nirva

Nirva STORAGE Service

Document Version: 1.23

Table of Contents

Overview

The storage service is a NIRVA external service that manages the persistent storing and retrieval of documents on physical medias.

Here are the major features:

- Persistent storage of documents
- Volume management
- Storage levels
- Synchronization between storage levels
- Storing of flat files or entire NIRVA containers

Volumes

A volume is the logical place where the storage service stores a document.

There are two kinds of volumes: directory or storage volume:

- The directory volume has a direct connection to a single physical mount point (directory). It allows direct control of the directory structure and content. One can decide the name of stored files and directories or let the service do that automatically.
- The storage volume is a collection of storage levels, each level corresponding to a physical mount point. The storage volume manages itself the structure of the physical data but provides more functionality than directory volume. This includes storage in several physical places at the same time, synchronization between these physical storages, etc…

A volume is uniquely identified by its name. The volume name is case independent and cannot contain the characters ‘ ‘ (space), ‘.’, ‘/’, ‘\’, ‘-‘.

The storage service configuration tools provides the necessary functionality to create or remove volumes:

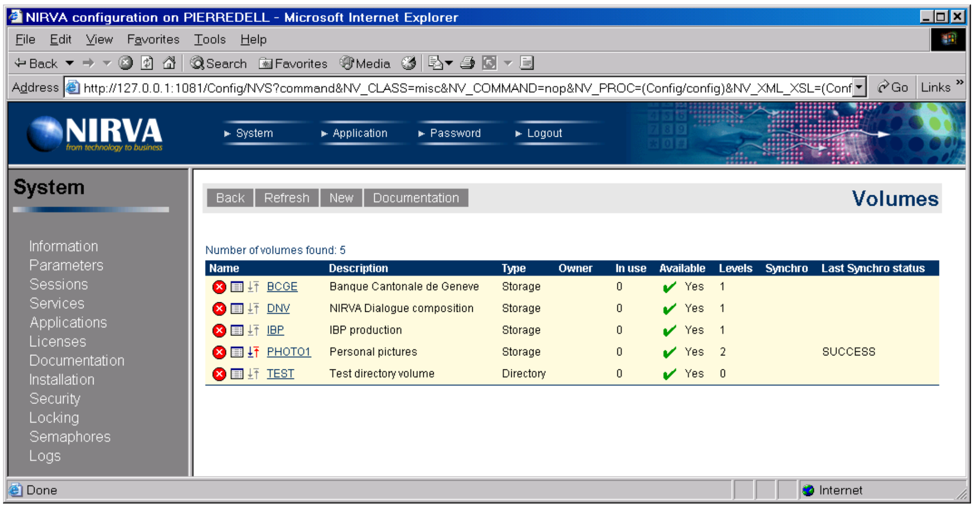
NIRVA configuration on PIERREDELL - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://127.0.0.1:1081/Config/NVS?command&NV_CLASS=misc&NV_COMMAND=nop&NV_PROC=(Config/config)&NV_XML_XSL=(Conf
Go Links
NIRVA
from technology to business
System
Application
Password
Logout
System
Information
Parameters
Sessions
Services
Applications
Licenses
Documentation
Installation
Security
Locking
Semaphores
Logs
Back
Refresh
New
Documentation
Done
Internet

Volumes

Number of volumes found: 5

Name	Description	Type	Owner	In use	Available	Levels	Synchro	Last Synchro status
BCGE	Banque Cantonale de Geneve	Storage		0	✓ Yes	1		
DNV	NIRVA Dialogue composition	Storage		0	✓ Yes	1		
IBP	IBP production	Storage		0	✓ Yes	1		
PHOTO1	Personal pictures	Storage		0	✓ Yes	2		SUCCESS
TEST	Test directory volume	Directory		0	✓ Yes	0		

Storage levels

Storage levels are available only for volumes of storage type. Each volume has its collection of storage levels. The storage level is the physical place where the document is stored. During write operation, the service writes the document to all the levels except if a level has a write mode set to “NEVER”.

A storage level is uniquely identified by its name. The storage level name is case independent.

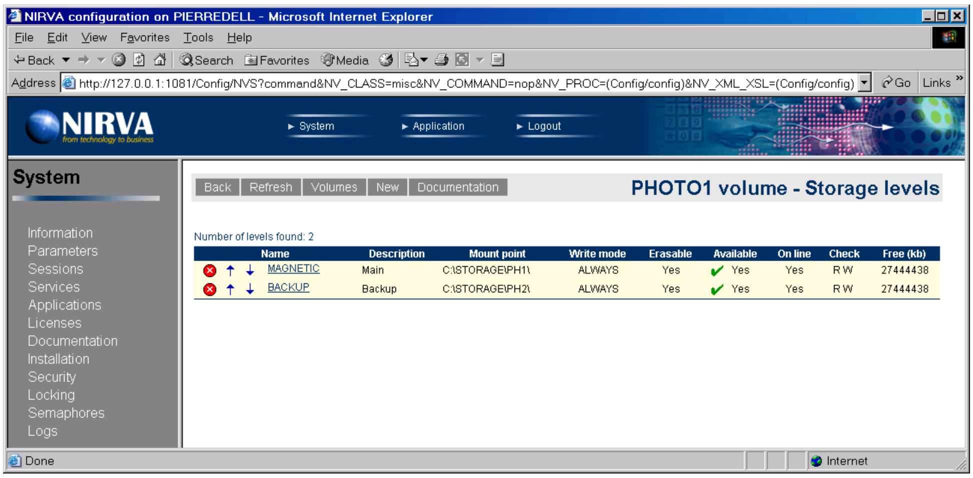
NIRVA configuration on PIERREDELL - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://127.0.0.1:1081/Config/NVS?command&NV_CLASS=misc&NV_COMMAND=nop&NV_PROC=(Config/config)&NV_XML_XSL=(Config/config)
Go Links
NIRVA
from technology to business
System
Application
Logout
System
Information
Parameters
Sessions
Services
Applications
Licenses
Documentation
Installation
Security
Locking
Semaphores
Logs
Back
Refresh
Volumes
New
Documentation
Done
Internet

PHOTO1 volume - Storage levels

Number of levels found: 2

Name	Description	Mount point	Write mode	Erasable	Available	On line	Check	Free (kb)
MAGNETIC	Main	C:\STORAGE\PH1\	ALWAYS	Yes	✓ Yes	Yes	R W	27444438
BACKUP	Backup	C:\STORAGE\PH2\	ALWAYS	Yes	✓ Yes	Yes	R W	27444438

In reading mode, the order of the storage levels is important. When reading a document, NIRVA first tries to get it from the first level. If not found here, it tries to get it on the second level. And so on…

The order of the storage levels can be changed by the configuration tool.

In writing mode, a level works differently following it’s write mode:

- ALWAYS means that NIRVA will always write the document on this level and will report an error if unsuccessful. The first level is often defined with the ALWAYS mode.
- TRY means that NIRVA will try to write the document on this level but will not report an error if not successful.

- NEVER means that NIRVA will not write the document during the WRITE command on this level. The level will then be written by using the level synchronization commands.

In any case, the write command will report an error if it cannot write the document on at least one level.

Documents

The storage service in fact stores documents on the volumes. A document can be a flat file of any kind or a complete NIRVA session hierarchical container.

Flat files

When storing a flat file, the storage service just requires the name of a NIRVA file object to store that must be in the input container. It then stores the attached file and gives back it’s reference as a NIRVA string object in the output container.

Containers

The storage service is able to store an entire NIRVA container. At this time it creates an export file that contains the entire container an stores this export file as a flat file on the volume. It then gives back it’s reference as a NIRVA string object in the output container.

Reference

When a document has been stored on a volume, it’s identified by a reference. A document reference has the following format:

```
Volume-[-]rpath_docpart
```

Volume is the volume name.

rpath is the relative path and filename on which the document is stored. The separator between the volume and the relative path is the ‘-‘ character in case of a flat file or 2 ‘-‘ characters (‘- -‘) in case of a container.

The reference may then be followed by a ‘_’ character (underscore) and by an object name (*docpart*). This can be used to access only a single object of a stored container.

Here are some valid document references:

- *PH1-200212/20/18/16/150001* -> Flat file
- *PH1--200212/20/18/16/150002* -> Container
- *PH1--200212/20/18/16/150003_obj1* -> object obj1 of stored container

- *PH1--200212/20/18/16/150003_sub1.obj1* -> object obj1 of sub-container sub1 of stored container
- *PH2-test/test1.pdf* -> Flat file (directory volume)

Repositories

When storing a document, the user can decide to put it in a repository. A repository is simply a subdirectory of the storage media. A repository name is case insensitive, it cannot contain any space characters and must have at least 3 characters.

The use of repositories is not mandatory. When not used, the storage service is storing documents in the default repository (blank repository).

Files and directories

On the storage media, the documents are stored differently following the type of volume

Directory volume

For a directory volume the documents are stored on a the following structure:

```
MountPoint[/Repository][/Subdir/]Filename
```

- *MountPoint* is the mount point of the storage level.
- *Repository* is the optional repository name.
- *Subdir* is the optional subdirectory name (can contain several subdirectories separated bya '/' character (ex test1/myfiles/pdf).
- *Filename* is the file name itself. It can b e directly given or automatically calculated by the service. At this time, its format is [Prefix]*YYYYMMDDHHMMSSI[.Extension]* where *YYYYMMDDHHMMSS* corresponds to the date and time, *Prefix* is an optional prefix, *I* is an index starting at 0 and *Extension* is an optional extension.

Storage volume

For a storage volume the documents are stored on a date based structure:

```
MountPoint/YYYYMM/DD[/Repository]/HH/MM/[Prefix]SSI
```

- *MountPoint* is the mount point of the storage level.
- *YYYY* is the year.
- *MM* is the month of the year from 1 to 12.
- *DD* is the day of the month from 1 to 31.
- *Repository* is the optional repository name.
- *HH* is the hour of the day from 0 to 23.
- *MM* is the minute of the hour from 0 to 59.
- *Prefix* is an optional file prefix.
- *SS* is the second of the minute from 0 to 59.
- *I* is an index staring at 0.

If several levels has been defined, the service will write a document on each level having a write mode set to “ALWAYS” or “TRY”. The service then warranty that the same file name is used for all levels.

Volume groups

Several volumes can be grouped together for writing purpose. When writing a document, one can give the name of a group instead of the volume name. At this time, the service automatically chooses one of the volumes of the group to write on. After each group write, the service marks it as the last written volume and simply chooses the next from the volume list of the group when the next write order arrives.

