time. For example, if numberOfCopies is 5, the sixth snapshot copy will overwrite the oldest of the previous five, thus keeping the total number of snapshot copies at a maximum of 5. The total maximum allowed number of copies for all Snapshot Schedules that belong to the same Snapshot Policy is 253.
timeOfDay	Integer	The time of the day the snapshot will be performed. Time is interpreted as UTC offset +00. Format: “hh:mm”. “hh” must be in the range 00 – 23 and “mm” must be in the range 00 – 59.
dayOfWeek	String	The day of the week the snapshot should be performed. Valid values: “monday”, “tuesday”, “wednesday”, “thursday”, “friday”, “saturday”, “sunday”
hoursOfDay	Array of integers	The hours of the day a snapshot should be performed. Hours are interpreted as UTC offset +00. Format: [h1, h2, h3, ...]. Valid values: Each entry must be in the range 0 – 23. Duplicates are ignored.

The Snapshot Policy object

An array of Snapshot Schedule objects, each having a distinct type (i.e. weekly, hourly, daily or manually). The sum of all numberOfCopies values of a Snapshot Policy cannot exceed 253.

```
[
 {
  "type": "weekly",
  "numberOfCopies": 5,
  "timeOfDay": "00:00",
  "dayOfWeek": "monday",
},
 {
  "type": "daily",
  "numberOfCopies": 7,
  "timeOfDay": "13:56",
 },
 {
  "type": "hourly",
  "numberOfCopies": 10,
  "hoursOfDay": [1,13]
 },
 ...
]
```

Create Snapshot Policy

POST */api/v1/fs/<path>?policy=snapshot*

Add a snapshot policy to a directory.

Path Parameters

path	The path of the directory the policy should be added to, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Request Body

The Snapshot Policy object.

Sample Request

```
POST /api/v1/fs/path/to/directory?policy=snapshot HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

[
 {
  "type": "weekly",
  "numberOfCopies": 8,
  "timeOfDay": "18:00",
  "dayOfWeek": "friday"
 },
 {
  "type": "hourly",
  "numberOfCopies": 10,
  "hoursOfDay": [1,13]
 }
]
```

Sample Response

```
HTTP/1.1 201 Created
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Modify Snapshot Policy

PUT */api/v1/fs/<path>?policy=snapshot*

Modify a directory's snapshot policy. If, as a consequence of the modification, the sum of all numberOfCopies of the Snapshot Policy would exceed 253, the modification request will be denied.

Path Parameters

path	The path of the directory that has the snapshot policy, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Request Body

The Snapshot Policy object. Any snapshot schedule or snapshot schedule parameter not provided, remains unchanged.

Sample Request

```
PUT /api/v1/fs/path/to/directory?policy=snapshot HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>
[
 {
  "type": "hourly",
  "numberOfCopies": 20
 }
]
```

This sample request changes number of copies for the hourly schedule to 20. The other snapshot policy configuration remains unchanged.

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Delete Snapshot Policy

DELETE /api/v1/fs/<path>?policy=snapshot

Delete a directory's snapshot policy.

Path Parameters

path	The path of the directory that has the snapshot policy, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Sample Request

```
DELETE /api/v1/fs/path/to/directory?policy=snapshot HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 204 No Content
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Get Snapshot Policy

GET /api/v1/fs/<path>?policy=snapshot

Retrieve a directory's snapshot policy configuration.

Path Parameters

path	The path of the directory that has the snapshot policy, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Response Body

The Snapshot Policy object.

Sample Request

```
GET /api/v1/fs/path/to/directory?policy=snapshot HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

[
 {
  "type": "weekly",
  "numberOfCopies": 8,
  "timeOfDay": "18:00",
  "dayOfWeek": "friday"
 },
 {
  "type": "hourly",
  "numberOfCopies": 20,
  "hoursOfDay": [1,13]
 }
]
```

List Snapshot Policies

GET */api/v1/fs?policy=snapshot&list*

List all snapshot policies in the file system.

Response Body

Name	Type	Description
path	String	The path of the directory that has the snapshot policy.
snapshotPolicy	Object	The Snapshot Policy object.

Sample Request

```
GET /api/v1/fs?policy=snapshot&list HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

[
  {
    "path": "/path/to/directory",
    "snapshotPolicy": [
     {
      "type": "weekly",
      "numberOfCopies": 5,
      "timeOfDay": "13:56",
      "dayOfWeek": "monday"
     },
     ...
    ]
  },
  ...
]
```

Antivirus Policies

The Antivirus Policy object

```
{
  "includePattern": "*",
  "excludePattern": "*.zip;*.exe;*.pdf",
  "maxFileSize": {
    "count": 2,
    "type": "GB"
  },
  "alerts": {
    "alertOnWarnings": true,
    "email": true,
    "snmp": true
  },
  "options": {
    "deleteInfectedFiles": false
  },
  "backgroundScan": {
    "enable": true,
    "maxFileSize": {
      "count": 10,
      "type": "GB"
    },
    "interval": {
      "type": "weekly",
      "timeOfDay": "14:00",
      "dayOfweek": "monday",
      "hoursOfDay": [1, 12, 23]
    }
  }
}
```

Attribute Definitions

Name	Type	Description
includePattern	String	The file name patterns, semicolon separated, for files that will be included in live scanning (at read time) and background scanning (example, "*.exe; *.zip"). If includePattern is empty, all file patterns are included.
excludePattern	String	The file name patterns, semicolon separated, for files that will be excluded from live scanning (at read time) and background scanning (E.g. “*.zip; *.pdf).
maxFileSize	Object	Files having this maximum size will be included in the live scanning (scanning at read time).
maxFileSize.count	Integer	The unit count for the maximum size of files that will be included in the live scanning (scanning at read time). Valid values: 1-999
maxFileSize.type	String	The unit type for the maximum size of files that will be included in the live scanning (scanning at read time). Valid values: “KB”, “MB”, “GB”
alerts	Object	A feature that alerts the user via email and/or SNMP traps on various antivirus-related events, e.g. when infected files are found (critical-level event), or when suspicious content is found (warning-level event). Note: the warning-level events must be separately enabled, see the “alerts.alertOnWarnings” attribute below.
alerts.alertOnWarnings	Bool	A feature that alerts the user even on warning-level events, e.g. suspicious content. If this attribute is disabled, and at least one of “alerts.email” or “alerts.snmp” is enabled, alerts will be sent only on critical-level events, i.e. when infected files are found.
alerts.email	Bool	Flag that signals if the email alerts are enabled (true) or disabled (false).
alerts.snmp	Bool	Flag that signals if the SNMP trap alerts are enabled (true) or disabled (false).
options.deleteInfectedFiles	Bool	If enabled, files will be automatically deleted if found infected. **Warning**: By using this option, there is a risk of unintentional loss of data.
backgroundScan	Object	A feature that does background antivirus scanning.
backgroundScan.enable	Bool	Flag that signals if the background scanning is enabled (true) or disabled (false).
backgroundScan.maxFileSize	Object	Files having this maximum size will be included in the background scanning.
backgroundScan.maxFileSize.count	Integer	The unit count for the maximum size of files that will be included in the background scanning. Valid values: 1-999.

