



# IBM Storwize FileSet Tasks

---

This chapter contains the following sections:

- [Delete IBM Storwize File Set, on page 2](#)
- [New IBM Storwize File Set, on page 3](#)

# Delete IBM Storwize File Set

## Summary

Delete IBM Storwize File Set

## Description

This task deletes the selected file set. If the file set is linked it will be unlinked. If the file set contains any linked child file sets they are also unlinked. The task fails if the file set contains snapshots, the snapshots must be deleted prior deleting the file set. The root file set cannot be deleted.

## Inputs

Input	Description	Mappable To Type	Mandatory
File Set	Choose an IBM Storwize file set to be deleted	ibmStorwizeFileSet Identity	Y
Deleting the file set also deletes all data that the file set contains Do you want to continue?			
* To confirm, type YES	Enter YES for recursive delete		Y

## Outputs

Output	Description	Type
IBM_STORWIZE_OUTPUT_FILESET_IDENTITY	Output of file set identity	ibmStorwizeFileSet Identity
IBM_STORWIZE_OUTPUT_ACCOUNT_NAME	Output of infra account name	accountName
IBM_STORWIZE_OUTPUT_FILESET_NAME	Output of file set name	gen_text_input
IBM_STORWIZE_OUTPUT_DATACENTER	Output of datacenter name on which selected operation was performed	datacenterName

# New IBM Storwize File Set

## Summary

New IBM Storwize File Set

## Description

This task creates a new file set using the specified name. The file sets root can be anywhere in the directory structure of the parent file system. The sets include all files and directories above that junction point. When creating a file set, the base directory path must exist. However, the directory (or junction point) that is being defined must not exist already, because it is to be created as part of the file set creation process. You must also define the file set as either dependent or independent. A dependent file set shares the same file system and inode definitions as the parent independent file set that contains it. If set to independent, the file set has its own inode space. That allows for independent management, such as quotas. When a file system is created, an initial file set is also created automatically in the root directory.

## Inputs

Input	Description	Mappable To Type	Mandatory
Select Account Name	Choose an IBM Storwize account	accountName	Y
File Set Type	Choose a file set type	ibmStorwizeFileSet Type	Y
Junction Path	Choose the path. The junction path must be on one of the file systems and not refer to any existing file or directory. Example of a valid path: file_system_path/existing_sub_directory/new_sub_directory. Example of an invalid path: file_system_path/existing_sub_directory/new_sub_directory1/new_sub_directory2. The linking process cannot create more than one new subdirectory.	ibmStorwizeFileSet JunctionPath Identity	Y
Subdirectory	Enter the subdirectory to which the file set belongs	gen_text_input	Y
Name	Enter the name of the file set	gen_text_input	Y
Owner	Enter the owner of the path. The owner can be a user name or a combination of domain and user name. For example: admin1 or domain1\admin1. Owner can be specified as 'root'. Changes are possible only while the file system is empty.	gen_text_input	Y
Group	Enter an associated group for the path. Explicitly configure access the control list if read or write access is required for the group.	gen_text_input	
Comments	Specify the comment that appears in the output of the "lsfset" command.	gen_text_input	
Type	Choose a file set type. Independent file sets allow for snapshots and group quotas. Dependent file sets support quotas only at the level of file system that contains the file set. There are other differences	ibmStorwizeFileSet SubType	Y

Input	Description	Mappable To Type	Mandatory
	between the two types for inode space, WAN caching, deletion, and linking.		
Independent parent file set that dependent file set belongs to		gen_text_input	
Maximum number of inodes	Enter a maximum number of inodes that the file set can use. This value depends on the available inodes in the file system. The minimum number of inodes that must be available to create a file set is 1024. You cannot reduce the maximum number of inodes to less than the number of allocated inodes. The system rounds up the value that you enter to fit into the blocks needed to accommodate the request. As a result, the maximum number of inodes can be higher than value that you entered.	gen_text_input	
Allocated number of inodes	Enter a number of inodes to be allocated when the file set is created. Allocating inodes increases the time it takes to add file sets to system; however, all subsequent writes are faster. The number of maximum allowed inodes cannot be decreased below the allocated number. To allocate all inodes for a files set, specify 0. The system rounds up to fit into the blocks needed to accommodate the request. As a result, the maximum number of inodes can be higher than the value that you entered.	gen_text_input	
Set a quota	Check to enable quota creation	gen_text_input	
Soft limit	Enter the soft quota limit. At the soft quota limit, a grace period starts. Data can be written until the grace period expires, or until hard quota limit is reached.	gen_text_input	
		ibmStorwizeQuota SoftHardLimitUnits	
Hard limit	Enter the hard limit value. At the hard quota limit, additional data cannot be stored until files are removed. Root user allocation has no quota limits.	gen_text_input	
		ibmStorwizeQuota SoftHardLimitUnits	
Snapshot	Choose one or more snapshots to associate with the file set	ibmStorageFileSet SnapshotRule Identity	

**Outputs**

<b>Output</b>	<b>Description</b>	<b>Type</b>
IBM_STORWIZE_OUTPUT_FILESET_IDENTITY	Output of file set identity	ibmStorwizeFileSet Identity
IBM_STORWIZE_OUTPUT_ACCOUNT_NAME	Output of infra account name	accountName
IBM_STORWIZE_OUTPUT_FILESET_NAME	Output of file set name	gen_text_input
IBM_STORWIZE_OUTPUT_DATACENTER	Output of datacenter name on which selected operation was performed	datacenterName