This feature allows to optimize performances when an important amount of information is to be written on the system. In fact the storage service locks access at volume levels. So if we want to write several files in parallel, this must be done onto different volumes. For example, if we have defined a group of 10 volumes, the system will be able to process 10 writes at the same time. If we have only one volume, all orders are processed one after the other.

Installation

The storage service is delivered as a NIRVA package and can be installed like any NIRVA service directly from the NIRVA configuration web site. Please see the NIRVA configuration chapter in the NIRVA user's guide for further information.

Configuration

The storage service configuration is entirely dynamic and available from a web browser.

The configuration of the storage service is accessible directly from the main list of services of the NIRVA configuration web site:

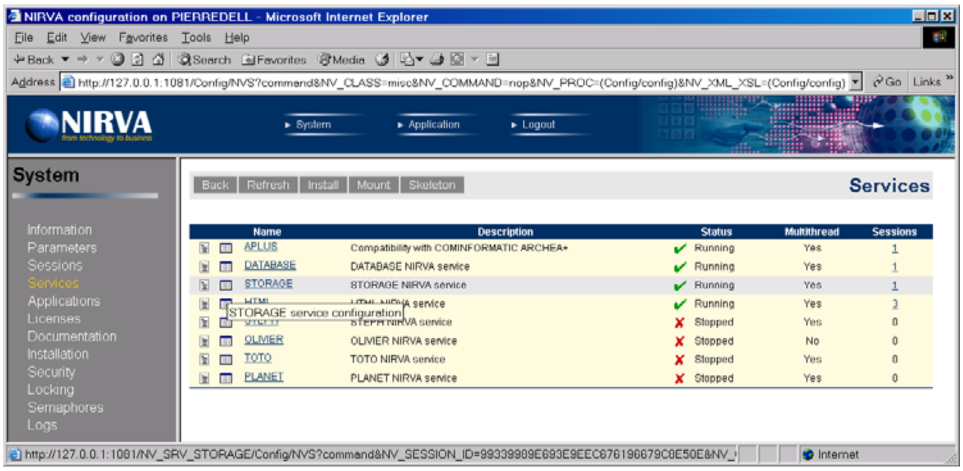
NIRVA configuration on PIERREDELL - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://127.0.0.1:1081/Config/NVS?command&NV_CLASS=misc&NV_COMMAND=nop&NV_PROC=(Config/config)&NV_XML_XSL=(Config/config)
Go Links
NIRVA
from technology to business
System
Application
Logout
System
Information
Parameters
Sessions
Services
Applications
Licenses
Documentation
Installation
Security
Locking
Semaphores
Logs
Back
Refresh
Install
Mount
Skeleton
Services
STORAGE service configuration
http://127.0.0.1:1081/NV_SRV_STORAGE/Config/NVS?command&NV_SESSION_ID=99339909E693E9EEC676196679C0E50E&NV_
Internet

Name	Description	Status	Multithread	Sessions
APLUS	Compatibility with COMINFORMATIC ARCHEA+	✓ Running	Yes	1
DATABASE	DATABASE NIRVA service	✓ Running	Yes	1
STORAGE	STORAGE NIRVA service	✓ Running	Yes	1
HTML	HTML NIRVA service	✓ Running	Yes	3
STEPH	STEPH NIRVA service	✗ Stopped	Yes	0
OLIVIER	OLIVIER NIRVA service	✗ Stopped	No	0
TOTO	TOTO NIRVA service	✗ Stopped	Yes	0
PLANET	PLANET NIRVA service	✗ Stopped	Yes	0

Volume management

When choosing the configuration icon of the storage service (▤), this starts the storage service configuration by displaying the list of storage volumes:

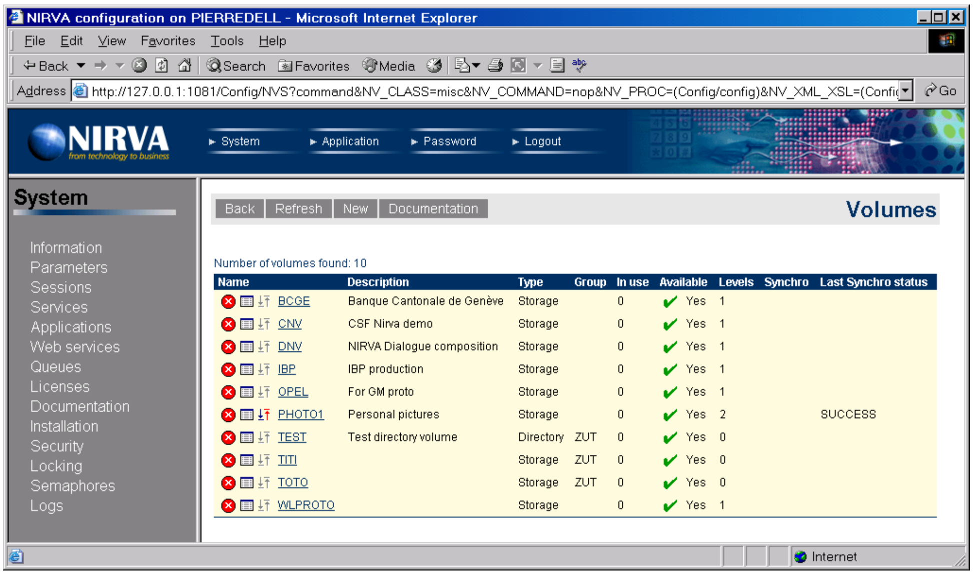
NIRVA configuration on PIERREDELL - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://127.0.0.1:1081/Config/NVS?command&NV_CLASS=misc&NV_COMMAND=nop&NV_PROC=(Config/config)&NV_XML_XSL=(Confic
Go
NIRVA
from technology to business
System
Application
Password
Logout
System
Information
Parameters
Sessions
Services
Applications
Web services
Queues
Licenses
Documentation
Installation
Security
Locking
Semaphores
Logs
Back
Refresh
New
Documentation
Volumes
Internet

Number of volumes found: 10

Name	Description	Type	Group	In use	Available	Levels	Synchro	Last Synchro status
BCGE	Banque Cantonale de Genève	Storage		0	✓ Yes	1		
CNV	CSF Nirva demo	Storage		0	✓ Yes	1		
DNV	NIRVA Dialogue composition	Storage		0	✓ Yes	1		
IBP	IBP production	Storage		0	✓ Yes	1		
OPEL	For GM proto	Storage		0	✓ Yes	1		
PHOTO1	Personal pictures	Storage		0	✓ Yes	2		SUCCESS
TEST	Test directory volume	Directory	ZUT	0	✓ Yes	0		
TITI		Storage	ZUT	0	✓ Yes	0		
TOTO		Storage	ZUT	0	✓ Yes	0		
WLPROTO		Storage		0	✓ Yes	1		

This is the list of actual volumes seen by the storage system. For each volume, the list displays the following information:

- *Name* is the volume name
- *Description* is the volume description.
- *Type* is the volume type (directory or storage).
- *Group* is the name of the volume group is there is one.
- *Owner* is the owner of the volume. The owner feature is reserved for future use.
- *In use* is a counter of actual sessions requesting the use of the volume. Only one session at a time can access the volume in read or write mode. The other sessions that want to access the volume are put in wait state by the storage service until the volume is free.
- *Available* tells if the volume is available or not. A disabled volume is still in the volume list but is not accessible.
- *Levels* gives the number of volume levels. The volume levels are the mount points for a storage volume. For a directory volume, the number of storage levels is always 0.
- *Synchro* gives the progression of the synchronization process (in percent) if one has been launched. When this information is not blank, the user can stop the synchronization process by clicking on the *Synchro* information. The synchronization feature is available only for a storage volume.
- *Last synchro status* gives the status of last ran synchronization process. This can be “STARTED”, “FAILED”, “STOPPED” or “SUCCESS”. The synchronization feature is available only for a storage volume.

Creating a new volume

Creating a new volume is done by clicking the button “New”. This displays the following screen:

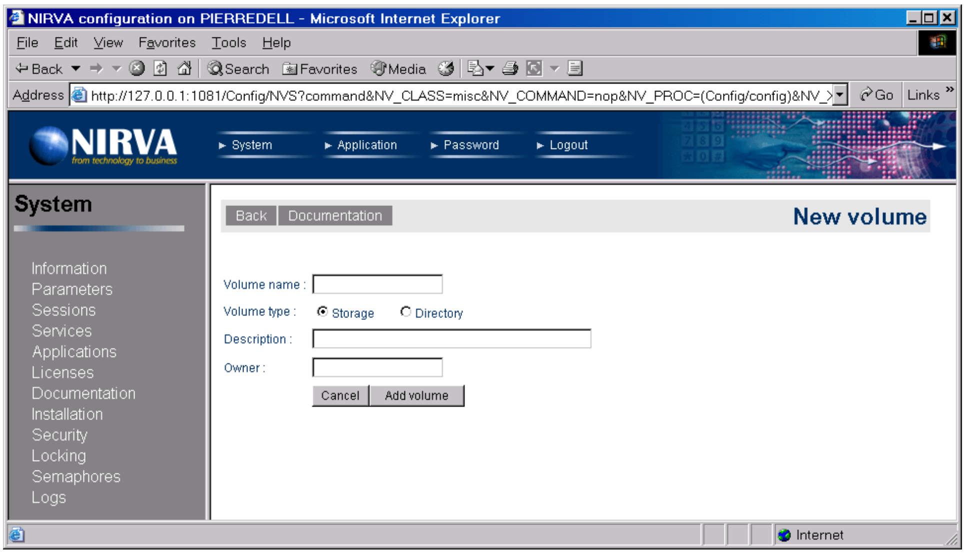
NIRVA configuration on PIERREDELL - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://127.0.0.1:1081/Config/NVS?command&NV_CLASS=misc&NV_COMMAND=nop&NV_PROC=(Config/config)&NV_
Go Links
NIRVA from technology to business
► System ► Application ► Password ► Logout
System
Information
Parameters
Sessions
Services
Applications
Licenses
Documentation
Installation
Security
Locking
Semaphores
Logs
Back Documentation
Cancel Add volume
Internet

New volume

Volume name :

Volume type : ● Storage ○ Directory

Description :

Owner :

The volume name is mandatory. The volume name is case independent and cannot contain the characters ' ' (space), '.', '/', '\', '-'.