backgroundScan.maxFileSize .type	String	The unit type for the maximum size of files that will be included in the background scanning. Valid values: “KB”, “MB”, “GB”
backgroundScan.interval	Object	Attributes to define intervals at which the background scanning is performed. If backgroundScan.enable is true, then this interval must be included.
backgroundScan.interval .type	String	The type of interval for the background scanning. The type determines how the timeOfDay, dayOfWeek and hoursOfDay attributes are used. Valid values: “hourly”, “daily”, “weekly” The type “hourly” uses hoursOfDay. The type “daily” uses timeOfDay. The type “weekly” uses timeOfDay and dayOfWeek.
backgroundScan.interval .timeOfDay	Integer	The time of day the background scan should be performed. Time is interpreted as UTC offset +00. Format: “hh:00”. hh must be in the range 00 – 23.
backgroundScan.interval .dayOfweek	String	The day of the week the background scan should be performed. Valid values: “monday”, “tuesday”, “wednesday”, “thursday”, “friday”, “saturday”, “sunday”
backgroundScan.interval .hoursOfDay	Array of integers	The hours of the day a background scan should be performed. Hours are interpreted as UTC offset +00. Format: [h1, h2, h3, ...]. Valid values: Each entry must be in the range 0 – 23. Duplicates are ignored.

Create Antivirus Policy

POST */api/v1/fs/<path>?policy=antivirus*

Add an antivirus policy to a directory.

Path Parameters

path	The path of the directory the policy should be added to, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Request Body

The Antivirus Policy object.

Sample Request

```
POST /api/v1/fs/path/to/directory?policy=antivirus HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "includePattern": "*",
  "excludePattern": "*.zip;*.exe;*.pdf",
  "maxFileSize": {
    "count": 2,
    "type": "GB"
  },
  "alerts": {
    "alertOnWarnings": true,
    "email": true,
    "snmp": true
  },
  "options": {
    "deleteInfectedFiles": false
  },
  "backgroundScan": {
    "enable": true,
    "maxFileSize": {
      "count": 10,
      "type": "GB"
    },
    "interval": {
      "type": "weekly",
      "timeOfDay": "14:00",
      "dayOfweek": "monday",
      "hoursOfDay": [1, 12, 23]
    }
  }
}
```

Sample Response

```
HTTP/1.1 201 Created
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Modify Antivirus Policy

PUT */api/v1/fs/<path>?policy=antivirus*

Modify a directory's antivirus policy.

Path Parameters

path	The path of the directory that has the antivirus policy, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Request Body

The Antivirus Policy object. Any parameters not provided remains unchanged.

Sample Request

```
PUT /api/v1/fs/path/to/directory?policy=antivirus HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "alerts": {
    "alertOnWarnings": false,
  }
}
```

This sample request changes the antivirus warnings to false. The other antivirus policy configuration remains unchanged.

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Delete Antivirus Policy

DELETE /api/v1/fs/<path>?policy=antivirus

Delete a directory's antivirus policy.

Path Parameters

path	The path of the directory that has the antivirus policy, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Sample Request

```
DELETE /api/v1/fs/path/to/directory?policy=antivirus HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 204 No Content
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Get Antivirus Policy

GET */api/v1/fs/<path>?policy=antivirus*

Retrieve a directory's antivirus policy configuration.

Path Parameters

path	The path of the directory that has the antivirus policy, relative to the current working directory.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path is relative to.

Response Body

The Antivirus Policy object.

Sample Request

```
GET /api/v1/fs/path/to/directory?policy=antivirus HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

{
  "includePattern": "*",
  "excludePattern": "*.zip;*.exe;*.pdf",
  "maxFileSize": {
    "count": 2,
    "type": "GB"
  },
  "alerts": {
    "alertOnWarnings": true,
    "email": true,
    "snmp": true
  },
  "options": {
    "deleteInfectedFiles": false
  },
  "backgroundScan": {
    "enable": true,
    "maxFileSize": {
      "count": 10,
      "type": "GB"
    },
    "interval": {
      "type": "weekly",
      "timeOfDay": "14:00",
      "dayOfweek": "monday",
      "hoursOfDay": [1, 12, 23]
    }
  }
}
```

List Antivirus Policies

GET */api/v1/fs?policy=antivirus&list*

List all antivirus policies in the file system.

Response Body

Name	Type	Description
path	String	The path of the directory that has the antivirus policy.
antivirusPolicy	Object	The Antivirus Policy object.

Sample Request

```
GET /api/v1/fs?policy=antivirus&list HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

[
  {
    "path": "/path/to/directory",
    "antivirusPolicy": {
      "includePattern": "*",
      "excludePattern": "*.zip;*.exe;*.pdf",
      "maxFileSize": {
        "count": 2,
        "type": "GB"
      },
      "alerts": {
        "alertOnWarnings": true,
        "email": true,
        "snmp": true
      },
      "options": {
        "deleteInfectedFiles": false
      },
      "backgroundScan": {
        "enable": true,
        "maxFileSize": {
          "count": 10,
          "type": "GB"
        },
        "interval": {
          "type": "weekly",
          "timeOfDay": "14:00",
          "dayOfweek": "monday",
          "hoursOfDay": [1, 12, 23]
        }
      }
    }
  },
  ...
]
```

Share Management

Operations for share management. Create, modify, delete, and get a file share; list file shares. A share is created by configuring an existing directory.

Shares

The Share object

```
{
  "name": "myshare",
  "path": "/path/to/directory",
  "smb": {
    "enable": true,
    "options": {
      "ALLOW_GUEST_ACCESS": true,
      "DISABLE_CLIENT_WRITE_CACHE": true,
      "ENCRYPTION": true,
      "SCALE_OUT": true,
      "ALLOW_OFFLINE_FILE_SUPPORT": true,
      "ACCESS_BASED_ENUMERATION": true,
      "DISTRIBUTED_FILESYSTEM": true
    },
    "ipFilter": "172.16.0.0/16"
  },
  "nfs": {
    "enable": true,
    "options": {
      "ALLOW_ROOT_ACCESS": true,
      "ALLOW_GUEST_ACCESS": true,
      "USE_32BIT_INODES": true,
      "DISABLE_SHARE_MODE_CHECK": true,
      "DISABLE_RW_DELEGATION": true
    },
    "authentication": {
      "sys": true,
      "krb5": true,
      "krb5i": true,
      "krb5p": true
    },
    "ipFilter": "192.168.1.0/24 192.168.2.0/24"
  }
}
```

Attribute Definitions

Name	Type	Description
name	String	The name of the new share.
path	String	The path to the directory that should be enabled as an IBM Spectrum NAS file share. The path is relative to the current working directory. If the path starts with "/", it is absolute (starts at the file system root).
smb	Object	Configuration for the SMB protocol.
smb.enable	Bool	Flag that signals if the new share will be accessible via the SMB protocol.
smb.options	Object	A set of attributes with Boolean values that configure the SMB-access on the new share. Valid attributes are: * ALLOW_GUEST_ACCESS (if true, enables the user to access the file share without user credentials) * DISABLE_CLIENT_WRITE_CACHE (if true, disables write caching on the client side and enables read only cache) * ENCRYPTION (if true, adds additional encryption on the data being sent to a client) * SCALE_OUT (if true, enables support for failover between nodes) * ALLOW_OFFLINE_FILE_SUPPORT (if true, enables the user to modify files in the file share even when the network is disconnected and merges the file changes when the network is available again) * ACCESS_BASED_ENUMERATION (if true, enables user access control based on the criteria present in the "User" section of the Management Tool) * DISTRIBUTED_FILESYSTEM (if true, enables DFS – Distributed file system)
smb.ipFilter	String	If present, only clients having the IP specified by this attribute will be able to connect to the SMB share. Both individual IP addresses and IP+subnet mask combinations can be used. (E.g. a value of 172.16.0.0/16 means that only clients with the IP in the range 172.16.0.0 – 172.16.255.255 will be able to connect to the share). Several IP filters can be included, separated by space.
nfs	Object	Configuration for the NFS protocol.
nfs.enable	Bool	Flag that signals if the new share will be accessible via the NFS protocol.

nfs.options	Object	A set of attributes with Boolean values that configure the NFS-access on the new share. Valid attributes are: * ALLOW_ROOT_ACCESS (if true, remote root users are able to access and change any file on the shared file system. This corresponds to no_root_squash. If false, the remote root user will become anonymous (uid -2, gid -2) and is either allowed or denied depending on the setting of "Allow Guest Access" - see below) * ALLOW_GUEST_ACCESS (if true, enables users that do not provide valid NFS credentials to access the file share. The user becomes anonymous - uid -2, gid -2) * USE_32BIT_INODES (if true, forces IBM Spectrum NAS to use 32-bit inodes internally instead of 64-bit inodes) * DISABLE_SHARE_MODE_CHECK (attribute for Mac only. If true, disables share mode check because Mac will always check if the file is open for read). * DISABLE_RW_DELEGATION (if true, disables read/write delegation, which allows clients to cache data locally for faster updates using less network traffic and improving response time. When enabled, can result in unwanted latency).
nfs.authentication	Object	A set of attributes with Boolean values that enable and control various authentication methods to the NFS share. Valid attributes are: sys (Unix uid), krb5 (Kerberos login), krb5i (Kerberos integrity), krb5p (Kerberos privacy / encryption of all traffic between the client and the server)
nfs.ipFilter	String	If present, only clients having the IP specified by this attribute will be able to connect to the NFS share. Both individual IP addresses and IP+subnet mask combinations can be used. (E.g. a value of 172.16.0.0/16 means that only clients with the IP in the range 172.16.0.0 – 172.16.255.255 will be able to connect to the share). Several IP filters can be included, separated by space.

Create Share

POST /api/v1/shares

Create a new IBM Spectrum NAS file share. Note that the directory specified in the path in the request body must already exist.

x-cv-cwd-id	Optional. The ID of the directory the path in the request body is relative to.

Request Body

The Share object. Name and path are required.

Sample Request

```
POST /api/v1/shares HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "name": "myshare",
  "path": "/path/to/directory",
  "smb": {
    "enable": true,
    "options": {
      "ALLOW_GUEST_ACCESS": true,
      "DISABLE_CLIENT_WRITE_CACHE": true,
      "ENCRYPTION": true,
      "SCALE_OUT": true,
      "ALLOW_OFFLINE_FILE_SUPPORT": true,
      "ACCESS_BASED_ENUMERATION": true,
      "DISTRIBUTED_FILESYSTEM": true
    },
    "ipFilter": "172.16.0.0/16"
  },
  "nfs": {
    "enable": true,
    "options": {
      "ALLOW_ROOT_ACCESS": true,
      "ALLOW_GUEST_ACCESS": true,
      "USE_32BIT_INODES": true,
      "DISABLE_SHARE_MODE_CHECK": true,
      "DISABLE_RW_DELEGATION": true
    },
    "authentication": {
      "sys": true,
      "krb5": true,
      "krb5i": true,
      "krb5p": true
    },
    "ipFilter": "192.168.1.0/24 192.168.2.0/24"
  }
}
```

Sample Response

```
HTTP/1.1 201 Created
Date: Fri, 14 Oct 2016 14:15:12 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
Location: http://<node>:81/api/v1/shares/myshare
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Modify Share

PUT */api/v1/shares/<share>*

Modify a file share. Enable or disable a protocol, change protocol settings, and/or change the name of the share.

Path Parameters

share	The name of the share.

Request Body

The Share object. Path is ignored and Name is optional. The name attribute is the new name of the share. Any parameters not provided remains unchanged.

Sample Request

```
PUT /api/v1/shares/myshare HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "nfs": {
    "authentication": {
      "sys": false
    }
  }
}
```

This sample request disables sys as a security flavor for NFS. The other share and protocol configuration remains unchanged.

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 14:17:33 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Delete Share

DELETE */api/v1/shares/<share>*

Delete a file share. Note that the directory is not deleted.

Path Parameters

share	The name of the share.

Sample Request

```
DELETE /api/v1/shares/myshare HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 204 No Content
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Get Share

GET /api/v1/shares/<share>

Retrieve the configuration for a file share.