If the user attempts to create a volume with the name of an existing one, the command fails.

The volume type must be defined here. After the volume creation, it's not possible to change its type.

Removing a volume

Removing a volume is quiet simple. Just select the ⊗ icon near the volume name in the volume list. The configuration tool will then request a confirmation and try to remove the volume. For that it will internally wait 10 seconds for the volume to be free. In fact, only a free volume can be removed.

If the volume cannot be removed in 10 seconds, an error message is displayed. At this time, the order to remove must be given again or a manual VOLUME:REMOVE command must be sent with a WAIT parameter set to a greater value than 10. See the command reference chapter for further information about the VOLUME:REMOVE command.

Changing volume parameters

For changing the volume parameters, just select the ▤ icon near the volume name in the volume list. This displays the following window:

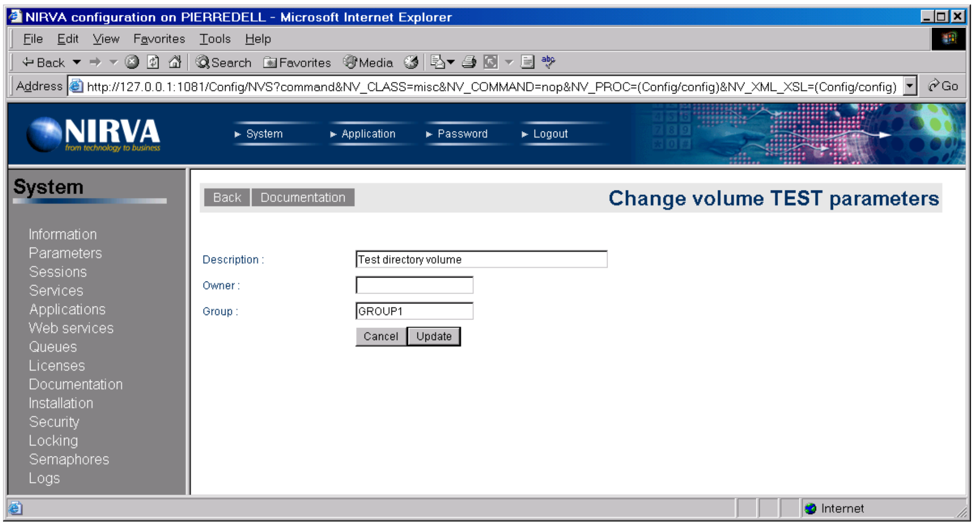
NIRVA configuration on PIERREDELL - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://127.0.0.1:1081/Config/NVS?command&NV_CLASS=misc&NV_COMMAND=nop&NV_PROC=(Config/config)&NV_XML_XSL=(Config/config)
Go
NIRVA
from technology to business
System
Application
Password
Logout
System
Information
Parameters
Sessions
Services
Applications
Web services
Queues
Licenses
Documentation
Installation
Security
Locking
Semaphores
Logs
Back
Documentation
Cancel
Update
Internet

Change volume TEST parameters

Description : Test directory volume

Owner :

Group : GROUP1

The volume name cannot be changed from this window. The only way to change the volume name is to remove the volume and to create it again.

The Owner parameter is reserved for future use.

The Group parameter tells if the volumes is inside a group. The group name must not contain space and is case independent. Please see the overview chapter for information about volume groups.

Setting volume groups

The volume groups can be defined when setting the volume parameters. See the previuos chapter.

Enabling or disabling a volume

This is simply done by clicking the red or green light just near the Available information in the volume list:

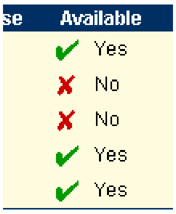

se	Available
	✓ Yes
	✗ No
	✗ No
	✓ Yes
	✓ Yes

Synchronizing a volume

This feature is available only to the volumes of storage type.

The synchronization process consists of copying all the documents from a source level to a destination level. Only the documents that are not found on the destination level are copied. The synchronization always occurs for a date range.

The synchronization is very useful for backup purposes or when working with a slow optical disk. At this time, the writing can occur on a magnetic disk and the synchronization is then launched during the night from the magnetic to the optical disk.

The synchronization can occur in a background task so the volume is still accessible in read and write mode during a synchronization process.

In order to start a volume synchronization, one just need to click on the synchronization icon ⇅ near the volume name in the volume list.

This then displays the following screen:

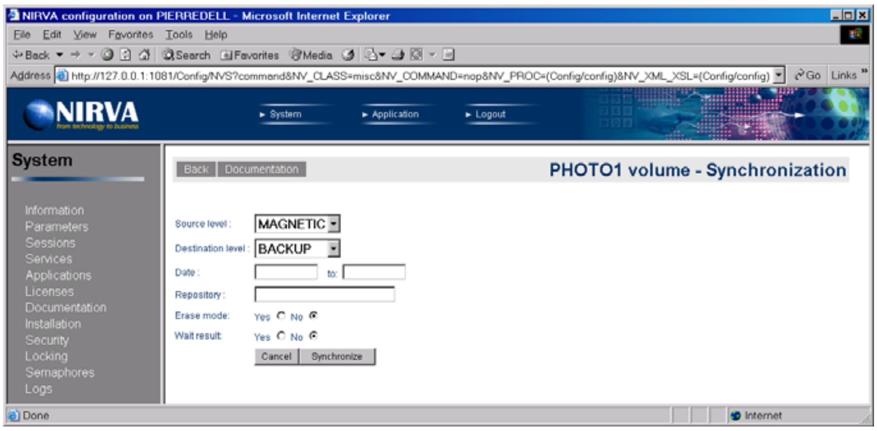

NIRVA configuration on PIERREDELL - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://127.0.0.1:1081/Config/NVS?command&NV_CLASS=misc&NV_COMMAND=nop&NV_PROC=(Config/config)&NV_XML_XSL=(Config/config)
Go Links
NIRVA
from technology to business
► System
► Application
► Logout
System
Information
Parameters
Sessions
Services
Applications
Licenses
Documentation
Installation
Security
Locking
Semaphores
Logs
Back
Documentation
Done
Internet

PHOTO1 volume - Synchronization

Source level : MAGNETIC

Destination level : BACKUP

Date : to:

Repository :

Erase mode: Yes ○ No ●

Wait result: Yes ○ No ●

The user must choose the volume source and destination levels. The default is the first level as source level and the second level as destination level.

The date range is mandatory. If the user doesn't enter it, the service is synchronizing the production of the day before the actual day. If only the first date is given, the service is synchronizing the production of the given day. The dates must be entered in the format dd/mm/yy or dd/mm/yyyy.

The repository, if given, allows to restrict the synchronization to the given repository.

Erase mode, when set to yes, allows to erase the documents from the source level. The removing of a document occurs only if the document is correctly copied on the destination level or if it ever exists on the destination level.

The wait result option allows for the user to wait for the synchronization to be completed. The default is to not wait so the synchronization occurs in background and the volume is then available for read and write commands.

Stopping a synchronization

This feature is available only to the volumes of storage type.

When a synchronization process has been launched in background, the configuration tool comes back to the volume list and displays the progression of the synchronization process in the "Synchro" column. The synchronization process can then be stopped by clicking on the icon near the progression information. If there is no icon in the "Synchro" columns, that means that the synchronization process has been finished.

Note: The web page is not refreshing itself automatically so the user must press the refresh button in order to actualize the synchronization progression information.

Level management

This feature is available only to the volumes of storage type.

The level management is accessible by directly clicking the volume name in the volume list.

This then displays the following screen:

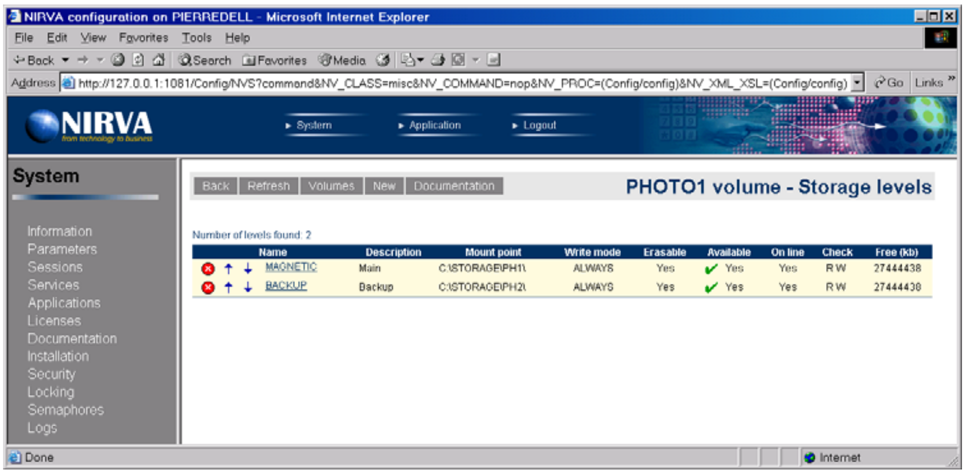
NIRVA configuration on PIERREDELL - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://127.0.0.1:1081/Config/NVS?command&NV_CLASS=misc&NV_COMMAND=nop&NV_PROC=(Config/config)&NV_XML_XSL=(Config/config)
Go Links
NIRVA
System
Application
Logout
System
Information
Parameters
Sessions
Services
Applications
Licenses
Documentation
Installation
Security
Locking
Semaphores
Logs
Back
Refresh
Volumes
New
Documentation
Done
Internet

PHOTO1 volume - Storage levels

Number of levels found: 2

	Name	Description	Mount point	Write mode	Erasable	Available	On line	Check	Free (kb)
✕ ↑ ↓	MAGNETIC	Main	C:\STORAGE\PH1\	ALWAYS	Yes	✓ Yes	Yes	R W	27444438
✕ ↑ ↓	BACKUP	Backup	C:\STORAGE\PH2\	ALWAYS	Yes	✓ Yes	Yes	R W	27444438

This is the list of actual levels defined for the volume. For each level, the list displays the following information:

- Name.
- Description.
- Path (mount point).
- Write mode.
- Erasable state.
- Available state.
- On line state.
- Check read / write mode.
- Free space.