Path Parameters

share	The name of the share.

Response Body

The Share object.

Sample Request

```
GET /api/v1/shares/myshare HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

{
  "name": "myshare",
  "path": "/path/to/directory",
  "smb": {
    "enable": true,
    "options": {
      "ALLOW_GUEST_ACCESS": true,
      "DISABLE_CLIENT_WRITE_CACHE": true,
      "ENCRYPTION": true,
      "SCALE_OUT": true,
      "ALLOW_OFFLINE_FILE_SUPPORT": true,
      "ACCESS_BASED_ENUMERATION": true,
      "DISTRIBUTED_FILESYSTEM": true
    },
    "ipFilter": "172.16.0.0/16"
  },
  "nfs": {
    "enable": true,
    "options": {
      "ALLOW_ROOT_ACCESS": true,
      "ALLOW_GUEST_ACCESS": true,
      "USE_32BIT_INODES": true,
      "DISABLE_SHARE_MODE_CHECK": true,
      "DISABLE_RW_DELEGATION": true
    },
    "authentication": {
      "sys": true,
      "krb5": true,
      "krb5i": true,
      "krb5p": true
    },
    "ipFilter": "192.168.1.0/24 192.168.2.0/24"
  }
}
```

List Shares

GET /api/v1/shares

List all file shares in the file system.

Response Body

A JSON array of Share objects.

Sample Request

```
GET /api/v1/shares HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

[
  {
    "name": "myshare",
    "path": "/path/to/directory",
    "smb": {
      "enable": true,
      "options": {
        "ALLOW_GUEST_ACCESS": true,
        "DISABLE_CLIENT_WRITE_CACHE": true,
        "ENCRYPTION": true,
        "SCALE_OUT": true,
        "ALLOW_OFFLINE_FILE_SUPPORT": true,
        "ACCESS_BASED_ENUMERATION": true,
        "DISTRIBUTED_FILESYSTEM": true
      },
      "ipFilter": "172.16.0.0/16"
    },
    "nfs": {
      "enable": true,
      "options": {
        "ALLOW_ROOT_ACCESS": true,
        "ALLOW_GUEST_ACCESS": true,
        "USE_32BIT_INODES": true,
        "DISABLE_SHARE_MODE_CHECK": true,
        "DISABLE_RW_DELEGATION": true
      },
      "authentication": {
        "sys": true,
        "krb5": true,
        "krb5i": true,
        "krb5p": true
      },
      "ipFilter": "192.168.1.0/24 192.168.2.0/24"
    }
  },
  ...
]
```

User & Group Management

Operations for user and group management. Create, modify, delete, and get a user; list existing users. Create, modify, delete, and get a group; list existing groups.

Permissions

Name	Type	Description	Status
SET_SHARE	Bool	Create/modify/delete directories and shares.	Optional, default false
GET_SHARE	Bool	Get/list directories and shares.	Optional, default false
SET_POLICY	Bool	Create/modify/delete snapshot, file and quota policies.	Optional, default false
GET_POLICY	Bool	Get/list snapshot, file and quota policies.	Optional, default false
TAKE_SNAPSHOT	Bool	Take manual snapshot.	Optional, default false
SET_USER_GROUP	Bool	Create/modify/delete users and groups	Optional, default false
GET_USER_GROUP	Bool	Get/list users and groups.	Optional, default false

Users

The User object

```
{
  "name": "myuser",
  "password": "mypassword",
  "enable": true,
  "linuxUID": 1000,
  "windowsSID": "S-1-5-21-1634518050-1170479241-3746985662-1000",
  "mainGroup": "Administrators",
  "additionalGroups": ["users"],
  "permissions": {
    "SET_SHARE": true,
    "GET_SHARE": true,
    "SET_POLICY": true,
    "GET_POLICY": true,
    "TAKE_SNAPSHOT": true,
    "SET_USER_GROUP": true,
    "GET_USER_GROUP": true,
    "GET_SYSTEM_REPORTS": true,
    "GET_FILESYSTEM_REPORTS": true
  }
}
```

Attribute Definitions

Name	Type	Description
name	String	The name of the user. Must be unique in the scope of the file system.
password	String	The password of the user.
enable	Bool	If true, the user is allowed to login.
linuxUID	Integer	The Linux user ID.
windowsSID	String	The Windows Security Identifier.
mainGroup	String	The name of the user's primary group.
additionalGroups	Array	A list of groups the user is a member of.
permissions	Object	A set of attributes with Boolean values that configure the user's access permissions on the file system. Valid attributes are: * SET_SHARE (create/modify/delete directories and shares) * GET_SHARE (get / list directories and shares) * SET_POLICY (create/modify/delete snapshot, file, quota and antivirus policies) * GET_POLICY (get/list snapshot, file, quota and antivirus policies) * TAKE_SNAPSHOT (take manual snapshot) * SET_USER_GROUP (create/modify/delete users and groups) * GET_USER_GROUP (get/list users and groups) * GET_SYSTEM_REPORTS (get System performance and audit reports) * GET_FILESYSTEM_REPORTS (get Filesystem audit reports)

Create User

POST */api/v1/users*

Create a user.

Request Body

The User object. Name, password and mainGroup are required. WindowsSID is ignored.

Sample Request

```
POST /api/v1/users HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "name": "myuser",
  "password": "mypassword",
  "enable": true,
  "linuxUID": 1000,
  "mainGroup": "Administrators",
  "additionalGroups": ["users"],
  "permissions": {
    "SET_SHARE": true,
    "GET_SHARE": true,
    "SET_POLICY": true,
    "GET_POLICY": true,
    "TAKE_SNAPSHOT": true,
    "SET_USER_GROUP": true,
    "GET_USER_GROUP": true,
    "GET_SYSTEM_REPORTS": true,
    "GET_FILESYSTEM_REPORTS": true
  }
}
```

Sample Response

```
HTTP/1.1 201 Created
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
Location: http://<node>:81/api/v1/users/myuser
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Modify User

PUT */api/v1/users/<user>*

Modify an existing user.

Path Parameters

user	The name of the user.

Request Body

The User object. Name, password and mainGroup are optional. Any parameters not provided remains unchanged.

Sample Request

```
PUT /api/v1/users/myuser HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "name": "myuser2",
  "password": "mypassword2"
}
```

This sample request renames the user from myuser to myuser2, and changes the password from mypassword to mypassword2. The other user configuration remains unchanged.

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Delete User

DELETE */api/v1/users/<user>*

Delete a user.

Path Parameters

user	The name of the user.

Sample Request

```
DELETE /api/v1/users/myuser HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 204 No Content
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Get User

GET */api/v1/users/<user>*

Retrieve a user.