Creating a new level

Creating a new level is done by clicking the button “New” on the level list screen. This displays the following screen:

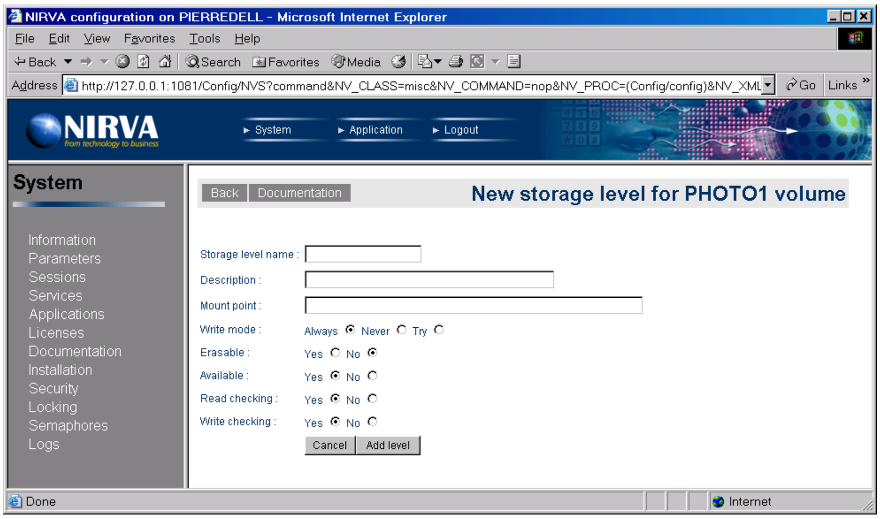
NIRVA configuration on PIERREDELL - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://127.0.0.1:1081/Config/NVS?command&NV_CLASS=misc&NV_COMMAND=nop&NV_PROC=(Config/config)&NV_XML
Go Links
NIRVA
from technology to business
System
Application
Logout
System
Information
Parameters
Sessions
Services
Applications
Licenses
Documentation
Installation
Security
Locking
Semaphores
Logs
Back
Documentation
Done
Internet

New storage level for PHOTO1 volume

Storage level name :

Description :

Mount point :

Write mode : Always ● Never ○ Try ○

Erasable : Yes ○ No ●

Available : Yes ● No ○

Read checking : Yes ● No ○

Write checking : Yes ● No ○

Cancel Add level

The level name is case independent. If the user attempts to create a level with the name of an existing one, the command fails. The level name is mandatory.

The level description is a simple string that describes the level. It's not mandatory.

The mount point is the physical mount point of the level. It's not necessary that the mount point exist for creating a level that points to it. At this time, the level will be created but it will not be on line. In any case, the level creation will not create the mount point directory so this one must be created manually. If the mount point is let blank, the service creates one in the "Volumes" directory of the storage service.

The level write mode defines the way to write new documents on the level. It can take the values "ALWAYS", "TRY" or "NEVER":

- ALWAYS means that NIRVA will always write the document on this level and will report an error if unsuccessful. The first level is often defined with the ALWAYS mode.
- TRY means that NIRVA will try to write the document on this level but will not report an error if not successful.
- NEVER means that NIRVA will not write the document during the WRITE command on this level. The level will then be written by using the level synchronization commands.

Erasable tells if the new level is erasable or not.

Available tells if the new level is available or not.

Read and write checking gives the level check mode respectively for read and write operations. Checking a level means verifying that it's mount point has a file named "volume.nvs" that contains the volume name. The checking can be enabled separately for read and write operations. It's recommended to use checking in both reading and writing mode in order to be sure to address the correct volume.

Removing a level

Removing a level is quiet simple. Just select the ⊗ icon near the level name in the level list. The configuration tool will then request a confirmation and try to remove the level (if not in use).

Changing the level order

The level order is important during read operations. When reading a document, the storage service first tries to get it on the first level, then on the second level and so on…

Changing a level order is made by selecting the up or down icons ↑ ↓ near the level name in the level list.

Changing level characteristics

This is simply done by clicking the level name in the level list. At this time, the following screen is displayed:

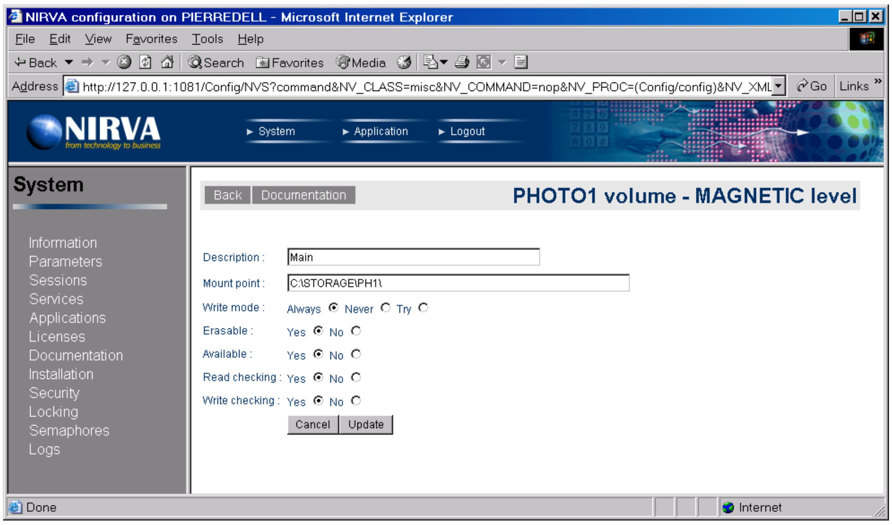
NIRVA configuration on PIERREDELL - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://127.0.0.1:1081/Config/NVS?command&NV_CLASS=misc&NV_COMMAND=nop&NV_PROC=(Config/config)&NV_XML
Go Links
NIRVA
from technology to business
System
Application
Logout
System
Information
Parameters
Sessions
Services
Applications
Licenses
Documentation
Installation
Security
Locking
Semaphores
Logs
Back
Documentation
Cancel
Update
Done
Internet

PHOTO1 volume - MAGNETIC level

Description : Main

Mount point : C:\STORAGE\PH1\

Write mode : Always ● Never ○ Try ○

Erasable : Yes ● No ○

Available : Yes ● No ○

Read checking : Yes ● No ○

Write checking : Yes ● No ○

All the level parameters can then be changed except its name.

The level description is a simple string that describes the level. It's not mandatory.

The mount point is the physical mount point of the level. It's not necessary that the mount point exist for creating a level that points to it. At this time, the level will be created but it will not be on line. In any case, the level creation will not create the mount point directory so this one must be created manually. If the mount point is let blank, the service creates one in the "Volumes" directory of the storage service.

The level write mode defines the way to write new documents on the level. It can take the values "ALWAYS", "TRY" or "NEVER".

- ALWAYS means that NIRVA will always write the document on this level and will report an error if unsuccessful. The first level is often defined with the ALWAYS mode.
- TRY means that NIRVA will try to write the document on this level but will not report an error if not successful.
- NEVER means that NIRVA will not write the document during the WRITE command on this level. The level will then be written by using the level synchronization commands.

Available tells if the new volume is available or not.

Read and write checking gives the level check mode respectively for read and write operations. Checking a level means verifying that it's mount point has a file named "volume.nvs" that contains the volume name. The checking can be enabled separately for read and write operations. It's recommended to use checking in both reading and writing mode in order to be sure to address the correct volume.

Enabling or disabling a level

This is simply done by clicking the red or green light just near the Available information in the level list:

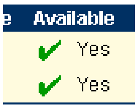

e Available
✓ Yes
✓ Yes

Directory management

This feature is available only to the volumes of directory type.

The directory management is accessible by directly clicking the volume name in the volume list.

This then displays the following screen:

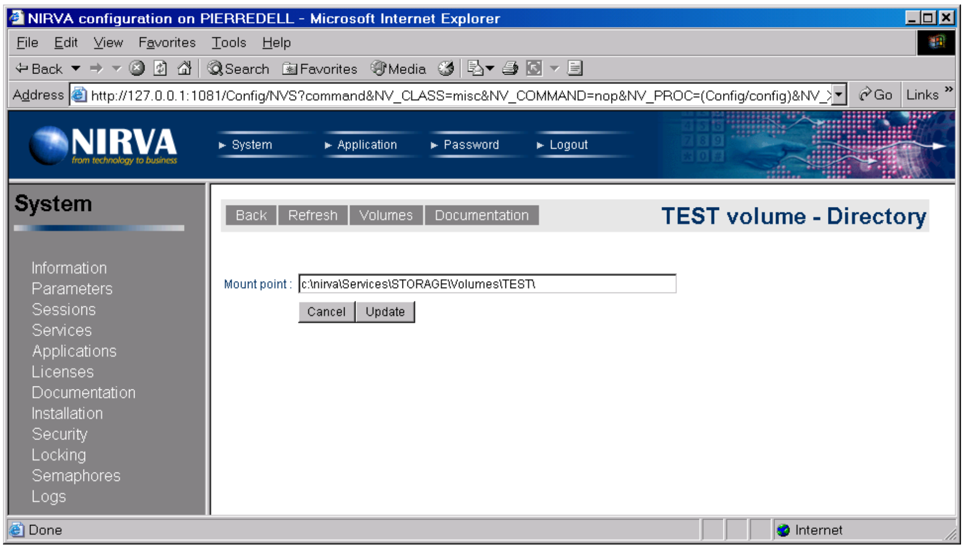
NIRVA configuration on PIERREDELL - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://127.0.0.1:1081/Config/NVS?command&NV_CLASS=misc&NV_COMMAND=nop&NV_PROC=(Config/config)&NV_
Go Links
NIRVA from technology to business
System
Application
Password
Logout
System
Information
Parameters
Sessions
Services
Applications
Licenses
Documentation
Installation
Security
Locking
Semaphores
Logs
Back
Refresh
Volumes
Documentation
Cancel
Update
Done
Internet

TEST volume - Directory

Mount point : c:\nirva\Services\STORAGE\Volumes\TEST\

The user can change here the mount point for a volume directory. The mount point must exist. After creating a volume the service creates a default mount point into the “Volumes” directory of the storage service.

Reference

This chapter gives the complete reference of all the storage service commands.

Classes

Here are the available STORAGE service classes:

STORAGE	Main class.
VOLUME	Volume management.
LEVEL	Volume levels management.
DOCUMENT	Document commands.