Path Parameters

user	The name of the user.

Response Body

The User object.

Sample Request

```
GET /api/v1/users/myuser HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

{
  "name": "myuser",
  "password": "mypassword",
  "enable": true,
  "linuxUID": 1000,
  "windowsSID": "S-1-5-21-1634518050-1170479241-3746985662-1000",
  "mainGroup": "Administrators",
  "additionalGroups": ["users"],
  "permissions": {
    "SET_SHARE": true,
    "GET_SHARE": true,
    "SET_POLICY": true,
    "GET_POLICY": true,
    "TAKE_SNAPSHOT": true,
    "SET_USER_GROUP": true,
    "GET_USER_GROUP": true,
    "GET_SYSTEM_REPORTS": true,
    "GET_FILESYSTEM_REPORTS": true
  }
}
```

List Users

GET */api/v1/users*

List all users.

Response Body

A JSON array of User objects.

Sample Request

```
GET /api/v1/users HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

[
  {
    "name": "myuser",
    "password": "mypassword",
    "enable": true,
    "linuxUID": 1000,
    "windowsSID": "S-1-5-21-1634518050-1170479241-3746985662-1000",
    "mainGroup": "Administrators",
    "additionalGroups": ["users"],
    "permissions": {
      "SET_SHARE": true,
      "GET_SHARE": true,
      "SET_POLICY": true,
      "GET_POLICY": true,
      "TAKE_SNAPSHOT": true,
      "SET_USER_GROUP": true,
      "GET_USER_GROUP": true,
      "GET_SYSTEM_REPORTS": true,
      "GET_FILESYSTEM_REPORTS": true
    }
  },
  ...
]
```

Groups

The Group object

```
{
  "name": "users",
  "linuxGID": 1000,
  "windowsSID": "S-1-5-21-1634518050-1170479241-3746985662-1000",
  "permissions": {
    "SET_SHARE": true,
    "GET_SHARE": true,
    "SET_POLICY": true,
    "GET_POLICY": true,
    "TAKE_SNAPSHOT": true,
    "SET_USER_GROUP": true,
    "GET_USER_GROUP": true,
    "GET_SYSTEM_REPORTS": true,
    "GET_FILESYSTEM_REPORTS": true
  }
}
```

Attribute Definitions

Name	Type	Description
name	String	The name of the group. Must be unique in the scope of the file system.
linuxGID	Integer	The Linux group ID.
windowsSID	String	The Windows Security Identifier.
permissions	Object	A set of attributes with Boolean values that configure the group's access permissions on the file system. Valid attributes are: * SET_SHARE (create/modify/delete directories and shares) * GET_SHARE (get / list directories and shares) * SET_POLICY (create/modify/delete snapshot, file, quota and antivirus policies) * GET_POLICY (get/list snapshot, file, quota and antivirus policies) * TAKE_SNAPSHOT (take manual snapshot) * SET_USER_GROUP (create/modify/delete users and groups) * GET_USER_GROUP (get/list users and groups) * GET_SYSTEM_REPORTS (get System reports via the Rest-API) * GET_FILESYSTEM_REPORTS (get Filesystem reports via the Rest-API)

Create Group

POST /api/v1/groups

Create a group.

Request Body

The Group object. Name is required. WindowsSID is ignored.

Sample Request

```
POST /api/v1/groups HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "name": "users",
  "linuxGID": 1000,
  "permissions": {
    "SET_SHARE": true,
    "GET_SHARE": true,
    "SET_POLICY": true,
    "GET_POLICY": true,
    "TAKE_SNAPSHOT": true,
    "SET_USER_GROUP": true,
    "GET_USER_GROUP": true,
    "GET_SYSTEM_REPORTS": true,
    "GET_FILESYSTEM_REPORTS": true
  }
}
```

Sample Response

```
HTTP/1.1 201 Created
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
Location: http://<node>:81/api/v1/groups/users
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Modify Group

PUT /api/v1/groups/<group>

Modify an existing user.

Path Parameters

group	The name of the group.

Request Body

The Group object. Name is optional. Any parameters not provided remains unchanged.

Sample Request

```
PUT /api/v1/groups/users HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "name": "users2"
}
```

This sample request renames the group from users to users2. The other user configuration remains unchanged.

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Delete Group

DELETE */api/v1/groups/<group>*

Delete a group.

Path Parameters

group	The name of the group.

Sample Request

```
DELETE /api/v1/groups/users HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 204 No Content
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Get Group

GET /api/v1/groups/<group>

Retrieve a group.

Path Parameters

group	The name of the group.

Response Body

The Group object.

Sample Request

```
GET /api/v1/groups/users HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

{
  "name": "users",
  "linuxGID": 1000,
  "windowsSID": "S-1-5-21-1634518050-1170479241-3746985662-1000",
  "permissions": {
    "SET_SHARE": true,
    "GET_SHARE": true,
    "SET_POLICY": true,
    "GET_POLICY": true,
    "TAKE_SNAPSHOT": true,
    "SET_USER_GROUP": true,
    "GET_USER_GROUP": true,
    "GET_SYSTEM_REPORTS": true,
    "GET_FILESYSTEM_REPORTS": true
  }
}
```

List Groups

GET */api/v1/groups*

List all groups.

Response Body

A JSON array of Group objects.

Sample Request

```
GET /api/v1/groups HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

[
  {
    "name": "users",
    "linuxGID": 1000,
    "windowsSID": "S-1-5-21-1634518050-1170479241-3746985662-1000",
    "permissions": {
      "SET_SHARE": true,
      "GET_SHARE": true,
      "SET_POLICY": true,
      "GET_POLICY": true,
      "TAKE_SNAPSHOT": true,
      "SET_USER_GROUP": true,
      "GET_USER_GROUP": true,
      "GET_SYSTEM_REPORTS": true,
      "GET_FILESYSTEM_REPORTS": true
    }
  },
  ...
]
```

Actions

Take Snapshot

POST */api/v1/actions/snapshot*

Take a manual snapshot of a directory. Note that the directory must have been configured with a snapshot policy, however it is not required that the schedule type “manually” is included.