Error codes

COMMAND Class

Value	Description
101	Invalid command
102	Invalid command parameter
103	Bad date format

DOCUMENT Class

Value	Description
101	Invalid document reference
102	Cannot read the document
103	Cannot write the document

Value	Description
104	Cannot remove the document
105	Cannot copy a document part
106	Cannot get document size

LEVEL Class

Value	Description
101	The level ever exists
102	Invalid level name
103	Invalid level mount point
104	Another level with the same mount point exist
105	The level is not available

STORAGE Class

Value	Description
101	Memory error (not enough)

VOLUME Class

Value	Description
101	The volume ever exists
102	Cannot load the volume list from registry
103	Cannot load the volume information from registry
104	Invalid volume name
105	Cannot get volume (probably in use)
106	The volume is not available
107	The repository name must have at least 3 characters
108	The synchronization has failed
109	The volume is in synchronization process
110	Cannot save the volume information to registry
111	Operation not available for the volume type
112	Invalid group name
113	No available volume found for this group

Permissions

Name	Description
VOLUME_ADMIN	Administrate volumes
DOCUMENT_READ	Read documents
DOCUMENT_WRITE	Write documents
DOCUMENT_REMOVE	Remove documents

Commands

For each command, the reference gives the command name, the sources for which the command may be used, the command description, the eventual command permissions, the parameter list and the eventual list of objects created by the command.

The parameters described in this chapter are command specific parameters. For general parameters, please refer to the Nirva command syntax chapter.

The available sources are:

- Client for all Nirva client interfaces including Nirva client library (nvc).
- Web for commands from a web browser.
- Procedure for commands from a Nirva procedure.
- Service for commands from service to service.

DOCUMENT class

The LEVEL class provides commands for reading and writing documents on a volume.

When a document has been stored on a volume, it’s identified by a reference. Please see the overview chapter for a description of the document reference.

COPY

STORAGE:DOCUMENT:COPY

Source	Use Input Container	Use Output Container

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	Yes

Description

This command writes a document from a source to a destination volume.

In case where the document to copy is a container, the command only allows to copy it as a flat file so it's not possible to copy only a part of a container document with the COPY command.

The physical writing occurs on all volume storage levels having a write mode set to “ALWAYS” or “TRY”. If the level write mode is “TRY”, the command will not fail if the writing is not successful.

In any case, the command will be successful only if the document has been written at least on one storage level.

The command creates an object named REFERENCE in the output container that contains the new document reference.

Permissions

DOCUMENT_READ

DOCUMENT_WRITE

Parameters

NAME

Destination volume name.
One of the NAME or GROUP parameter is mandatory.

GROUP

Destination group name.
One of the NAME or GROUP parameter is mandatory.
Several volumes can be grouped together for writing purpose. When writing a document, one can give the name of a group instead of the volume name. At this time, the service automatically chooses one of the volumes of the group to write on. After each group write, the service marks it as the last written volume and simply chooses the next from the volume list of the group when the next write order arrives.

WAIT

Wait time in seconds for the source and destination volumes to be free. For this command, the STORAGE service must wait for the volumes to be free. This parameters allows to define the maximum wait time. If the volumes are still in use after the wait time, the command fails.
The default is “0” (infinite wait).

REF

Document reference as returned by the write command. If the REF parameter points to a specific object of the document, the command fails.

REPOS

Repository name. This parameter allows to copy the document in the specified repository. A repository is simply a subdirectory of the storage media. A repository name is case insensitive, it cannot contain any space characters (or special characters used for directory names) and must have at least 3 characters.

If the REPOS parameter is not given, the command copies the document in the default repository (blank repository).

PREFIX

When this parameter is provided, the service uses its content as prefix for the physical file name on which the document is stored. The default is no prefix.

For a destination volume of directory type, this parameter has meaning only when the FILENAME parameter is not given.

The PREFIX parameter cannot contain the characters ‘.’, ‘/’, ‘\’, ‘-‘, ‘_’ or space. The service doesn’t generate errors if one of these characters is found but automatically eliminates it.

See the overview chapter for a description of the document file names.

FILENAME

This has meaning only for a destination volume of directory type. When this parameter is provided, the service uses it as the physical file name for the stored object. If the file ever exists, the service may generate an error if the REPLACE parameter is not set to “YES”. When the FILENAME parameter is not given, the service automatically generates it following information given in PREFIX, SUBDIR and EXTENSION parameters.

The FILENAME parameter cannot contain the characters ‘.’, ‘/’, ‘\’, ‘-‘, ‘_’ or space. The service doesn’t generate errors if one of these characters is found but automatically eliminates it.

See the overview chapter for a description of the document file names.

REPLACE

This has meaning only for a destination volume of directory type when the FILENAME parameter has been given. If the given filename exists, the REPLACE parameter controls if the service must replace it or generate an error. The possible values are “YES” and “NO”. The default is “NO” (generates an error if the file ever exists).

SUBDIR

This has meaning only for a destination volume of directory type when the FILENAME parameter has not been given. This is the subdirectory on which the service will write the document.

The SUBDIR parameter cannot contain the characters ‘.’, ‘-‘, ‘_’, or space. The service doesn’t generate errors if one of these characters is found but automatically eliminates it.

See the overview chapter for a description of the document file names.

EXTENSION

This has meaning only for a destination volume of directory type when the FILENAME parameter has not been given. This is the optional extension of the calculated filename.

The EXTENSION parameter cannot contain the characters ‘/’, ‘\’, ‘-‘, ‘_’ or space. The service doesn’t generate errors if one of these characters is found but automatically eliminates it.

See the overview chapter for a description of the document file names.

OFFSET

This parameter gives the offset in bytes of the source file to store. Combined with the *NUM_BYTES* parameter, it allows to store only a part of the file.
Since this parameter can be used while the source document is a container, it's dangerous to use it because the document structure is then not preserved.
The default value is ”0” (beginning of file).

NUM_BYTES

This parameter gives the number of bytes of the file to store. Combined with the OFFSET parameter, it allows to store only a part of the file.
Since this parameter can be used while the source document is a container, it's dangerous to use it because the document structure is then not preserved.
The default value is ”-1” (until the end of file).

Objects created

REFERENCE

This is a Nirva string object containing the document reference. The document reference is used to retrieve the document with the DOCUMENT:READ command. Please see the overview chapter for a complete description of the reference.

READ

STORAGE:DOCUMENT:READ

Source	Use Input Container	Use Output Container
Client Web Procedure Service	Yes	Yes

Description

This command reads a document or a part of it from the given volume. The document can be a flat file or a complete NIRVA container.

The result is written in the output container. If the document is a container, the reading is in fact an importation of the document into the output container (see the CONTAINER:IMPORT command in the NIRVA SYSTEM service in the NIRVA user's guide).

When reading a document, the storage service first tries to get it on the first storage level, then on the second level and so on…

Despite the number of it's parameters, this command is quiet simple to use because many of the parameters are optional. Generally only 2 of them will be used: REF (reference) and eventually REPLACE (replace mode).

Permissions

DOCUMENT_READ

Parameters

WAIT

Wait time in seconds for the volume to be free. For this command, the STORAGE service must wait for the volume to be free. This parameters allows to define the maximum wait time. If the volume is still in use after the wait time, the command fails.
The default is “0” (infinite wait).

REF

Document reference as returned by the write command.
If the document is a container, the REF parameter can point to a specific object of this container. At this time, the reference must be followed by a ‘_’ character (underscore) and by an object name (*docpart*). The docpart itself may point to an object of a subcontainer. For example, “_sub1.sub2.obj1” will point to the object “obj1” of the subcontainer “sub2” of the subcontainer “sub1”.
Always in the case where the document is a container, it’s possible to get it as a flat file by removing one of the 2 ‘-‘ characters (‘--‘) separating the volume name from the physical file in the reference.
See the overview chapter for further information about the reference.

FLAT

Force flat file.
The service knows that it has to access a flat file by checking if the REF parameter contains 2 ‘-‘ characters (‘--‘) as separator between the volume name and the physical file. If there is only one ‘-‘ character, this is a flat file access.
Now, when set to “YES”, the FLAT parameter allows to force to access a flat file in any case.

FILE

This parameter is used only for a flat file. It’s the name of a NIRVA file object that will contains the retrieved document. This object is created by the command in the output container if it doesn’t exist.
The default value is “FILE”.

APPEND

This parameter is used only for a container. If this parameter is set to “YES”, the command will append the document objects and subcontainers to the output container. In this case, the eventual replacement of existing objects is controlled by the parameter REPLACE.
When not in append mode, the output container is first released before the document is copied to it. This is the default.

REPLACE

When used with a container, this parameter has a meaning only in append mode (APPEND parameter set to “YES”). If REPLACE is set to “YES” all document objects will replace output container objects having the same name.
The default is “NO” so the container objects having the same name than the document objects will not be replaced.

When used with a flat file, and this parameter is set to “YES”, the service replaces the result file object with the same name if it exists. Otherwise, the service sends an error message if the result file object ever exist. The default is “NO”.

PERSIST

This parameter gives the new persistent value for persistent file objects that have to be imported in the output container. By default, the new file objects will not be persistent.

The parameter can take value “0” for temporary files, “-1” for persistent files and any other value (in seconds) for cached files.

For a flat file, it gives the new persistent value of the result file object.

FILES

This parameter is used only for a container. If it’s set to “YES”, it tells the service to also import files with file objects. If this parameter is set to “NO”, the file objects are imported but not the file they point to.

When files are really imported, the service automatically creates the new files in the application file directory if the file object is persistent and in the application work directory otherwise. This destination directory can be changed by the FILEDIR parameter if the command has been sent from a procedure or from an external service.

The default value is “YES”.

FILEDIR

This parameter is used only for a container. This parameter has meaning only if the FILES parameter has been set to “YES”. It allows to change the destination directory for files of file objects.

If the FILEDIR is set to the application work directory and the PERSIST parameter is set to “-1” (persistent), the command will return an error because the application work directory should not contain any persistent file.

The FILEDIR parameter can be used only if the command source is procedure or service. Otherwise, it’s ignored.

OBJECT

This parameter is used only for a container. If this parameter is not blank, it must contain the name of a single object to import from the import file to the output container.

The single object may point to a subcontainer of the import file. At this time it must contain all the path to the requested object. For example: `“Sub1.Sub11.MyObject”` tells the service to get only the object `“MyObject”` of the subcontainer `“Sub11”` of the subcontainer `“Sub1”`. The subcontainer and object names are case insensitive.

This parameter doesn’t have priority on the eventual object name given in the document reference.

FILENAME

This parameter is used only for a flat file. It gives the file name of the resulting file object.

If the command comes from client or browser, the FILENAME parameter is automatically reset to “” (blank).

If this parameter is not provided (or is blank), the service creates itself a file name following information of the EXTENSION, PREFIX, SUFFIX and DIRECTORY parameters.

DIRECTORY	This parameter is used only for a flat file. It's the name of a directory where the service will write the resulting file object if the FILENAME parameter is not given. For information, the service generates an error message when trying to create a persistent file in the application work directory. If this parameter is not provided or is blank, the service uses the application work directory for temporary and cached files and the application file directory for persistent files.
EXTENSION	This parameter is used only for a flat file. This is the file extension to use for the resulting file object. If this parameter is not provided or is blank, the service uses ".obj" for persistent files and ".tmp" for temporary and cached files. If the point character is omitted in the EXTENSION parameter, the service adds it.
PREFIX	This parameter is used only for a flat file. This is the file prefix to use for the resulting file object. The prefix, if provided is used by the service to automatically creates the resulting object file name.
SUFFIX	This parameter is used only for a flat file. This is the file suffix to use for the resulting file object. The suffix, if provided is used by the service to automatically creates the resulting object file name.
OFFSET	Offset. When the document is a container, this parameter is used only for file and binary objects to import. It gives the offset in bytes from which the copy will start. The default value is "0" (beginning of file or binary data).
NUM_BYTES	Number of bytes to copy. When the document is a container, this parameter is used only for file and binary objects to import. It gives the number of bytes of the object to import. The default value is "-1" (until the end of file or binary data).

REMOVE

STORAGE:DOCUMENT:REMOVE

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	No

Description

This command removes a document or a part of it from the given volume. The document is physically removed from all on line levels that are erasable.

After erasing the document, the command eventually removes the empty directories.

Permissions

DOCUMENT_REMOVE

Parameters

WAIT

Wait time in seconds for the volume to be free. For this command, the STORAGE service must wait for the volume to be free. This parameters allows to define the maximum wait time. If the volume is still in use after the wait time, the command fails.
The default is “0” (infinite wait).

REF

Document reference as returned by the write command.
If the document is a container, the REF parameter can point to a specific object of this container. At this time, the reference must be followed by a ‘_’ character (underscore) and by an object name (*docpart*). The docpart itself may point to an object of a subcontainer. For example, “_sub1.sub2.obj1” will point to the object “obj1” of the subcontainer “sub2” of the subcontainer “sub1”.
Always in the case where the document is a container, it’s possible to get it as a flat file by removing one of the 2 ‘-‘ characters (‘--‘) separating the volume name from the physical file in the reference.
See the overview chapter for further information about the reference.
If a specific object is given, only this object will be removed. If after the removing of the object, the document is empty, it’s physically removed.

OBJECT

This parameter is used only for a container. If this parameter is not blank, it must contain the name of a single document object to remove.
The single object may point to a subcontainer of the document. At this time it must contain all the path to the requested object. For example: `“Sub1.Sub11.MyObject”` tells the service to remove only the object `“MyObject”` of the subcontainer `“Sub11”` of the subcontainer `“Sub1”`. The subcontainer and object names are case insensitive.
This parameter doesn’t have priority on the eventual object name given in the document reference.

SIZE

STORAGE:DOCUMENT:SIZE

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	Yes

Description

This command returns the size of a document or a part of it if a single object of a container document is requested:

- For a flat file document, the size is the file size in bytes.
- For a container document, the size is the file size in bytes.
- For a specific object of a container document, the size depends of the object type:
- For a boolean object the size is always 1.
- For an integer object the size is the size in bytes of a C++ integer (generally 4).
- For a string object, the size is the string length.
- For a string list object, the size is the list size.
- For an indexed string list object, the size is the list size.
- For an table object, the size is the number of rows.
- For an file object, the size is the file length. The IMPORT command allows to not really import files if requested (see the FILES parameters). At this time, the size information is also the real file length. This allows to use the import function just to require the length of a file embedded in the importation file. In the same way, if only a part of the files has to be imported (see OFFSET and NUM_BYTES parameters), the returned size is still the real file size.
- For a binary object, the size is the size of the binary data.

The size information is returned into an object named SIZE in the output container.

Permissions

DOCUMENT_READ

Parameters

WAIT	Wait time in seconds for the volume to be free. For this command, the STORAGE service must wait for the volume to be free. This parameters allows to define the maximum wait time. If the volume is still in use after the

wait time, the command fails.
The default is “0” (infinite wait).

REF

Document reference as returned by the write command.
If the document is a container, the REF parameter can point to a specific object of this container. At this time, the reference must be followed by a ‘_’ character (underscore) and by an object name (*docpart*). The docpart itself may point to an object of a subcontainer. For example, “_sub1.sub2.obj1” will point to the object “obj1” of the subcontainer “sub2” of the subcontainer “sub1”.
Always in the case where the document is a container, it’s possible to get it as a flat file by removing one of the 2 ‘-‘ characters (‘--‘) separating the volume name from the physical file in the reference.
See the overview chapter for further information about the reference.

FLAT

Force flat file.
The service knows that it has to access a flat file by checking if the REF parameter contains 2 ‘-‘ characters (‘--‘) as separator between the volume name and the physical file. If there is only one ‘-‘ character, this is a flat file access.
Now, when set to “YES”, the FLAT parameter allows to force to access a flat file in any case.

OBJECT

This parameter is used only for a container. If this parameter is not blank, it must contain the name of a valid single object of the container document.
The single object may point to a subcontainer of the document. At this time it must contain all the path to the requested object. For example: `“Sub1.Sub11.MyObject”` tells the service to get only the object `“MyObject”` of the subcontainer `“Sub11”` of the subcontainer `“Sub1”`.
The subcontainer and object names are case insensitive.
This parameter doesn’t have priority on the eventual object name given in the document reference.

Objects created

SIZE

Document size. This is a NIRVA integer object. See the description of the command for information about the document size.

WRITE

STORAGE:DOCUMENT:WRITE

Source	Use Input Container	Use Output Container

Source	Use Input Container	Use Output Container
Client Web Procedure Service	Yes	Yes

Description

This command writes a document on the given volume. The document can be a flat file or a complete NIRVA container (the input container).

For a volume of storage type, the physical writing occurs on all volume storage levels having a write mode set to “ALWAYS” or “TRY”. If the level write mode is “TRY”, the command will not fail if the writing is not successful. In any case, the command will be successful only if the document has been written at least on one storage level.

The command creates an object named REFERENCE in the output container that contains the new document reference.

Permissions

DOCUMENT_WRITE

Parameters

NAME
Volume name.
One of the NAME or GROUP parameter is mandatory.

GROUP
Group name.
One of the NAME or GROUP parameter is mandatory.
Several volumes can be grouped together for writing purpose. When writing a document, one can give the name of a group instead of the volume name. At this time, the service automatically chooses one of the volumes of the group to write on. After each group write, the service marks it as the last written volume and simply chooses the next from the volume list of the group when the next write order arrives.

WAIT
Wait time in seconds for the volume to be free. For this command, the STORAGE service must wait for the volume to be free. This parameter allows to define the maximum wait time. If the volume is still in use after the wait time, the command fails.
The default is “0” (infinite wait).

FLAT
This parameter must be set to “NO” in order to write the input container instead of a given flat file.
The default is “YES” (flat file).

REPOS
Repository name. This parameter allows to write the document in the specified repository. A repository is simply a subdirectory of the storage

media. A repository name is case insensitive, it cannot contain any space characters (or special characters used for directory names) and must have at least 3 characters.

If the REPOS parameter is not given, the command writes the document in the default repository (blank repository).

PREFIX

When this parameter is provided, the service uses its content as prefix for physical file name on which the document is stored. The default is no prefix.

For a volume of directory type, this parameter has meaning only when the FILENAME parameter is not given.

The PREFIX parameter cannot contain the characters ‘.’, ‘/’, ‘\’, ‘-‘, ‘_’ or space. The service doesn’t generate errors if one of these characters is found but automatically eliminates it.

See the overview chapter for a description of the document file names.

FILENAME

This has meaning only for a volume of directory type. When this parameter is provided, the service uses it as the physical file name for the stored object. If the file ever exists, the service may generate an error if the REPLACE parameter is not set to “YES”. When the FILENAME parameter is not given, the service automatically generates it following information given in PREFIX, SUBDIR and EXTENSION parameters.

The FILENAME parameter cannot contain the characters ‘.’, ‘/’, ‘\’, ‘-‘, ‘_’ or space. The service doesn’t generate errors if one of these characters is found but automatically eliminates it.

See the overview chapter for a description of the document file names.

REPLACE

This has meaning only for a volume of directory type when the FILENAME parameter has been given. If the given filename exists, the REPLACE parameter controls if the service must replace it or generate an error. The possible values are “YES” and “NO”. The default is “NO” (generates an error if the file ever exists).

SUBDIR

This has meaning only for a volume of directory type when the FILENAME parameter has not been given. This is the subdirectory on which the service will write the document.

The SUBDIR parameter cannot contain the characters ‘.’, ‘-‘, ‘_’, or space. The service doesn’t generate errors if one of these characters is found but automatically eliminates it.

See the overview chapter for a description of the document file names.

EXTENSION

This has meaning only for a volume of directory type when the FILENAME parameter has not been given. This is the optional extension of the calculated filename.

The EXTENSION parameter cannot contain the characters ‘/’, ‘\’, ‘-‘, ‘_’ or space. The service doesn’t generate errors if one of these characters is found but automatically eliminates it.

See the overview chapter for a description of the document file names.

FILE

This parameter is used only for a flat file. It’s the name of a NIRVA file object that contains the file to store. This object must exist in the input container

	and must point to a real file. The default value is “FILE”.
OFFSET	This parameter is used only for a flat file. It gives the offset in bytes of the flat file to store. Combined with the *NUM_BYTES* parameter, it allows to store only a part of the file. The default value is ”0” (beginning of file).
NUM_BYTES	This parameter is used only for a flat file. It gives the number of bytes of the flat file to store. Combined with the OFFSET parameter, it allows to store only a part of the file. The default value is ”-1” (until the end of file).
SUBCONTAINERS	This parameter is used only when storing the input container. It allows to also write the subcontainers of the input container. To write subcontainers, the parameter must be set to “YES”. The default value is “NO”.
OBJECTS	This parameter is used only when storing the input container. It allows to specify only some of the input container objects to store. If used, the parameter should contain the names of valid input container objects, separated by semicolon character (‘;’). The default is to include all input container objects.