Request Headers

x-cv-cwd-id	Optional. The ID of the directory the path in the request body is relative to.

Request Body

Name	Type	Description	Status
path	String	The path of the directory the snapshot should be applied to, relative to the current working directory.	Required
expireDate	String	When the snapshot will expire, as UTC time (ex. 2016-10-31 00:00:00). After the snapshot expires, it is removed.	Required

Sample Request

```
POST /api/v1/actions/snapshot HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
Content-Length: <length>

{
  "path": "path/to/directory",
  "expireDate": "2016-10-31 00:00:00"
}
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Length: 0
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e
```

Reports

The Report

A report is a collection of CSV-formatted text data, describing the state of an IBM Spectrum NAS system (storage node) or of an IBM Spectrum NAS filesystem, at a certain date.

Each report is generated by a specific storage node in the cluster (the one that the REST-API request is directed to), and contains information generated by that node only.

There are two categories of reports:

- **System report** - contains information relevant only to the node that sends the report, e.g. the amount of free memory on the node, or the fact that a storage disk of the node was brought offline.
- **Filesystem report** - contains information relevant to a filesystem, e.g. a file share was created, or a quota policy on a folder was modified. All nodes in the cluster have the same view of this type of information.

System reports

A System report is a collection of CSV-formatted text data that describes the state of an IBM Spectrum NAS system (storage node), at a certain date.

The report contains an initial header line and a number of rows. The values are delimited by the comma character (',').

The report is produced from the log files generated by each IBM Spectrum NAS storage node. Note that deleting the log files for a specific date from the storage node will result in reports for that date becoming permanently unavailable.

There are two sub-types of System reports: Performance and Audit.

Performance System report

A Performance System report describes the performance state of a storage node, at a certain date. Each row in the report shows the performance state of the system at a certain time, e.g. the instant CPU usage on the node, or the number of Read/Write operations performed since the previous row in the report was generated. The IBM Spectrum NAS logging system generates every 5 minutes a new row that describes the current performance state of the node.

GET */api/v1/reports/system/performance/<date>*

Retrieve the Performance System report for the respective date.

Path Parameters

date	The date of the report, in the format: yyyyMMdd

Response Body

The Performance System report.

Sample Request

```
GET /api/v1/reports/system/performance/20180115 HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Mon, 15 Jan 2018 12:01:37 GMT
Server: IBM Spectrum NAS Management
Content-Type: application/octet-stream
Content-Length: <length>
x-cv-request-id: 880c0b32-9ad0-4fbe-b338-3e48bacdd75e

Date,Time,Private IP,CPU Usage (%),Memory Total (MB),Memory Free (MB),Swap Usage (MB),Process
Used Memory (MB),Process Used Swap (MB),Network Utilization (%),Gateway Ops,Read Ops,Read
Data (MB),Reads/sec,Read Avg Latency (ms),Read Max Latency (ms),Write Ops,Write Data
(MB),Writes/sec,Write Avg Latency (ms),Write Max Latency (ms),NFS Ops,NFS Ops/sec,NFS Avg
Latency (ms),NFS Max Latency (ms),SMB Ops,SMB Ops/sec,SMB Avg Latency (ms),SMB Max Latency
(ms)
2018-02-
16,13:40,172.16.1.129,80,7986,4798,0,711,0,1,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0
2018-02-
16,13:45,172.16.1.129,43,7986,4665,0,774,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0
2018-02-
16,13:50,172.16.1.129,29,7986,4462,0,856,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0
2018-02-
16,13:55,172.16.1.129,44,7986,3979,0,877,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0
```

The significance of each column in the Performance System report is described in the following table.

Name	Description
Date	The date of the report
Time	The time when the row was generated by the storage node
Private IP	The Private IP address of the storage node that generated the report
CPU Usage (%)	The CPU Usage value (%) on the node, at the time when the row was generated
Memory Total (MB)	The total RAM memory installed on the node, in MB
Memory Free (MB)	The free RAM memory on the node, in MB, at the time when the row was generated
Swap Usage (MB)	The swap memory usage on the node, in MB, at the time when the row was generated
Process Used Memory (MB)	The RAM memory used by the IBM Spectrum NAS Gateway process, in MB, at the time when the row was generated
Process Used Swap (MB)	The swap memory used by the IBM Spectrum NAS Gateway process, in MB, at the time when the row was generated
Network Utilization (%)	The utilization of the network interfaces on the node (%), at the time when the row was generated
Gateway Operations	The amount of operations that were performed via all active protocols (e.g. SMB, NFS), since the previous row had been generated (e.g. list directories, create files, delete files, read/write to files etc.)
Read Operations	The amount of read operations that were performed via all active protocols (e.g. SMB, NFS), since the previous row had been generated
Read Data (MB)	The amount of data, in MB, that was transferred as a result of read operations via all active protocols (e.g. SMB, NFS), since the previous row had been generated

Reads/sec	The amount of read operations per second that were performed via all active protocols (e.g. SMB, NFS), since the previous row had been generated
Read Avg Latency (ms)	The average response time, in milliseconds, of all read operations that were performed via all active protocols (e.g. SMB, NFS), since the previous row had been generated
Read Max Latency (ms)	The maximum response time, in milliseconds, of all read operations that were performed via all active protocols (e.g. SMB, NFS), since the previous row had been generated
Write Operations	The amount of write operations that were performed via all active protocols (e.g. SMB, NFS), since the previous row had been generated
Write Data (MB)	The amount of data, in MB, that was transferred as a result of write operations via all active protocols (e.g. SMB, NFS), since the previous row had been generated
Writes/sec	The amount of write operations per second that were performed via all active protocols (e.g. SMB, NFS), since the previous row had been generated
Write Avg Latency (ms)	The average response time, in milliseconds, of all write operations that were performed via all active protocols (e.g. SMB, NFS), since the previous row had been generated
Write Max Latency (ms)	The maximum response time, in milliseconds, of all write operations that were performed via all active protocols (e.g. SMB, NFS), since the previous row had been generated
NFS Ops	The amount of operations that were performed via the NFS protocol, since the previous row had been generated
NFS Ops/sec	The amount of operations per second that were performed via the NFS protocol, since the previous row had been generated
NFS Avg Latency (ms)	The average response time, in milliseconds, of all operations that were performed via the NFS protocol, since the previous row had been generated
NFS Max Latency (ms)	The maximum response time, in milliseconds, of all operations that were performed via the NFS protocol, since the previous row had been generated
SMB Ops	The amount of operations that were performed via the SMB protocol, since the previous row had been generated
SMB Ops/sec	The amount of operations per second that were performed via the SMB protocol, since the previous row had been generated
SMB Avg Latency (ms)	The average response time, in milliseconds, of all operations that were performed via the SMB protocol, since the previous row had been generated
SMB Max Latency (ms)	The maximum response time, in milliseconds, of all operations that were performed via the SMB protocol, since the previous row had been generated