Objects created

REFERENCE	This is a Nirva string object containing the document reference. The document reference is used to retrieve the document with the DOCUMENT:READ command. Please see the overview chapter for a complete description of the reference.

LEVEL class

The LEVEL class provides commands for managing volume storage levels. One can create, remove, enable, disable or list levels.

Each volume of storage type has its collection of storage levels. The storage level is the physical place where the document is stored.

The LEVEL class command are available only for volumes of storage type except the LIST command that can also be used with a volume of directory type (this is only useful when the free space information has to be retrieved).

CHANGE

STORAGE:LEVEL:CHANGE

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	No

Description

This command is available only for a volume of storage type.

This command changes the storage level parameters. The storage level must exist.

Any parameter can be changed except the level name. If a parameter is not provided, it's value is unchanged.

It's not possible to change storage level parameters while it's in a synchronization operation. At this time, the command fails.

Permissions

VOLUME_ADMIN

Parameters

NAME

Volume name.
The NAME parameter is mandatory.

WAIT

Wait time in seconds for the volume to be free. For this command, the STORAGE service must wait for the volume to be free. This parameters allows to define the maximum wait time. If the volume is still in use after the wait time, the command fails.
The default is "0" (infinite wait).

LNAME

Level name. The level name is case independent. The level name must correspond to an existing storage level. The level name is mandatory.

DESCRIPTION

Level description. This is just a string that describes the level. If this parameter is not provided, the actual value is not changed.

PATH

Physical mount point of the level. It's not necessary that the mount point exist for a level to point to it. At this time, the level will not be on line. In any case, the level creation will not create the mount point directory so this one must be created manually.
If this parameter is not provided, the actual value is not changed.

WRITE

Level write mode. The level write mode defines the way to write new documents on the level. It can take the values "ALWAYS", "TRY" or

	"NEVER". ALWAYS means that NIRVA will always write the document on this level and will report an error if unsuccessful. The first level is often defined with the ALWAYS mode. TRY means that NIRVA will try to write the document on this level but will not report an error if not successful. NEVER means that NIRVA will not write the document during the WRITE command on this level. The level will then be written by using the level synchronization commands. If this parameter is not provided, the actual value is not changed.
ERASABLE	Tells if the level is erasable or not. If set to "YES", the level will be erasable. If this parameter is not provided, the actual value is not changed.
AVAILABLE	Tells if the new level is available or not. If set to "YES", the level will be available. If set to "NO", it will not be available. If a level is not available, it's still in the list but is not accessible. If this parameter is not provided, the actual value is not changed.
CHECKR	Read check mode. This parameter gives the level check mode for read operations. Checking a level means verifying that it's mount point has a file named "volume.nvs" that contains the volume name. The checking can be enabled separately for read and write operations. It's recommended to use checking in both reading and writing mode in order to be sure to address the correct volume. This parameter can take values "YES" or "NO" for respectively enabling or disabling read checking. If this parameter is not provided, the actual value is not changed.
CHECKW	Write check mode. This parameter gives the level check mode for write operations. Checking a level means verifying that it's mount point has a file named "volume.nvs" that contains the volume name. The checking can be enabled separately for read and write operations. It's recommended to use checking in both reading and writing mode in order to be sure to address the correct volume. This parameter can take values "YES" or "NO" for respectively enabling or disabling write checking. If this parameter is not provided, the actual value is not changed.

CLEAN

STORAGE:LEVEL:CLEAN

Source	Use Input Container	Use Output Container

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	Yes

Description

This command is available only for a volume of storage type.

This command is used to check and remove empty subdirectories for a given level and date range.

Generally, the removing of empty directories is also assumed by the DOCUMENT:REMOVE and VOLUME:SYNCHRONIZE commands so this command will be useful only in case of external removing of some level documents.

Permissions

VOLUME_ADMIN

Parameters

NAME

Volume name.
The NAME parameter is mandatory.

WAIT

Wait time in seconds for the volume to be free. For this command, the STORAGE service must wait for the volume to be free. This parameters allows to define the maximum wait time. If the volume is still in use after the wait time, the command fails.
The default is “0” (infinite wait).

LNAME

Level name. The level name is case independent.
If the level name is not given or is blank, the command uses the first level.

DATE

Starting date. The DATE parameter is the starting date. It must be on the format dd/mm/yy or dd/mm/yyyy. If the DATE parameter is not provided, the service uses the date of the previous day.

TO_DATE

Ending date. The TO_DATE parameter is the ending date. It must be on the format dd/mm/yy or dd/mm/yyyy. If the TO_DATE parameter is not provided, the service uses the same value than the DATE parameter.

EPOS

Repository name. This parameter allows to restrict the cleaning to a given repository.

CREATE

STORAGE:LEVEL:CREATE

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	No

Description

This command is available only for a volume of storage type.

This command creates a storage level for the given volume. If the storage level ever exists, the command fails.

Permissions

VOLUME_ADMIN

Parameters

NAME Volume name.
The NAME parameter is mandatory.

WAIT Wait time in seconds for the volume to be free. For this command, the STORAGE service must wait for the volume to be free. This parameters allows to define the maximum wait time. If the volume is still in use after the wait time, the command fails.
The default is “0” (infinite wait).

LNAME Level name. The level name is case independent. If the user attempts to create a level with the name of an existing one, the command fails. The level name is mandatory.

DESCRIPTION Level description. This is just a string that describes the level. It’s not mandatory.

PATH Physical mount point of the level. It’s not necessary that the mount point exist for creating a level that points to it. At this time, the level will be created but it will not be on line. In any case, the level creation will not create the mount point directory so this one must be created manually.
This parameter is mandatory.

WRITE Level write mode. The level write mode defines the way to write new documents on the level. It can take the values “ALWAYS”, “TRY” or “NEVER”.
ALWAYS means that NIRVA will always write the document on this level and

will report an error if unsuccessful. The first level is often defined with the ALWAYS mode. The ALWAYS mode is the default.

TRY means that NIRVA will try to write the document on this level but will not report an error if not successful.

NEVER means that NIRVA will not write the document during the WRITE command on this level. The level will then be written by using the level synchronization commands.

ERASABLE

Tells if the new level is erasable or not. If set to “YES”, the level will be erasable. The default is “NO”.

AVAILABLE

Tells if the new level is available or not. If set to “YES”, the level will be available. If set to “NO”, it will not be available. If a level is not available, it’s still in the list but is not accessible. The default is “YES”.

CHECKR

Read check mode. This parameter gives the level check mode for read operations. Checking a level means verifying that it’s mount point has a file named “volume.nvs” that contains the volume name. The checking can be enabled separately for read and write operations. It’s recommended to use checking in both reading and writing mode in order to be sure to address the correct volume.

This parameter can take values “YES” or “NO” for respectively enabling or disabling read checking.

The default is “YES”.

CHECKW

Write check mode. This parameter gives the level check mode for write operations. Checking a level means verifying that it’s mount point has a file named “volume.nvs” that contains the volume name. The checking can be enabled separately for read and write operations. It’s recommended to use checking in both reading and writing mode in order to be sure to address the correct volume.

This parameter can take values “YES” or “NO” for respectively enabling or disabling write checking.

The default is “YES”.

DOC_LIST

STORAGE:LEVEL:DOC_LIST

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	Yes

Description

This command is available only for a volume of storage type.

This command returns the list of documents for a given level and date range.

The command creates a table object returning following information for each document:

- Document reference.
- Kind of document (flat file or container).
- Document size in bytes.

Permissions

VOLUME_ADMIN

Parameters

NAME

Volume name.
The NAME parameter is mandatory.

WAIT

Wait time in seconds for the volume to be free. For this command, the STORAGE service must wait for the volume to be free. This parameters allows to define the maximum wait time. If the volume is still in use after the wait time, the command fails.
The default is “0” (infinite wait).

LNAME

Level name. The level name is case independent.
If the level name is not given or is blank, the command uses the first level.

DATE

Starting date. The DATE parameter is the starting date. It must be on the format dd/mm/yy or dd/mm/yyyy. If the DATE parameter is not provided, the service uses the date of the previous day.

TO_DATE

Ending date. The TO_DATE parameter is the ending date. It must be on the format dd/mm/yy or dd/mm/yyyy. If the TO_DATE parameter is not provided, the service uses the same value than the DATE parameter.

REPOS

Repository name. This parameter allows to restrict the list to a given repository. If it’s not provided, all the production of all repositories is returned.

Objects created

DOC_LIST

This is a Nirva table object containing following columns:

- “REFERENCE” is the document reference. The document reference is used in the DOCUMENT:READ command.
- “TYPE is the document type. It takes the value “CONTAINER” for a container and “FLAT” for a flat file.
- “SIZE” is the document size in bytes.

DOWN

STORAGE:LEVEL:DOWN

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	No

Description

This command is available only for a volume of storage type.

This command changes the order of the storage level list associated to the volume.

The storage level order is important for the read operations. When reading, the service first tries on the first level, then on the second one and so on…

The DOWN command moves down the given storage level in the level list. If the level is ever at the bottom of the list, the DOWN command does nothing.

Permissions

VOLUME_ADMIN

Parameters

NAME	Volume name. The NAME parameter is mandatory.
WAIT	Wait time in seconds for the volume to be free. For this command, the STORAGE service must wait for the volume to be free. This parameters allows to define the maximum wait time. If the volume is still in use after the wait time, the command fails. The default is “0” (infinite wait).
LNAME	Level name. The level name is case independent. The level name is mandatory.

LIST

STORAGE:LEVEL:LIST

Source	Use Input Container	Use Output Container

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	Yes

Description

This command returns the list of levels for a given volume.

The command creates a table object returning complete level information:

- Name.
- Description.
- Path (mount point).
- Write mode.
- Erasable state.
- Available state.
- On line state.
- Check read mode.
- Check write mode.
- Free space.

This command is also available for a volume of directory type. In this case, the only significant information is path and free space.

Permissions

VOLUME_ADMIN

Parameters

NAME

Volume name.
The NAME parameter is mandatory.

WAIT

Wait time in seconds for the volume to be free. For this command, the STORAGE service must wait for the volume to be free. This parameters allows to define the maximum wait time. If the volume is still in use after the wait time, the command fails.
The default is “0” (infinite wait).