Audit System report

An Audit System report describes administrative events that occurred on the IBM Spectrum NAS system (storage node), at a certain date. Each row in the report is an administrative event (e.g. a node was taken online/offline, a disk was taken online/offline). The IBM Spectrum NAS logging system generates a new row when the respective event occurs.

GET */api/v1/reports/system/audit/<date>*

Retrieve the Audit System report for the respective date.

Path Parameters

date	The date of the report, in the format: yyyyMMdd

Response Body

The Audit System report.

Sample Request

```
GET /api/v1/reports/system/audit/20180115 HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
Date,Time,Private IP,Event
2018-01-15,08:38,172.16.1.129,A request to take the node offline was received from
172.16.0.42 (grace period 00:30:00).
2018-01-15,08:41,172.16.1.129,A request to take the node online was received from
172.16.0.42.
2018-01-15,15:11,172.16.1.129,A request to take online the storage disk in slot 1 was
received from 172.16.0.42.
```

The significance of each column in the Audit System report is described in the following table.

Name	Description
Date	The date of the report
Time	The time when the row was generated by the storage node
Private IP	The Private IP address of the storage node that generated the report
Event	The administrative event that occurred on the storage node

Audit System report Events

Node was taken offline
Node was taken online
Disk was taken offline
Disk was taken online
Disk was retired

Filesystem reports

A Filesystem report describes the state of an IBM Spectrum NAS filesystem, at a certain date.

There are two types of Filesystem reports, the Audit Filesystem report and the DiskUsage Filesystem report.

Audit Filesystem report

An Audit Filesystem report is a collection of CSV-formatted text data that describes administrative events that occurred on the IBM Spectrum NAS filesystem, at a certain date.

The report contains an initial header line and a number of rows. The values are delimited by the comma character (',').

Each row represents an administrative event (creation, modification or deletion of file shares, file/quota/snapshot/antivirus policies, users, user groups).

The report is produced from the log files generated by each IBM Spectrum NAS storage node. Note that deleting the log files for a specific date from the storage node will result in reports for that date becoming permanently unavailable.

GET */api/v1/reports/filesystem/audit/<date>*

Retrieve the Audit Filesystem report for the respective date.

Path Parameters

date	The date of the report, in the format: yyyyMMdd

Response Body

The Audit Filesystem report.

Sample Request

```
GET /api/v1/reports/filesystem/audit/20180115 HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
Date,Time,Private IP,Event
2018-01-15,10:20,172.16.1.129,Administrator 'user' from 192.168.4.42 created share 'share1'.
Protocols: 'SMB'. Absolute path: '/rest-shares/share1'
2018-01-15,10:24,172.16.1.129,Administrator from 172.16.0.42 deleted file policy on folder
'subfolder1' at absolute path '/rest-shares/share1/subfolder1'
```

The significance of each column in the Audit Filesystem report is described in the following table.

Name	Description
Date	The date of the report
Time	The time when the row was generated by the storage node
Private IP	The Private IP address of the storage node that generated the report
Event	The administrative event that occurred on the filesystem

Audit Filesystem report Event Summary

File share was created/modified/deleted
File/Quota/Snapshot/Antivirus policy was created/modified/deleted
User was created/modified/deleted
User Group was created/modified/deleted

DiskUsage Filesystem report

A DiskUsage Filesystem report gives information about the disk space occupied by a folder protected with a snapshot policy in the IBM Spectrum NAS filesystem.

The Snapshot DiskUsage object

```
{
  "ordinal": 1,
  "name": "days.0",
  "createDate": "2018-02-12 09:34:02",
  "expireDate": "2018-03-19 09:34:00",
  "diskSize": 1404,
  "status": "complete"
}
```

Attribute Definitions

Name	Type	Description
ordinal	Integer	A numerical label of the snapshot copy. Among all snapshot copies taken for a folder, the one with the smallest ordinal is the most recent.
name	String	The name of the folder where the snapshot copy can be accessed.
createDate	String	When the snapshot was created, as UTC time (ex. 2016-10-31 00:00:00).
expireDate	String	When the snapshot will expire, as UTC time (ex. 2016-10-31 00:00:00). After the snapshot expires, it is removed.
diskSize	Integer	For a file, this attribute indicates the actual size, in bytes, of the file (not-written areas in the file excluded, erasure coding data included) and of the metadata associated with the file, as stored on the storage disk. For a directory, this attribute indicates the sum of the “diskSize” values for each file in the directory and all its subdirectories. Note: This attribute refers to a subdirectory under the .snapshot directory and indicates the total “diskSize” used by the respective subdirectory in order to store the required snapshot data. Note: Each snapshot is differential (only stores the data that differs from the previous snapshot).
status	String	An attribute that describes the current status of the snapshot. Possible values are “complete” and “inprogress”. A snapshot that is in progress is currently being built and cannot yet be used for reverting data.