Objects created

LEVEL_LIST

This is a Nirva table object containing following columns:

- “NAME” is the level name.
- “DESCRIPTION is the level description.
- “PATH” is the level physical mount point.
- “WMODE” is the write mode for the level. It can take values “ALWAYS”, “TRY” or “NEVER”.
- “ERASABLE” is set to “YES” if the level is erasable and to “NO” otherwise.
- “AVAILABLE” is set to “YES” if the level is enabled and to “NO” if the level is disabled.
- “ONLINE” is set to “YES” if the level is on line and to “NO” otherwise. An on line volume means that the mount point is accessible and has a file named “volume.nvs” that itself contains the name of the volume.
- “CHECKR” is set to “YES” if the level is to be checked for each read access and to “NO” otherwise.
- “CHECKW” is set to “YES” if the level is to be checked for each write access and to “NO” otherwise.
- “FREE” gives the level free space in kilobytes. If the volume is not on line, the free space is always 0.

REMOVE

STORAGE:LEVEL:REMOVE

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	No

Description

This command is available only for a volume of storage type.

This command removes the given storage level.

In any case, the command doesn’t remove any data physically stored on the storage level.

After removing, the level can be mounted again by using the LEVEL:CREATE command.

It’s not possible to remove a level while it’s in a synchronization operation. At this time, the command fails.

Permissions

VOLUME_ADMIN

Parameters

NAME	Volume name. The NAME parameter is mandatory.
WAIT	Wait time in seconds for the volume to be free. For this command, the STORAGE service must wait for the volume to be free. This parameters allows to define the maximum wait time. If the volume is still in use after the wait time, the command fails. The default is “0” (infinite wait).
LNAME	Level name. The level name is case independent. The level name is mandatory.

UP

STORAGE:LEVEL:UP

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	No

Description

This command is available only for a volume of storage type.

This command changes the order of the storage level list associated to the volume.

The storage level order is important for the read operations. When reading, the service first tries on the first level, then on the second one and so on…

The UP command moves up the given storage level in the level list. If the level is ever at the top of the list, the UP command does nothing.

Permissions

VOLUME_ADMIN

Parameters

NAME
Volume name.
The NAME parameter is mandatory.

WAIT
Wait time in seconds for the volume to be free. For this command, the STORAGE service must wait for the volume to be free. This parameters allows to define the maximum wait time. If the volume is still in use after the wait time, the command fails.
The default is “0” (infinite wait).

LNAME
Level name. The level name is case independent.
The level name is mandatory.

STORAGE class

This is the standard service class that provides just one command for testing purpose.

NOP

STORAGE:STORAGE:NOP

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	No

Description

This command does nothing but allows to test that the storage service is on line and answers correctly.

If the service is not on line, this command returns an error.

Parameters

None

VOLUME class

The VOLUME class provides commands for managing volumes. One can create, remove, enable, disable or list volumes.

A volume is the logical place where the storage service stores a document.

There are two kinds of volumes: directory or storage volume:

- The directory volume has a direct connection to a single physical mount point (directory). It allows direct control of the directory structure and content. One can decide the name of stored files and directories or let the service do that automatically.
- The storage volume is a collection of storage levels, each level corresponding to a physical mount point. The storage volume manages itself the structure of the physical data but provides more functionality than directory volume. This includes storage in several physical places at the same time, synchronization between these physical storages, etc…

CHANGE

STORAGE:VOLUME:CHANGE

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	No

Description

This command changes the volume parameters. The volume must exist.

Any parameter can be changed except the volume name. If a parameter is not provided, it's value is unchanged.

Permissions

VOLUME_ADMIN

Parameters

NAME Volume name.
The NAME parameter is mandatory.

WAIT Wait time in seconds for the volume to be free. For this command, the STORAGE service must wait for the volume to be free. This parameters allows to define the maximum wait time. If the volume is still in use after the wait time, the command fails.
The default is "0" (infinite wait).

DESCRIPTION Volume description. This is just a string that describes the volume. It's not mandatory.

PATH Volume path. This parameter has meaning only for a directory volume. It gives the physical mount point of the directory volume.

OWNER Volume owner. This information is reserved for future use.

GROUP Volume group. The group name must not contain space and is case independent. Please see the overview chapter for information about volume groups.

CREATE

STORAGE:VOLUME:CREATE

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	No

Description

This command creates a new volume in the STORAGE service volume list. When a new volume is created, it doesn't contain any storage levels.

Permissions

VOLUME_ADMIN

Parameters

NAME Volume name. A volume is uniquely identified by its name. The volume name is case independent and cannot contain the characters ' ' (space), '.', '/', '\', '-'. If these characters appears in the NAME parameters, the command automatically removes them.

The NAME parameter is mandatory. If a volume with the same name ever exist, the command fails.

TYPE Volume type. Can be "DIRECTORY" for a directory volume or "STORAGE" for a storage volume. The default is "STORAGE".

DESCRIPTION Volume description. This is just a string that describes the volume. It's not mandatory.

OWNER Volume owner. This information is reserved for future use.

ENABLE

STORAGE:VOLUME:ENABLE

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	No

Description

This command enables or disables a volume. When a volume is disabled, it’s not available for any read or write operations.

It’s not possible to disable a volume while it’s in a synchronization operation. At this time, the command fails.

Permissions

VOLUME_ADMIN

Parameters

NAME	Volume name. The NAME parameter is mandatory.
ENABLE	This parameter must be set to “YES” for enabling the volume and to “NO” for disabling it. The default is “YES” (enable).

LIST

STORAGE:VOLUME:LIST

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	Yes

Description

This command returns the list of STORAGE service volumes.

The command creates a table object returning complete volume information:

- Name.
- Description.
- Type.
- Owner.
- Path.
- Number of sessions actually using the volume.
- Available state.
- Number of volume storage levels.
- Percentage of synchronization if the volume is in synchronization process.
- Status of the last synchronization process.
- Group name

Permissions

VOLUME_ADMIN

Parameters

None.

Objects created

VOLUME_LIST

This is a Nirva table object containing following columns:

- “NAME” is the volume name.
- “DESCRIPTION” is the volume description.
- “TYPE” is the volume type. It can be “DIRECTORY” or “STORAGE”.
- “OWNER” is the volume owner (not use for now).
- “PATH” is the volume physical mount point for a directory volume. This information is empty for a storage volume.
- “IN_USE” is the number of sessions actually using the volume.
- “AVAILABLE” is set to “YES” if the volume is enabled and to “NO” if the volume is disabled..
- “LEVELS” gives the number of storage levels defined for this volume.
- “SYNCHRO” is the percentage of synchronization if the volume is in synchronization process.
- “LAST_SYNCHRO_STATUS” is the status of the last started synchronization process. It can take values “STARTED”, “SUCCESS”, “FAILED” or “STOPPED”.

- “GROUP” is the optional volume group name.

SYNCHRO

STORAGE:VOLUME:SYNCHRO

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	No

Description

This command is available only for a volume of storage type.

This command starts the synchronization process from one storage level to another storage level.

The command can wait for the synchronization or can just starts it without waiting. At this time, the synchronization process occurs in a background task and the volume is still available for reading and writing operations.

The synchronization process consists of copying all the documents from a source level to a destination level. Only the documents that are not found on the destination level are copied. The controlling of an existing document is made only on the file name and not on the file content. The synchronization always occurs for a date range.

Before starting a synchronization process, the command always checks the source and destination levels. The checking consists of verifying that the level mount points are on line and corresponds to the volume.

Permissions

VOLUME_ADMIN

Parameters

NAME

Volume name.
The NAME parameter is mandatory.

WAIT

Wait time in seconds for the volume to be free. For synchronizing a volume, the STORAGE service must wait for the volume to be free. This parameters allows to define the maximum wait time. If the volume is still in use after the wait time, the command fails.
The default is “0” (infinite wait).

SLNAME

Source storage level name. The command copies the documents from the source level to the destination level. This must correspond to an existing storage level for the volume. If this parameter is not provided or is blank, the service uses the first storage level.

DLNAME Destination storage level name. The command copies the documents from the source level to the destination level. This must correspond to an existing storage level for the volume. If this parameter is not provided or is blank, the service uses the second storage level. If the destination level is the same than the source level, the command returns immediately without error.

ERASE Erase option. If this parameter is set to “YES”, the command removes the source documents if they are on the destination level after the copy. This allows to move some production data from one level to another level. The default is “NO”.
If during remove operation, a subdirectory of a level is empty, this directory is automatically removed by the command.

WAIT_RESULT Wait for result option. If this parameter is set to “YES”, the command immediately starts the synchronization and waits for its completion. Otherwise, the synchronization is launched in a background task and the command returns immediately. Then, the VOLUME:LIST command allows to see the progression and status of the synchronization process.
The default is “NO” (synchronization in background).

DATE Starting date. The synchronization process occurs always for a given date range. The DATE parameter is the starting date. It must be on the format dd/mm/yy or dd/mm/yyyy. If the DATE parameter is not provided, the service uses the date of the previous day.

TO_DATE Ending date. The synchronization process occurs always for a given date range. The TO_DATE parameter is the ending date. It must be on the format dd/mm/yy or dd/mm/yyyy. If the TO_DATE parameter is not provided, the service uses the same value than the DATE parameter.

REPOS Repository name. This parameter allows to restrict the synchronization process to a given repository. If it’s not provided, all the repositories are synchronized.

REMOVE

STORAGE:VOLUME:REMOVE

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	No

Description

This command removes an existing volume from the STORAGE service volume list. In any case, the command doesn't remove any data physically stored on the volume. This command may fail if the volume is in use (see the WAIT parameter).

Permissions

VOLUME_ADMIN

Parameters

NAME Volume name.
The NAME parameter is mandatory.

WAIT Wait time in seconds for the volume to be free. For removing a volume, the STORAGE service must wait for the volume to be free. This parameters allows to define the maximum wait time. If the volume is still in use after the wait time, the command fails.
The default is "0" (no wait time).

STOP_SYNCHRO

STORAGE:VOLUME:STOP_SYNCHRO

Source	Use Input Container	Use Output Container
Client Web Procedure Service	No	No

Description

This command is available only for a volume of storage type.

This command stops and eventual synchronization process that has been started in background by the VOLUME:SYNCHRONIZE command.

The synchronization is stopped where it is but always after a complete document copy.

There is no rollback procedure to come back to the state before the synchronization so all ever copied documents stay copied (or moved) when the stop occurs.

Permissions

VOLUME_ADMIN

Parameters

NAME

Volume name.
The NAME parameter is mandatory.