The Folder DiskUsage object

```
{
  "path": "/path/to/directory",
  "size": 4145493381,
  "usedSize": 4145493360,
  "diskSize": 6368710686,
  "snapshotsDiskSize": 1695,
  "totalDiskSize": 6368712381,
  "snapshots": [
    {
             "ordinal": 1,
             "name": "days.0",
             "createDate": "2018-02-12 09:34:02",
             "expireDate": "2018-03-19 09:34:00",
             "diskSize": 1404,
             "status": "complete"
    },
    {
             "ordinal": 2,
             "name": "days.1",
             "createDate": "2018-02-11 08:11:01",
             "expireDate": "2018-03-12 09:34:00",
             "diskSize": 291,
             "status": "complete"
    }
  ]
}
```

Attribute Definitions

Name	Type	Description
path	String	The path to the directory that is protected by the snapshot policy. The path is relative to the current working directory. If the path starts with “/”, it is absolute (starts at the file system root).
size	Integer	For a file, this attribute indicates the total size, in bytes, of the file. For a directory, this attribute indicates the sum of the “size” values for each file in the directory and all its subdirectories. If the directory has a .snapshot subfolder, the “size” of the .snapshot subfolder is not included in the value of this attribute.
usedSize	Integer	For a file, this attribute indicates the actual size, in bytes, of the data in the file (not-written areas in the file excluded). For a directory, this attribute indicates the sum of the “usedSize” values for each file in the directory and all its subdirectories. Note: due to thin-provisioning, “usedSize” is usually less than “size”. If the directory has a .snapshot subfolder, the “usedSize” of the .snapshot subfolder is not included in the value of this attribute.
diskSize	Integer	For a file, this attribute indicates the actual size, in bytes, of the file (not-written areas in the file excluded, erasure coding data included) and of the metadata associated with the file, as stored on the storage disk. For a directory, this attribute indicates the sum of the “diskSize” values for each file in the directory and all its subdirectories. If the directory has a .snapshot subfolder, the “diskSize” of the .snapshot subfolder is not included in the value of this attribute.

snapshotsDiskSize	Integer	The sum of the “diskSize” values for each existent snapshot (each snapshot is a subfolder of the .snapshot folder).
totalDiskSize	Integer	The sum of the “diskSize” and the “snapshotsDiskSize” for the folder.
snapshots	Array	An array of Snapshot DiskUsage objects. Each element of the array represents an existent snapshot (a subfolder of the .snapshot folder).

GET /api/v1/reports/filesystem/diskusage/<path>

Retrieve the DiskUsage Filesystem report for the respective folder.

Path Parameters

path	The path of the directory protected with a snapshot policy, relative to the current working directory.

Response Body

The Folder DiskUsage object.

Sample Request

```
GET /api/v1/reports/filesystem/diskusage/path/to/directory HTTP/1.1
Host: <node>
Authorization: Basic <base64_encoded_string>
```

Sample Response

```
HTTP/1.1 200 OK
Date: Fri, 14 Oct 2016 12:05:32 GMT
Server: IBM Spectrum NAS Management
Content-Length: <length>
x-cv-request-id: c69a0b0c-236c-4372-8fc6-2caaf85551b8
x-cv-parent-id: 3b100000-0000-0000-1956-56df2b153a43

{
  "size": 4145493381,
  "usedSize": 4145493360,
  "diskSize": 6368710686,
  "snapshotsDiskSize": 1695,
  "totalDiskSize": 6368712381,
  "snapshots": [
    {
            "ordinal": 1,
            "name": "days.0",
            "createDate": "2018-02-12 09:34:02",
            "expireDate": "2018-03-19 09:34:00",
            "diskSize": 1404,
            "status": "complete"
    },
    {
            "ordinal": 2,
            "name": "days.1",
            "createDate": "2018-02-11 08:11:01",
            "expireDate": "2018-03-12 09:34:00",
            "diskSize": 291,
            "status": "complete"
    }
  ]
}
```

Appendix

Common Request Headers

x-cv-request-id	The unique ID of the request. Used for troubleshooting.

Errors

Any response with a 4xx or 5xx HTTP status code includes an Error object in the response.

The Error object

```
{
  "code": 72,
  "error": "UserNotFound",
  "message": "UserNotFound"
}
```

Attribute Definitions

Name	Type	Description
code	Integer	Uniquely identifies the error.
error	String	Human readable string of the numerical code.
message	String	Description of the error. For some errors, this has the same value as the error attribute.

The following table is a list of the errors:

Code	Error	Description	HTTP Status Code
53	AccessDenied	Access denied. The user doesn't have the required permissions, or the ID in the x-cv-cwd-id header is not valid for the current domain / tenant.	403 Forbidden
11	DirectoryNotEmpty	The specified directory is not empty.	409 Conflict
101 or 102	GroupAlreadyExists	A group with the specified name / GID already exists.	409 Conflict
100	GroupNotFound	The group does not exist.	404 Not Found
Depends	InternalError	An internal server error occurred. The code varies, depending on the type of server error.	500 Internal Server Error
5	InvalidArgument	The request contains invalid arguments. Malformed JSON document, JSON attributes that doesn't conform to the specification, or invalid query parameters.	400 Bad Request
17	InvalidCredentials	User authentication failed.	403 Forbidden

15	InvalidMessage	There isn't an operation that matches the request, e.g. the combination of HTTP method and request URI is invalid.	400 Bad Request
51	PathAlreadyExists	A directory already exists at the specified path.	409 Conflict
50	PathNotFound	The directory does not exist.	404 Not Found
91	PolicyAlreadyExists	The directory already has a configuration for the specified policy.	409 Conflict
90	PolicyNotFound	The directory does not have a configuration for the specified policy.	404 Not Found
10	RequestBodyTooLarge	The size of the request body is larger than what is allowed.	400 Bad Request
81	ShareAlreadyExists	A file share with the specified name already exists.	409 Conflict
80	ShareNotFound	The file share does not exist.	404 Not Found
73 or 74	UserAlreadyExists	A user with the specified name / UID already exists.	409 Conflict
72	UserNotFound	The user does not exist.	404 Not Found

File encodings

DEFAULT
COPIES_3
COPIES_5
ERASURE_2_1
ERASURE_2_2
ERASURE_3_1
ERASURE_3_2
ERASURE_4_1
ERASURE_4_2
ERASURE_5_1
ERASURE_5_2
ERASURE_6_1
ERASURE_6_2
ERASURE_8_1
ERASURE_8_2
METRO_COPIES_3
METRO_ERASURE_2_1
METRO_ERASURE_2_2

METRO_ERASURE_3_1
METRO_ERASURE_3_2
METRO_ERASURE_4_1
METRO_ERASURE_4_2
METRO_ERASURE_5_1
METRO_ERASURE_5_2
METRO_ERASURE_6_1
METRO_ERASURE_6_2

Tiers

ANY
DEFAULT
TIER_0
TIER_1
TIER_2
TIER_3
TIER_4

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM® product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing Legal and Intellectual Property Law IBM Japan Ltd. 19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).
Portions of this code are derived from IBM Corp.
Sample Programs. © Copyright IBM Corp. _enter the year or years_.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Intel is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of the Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to

collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Printed in USA

SC27-9227-01