IBM Storwize FileSystems Tasks

This chapter contains the following sections:

- Add Pool to IBM Storwize File System, on page 2
- Create IBM Storwize Compressed File System, on page 3
- Create IBM Storwize File System, on page 5
- Create IBM Storwize Migration-ILM File System, on page 6
- Delete IBM Storwize File System, on page 8
- Edit IBM Storwize File System, on page 9
- Mount IBM Storwize File System, on page 10
- Remove Pool from IBM Storwize File System, on page 11
- Replicate IBM Storwize File System, on page 12
- Unmount IBM Storwize File System, on page 13

Add Pool to IBM Storwize File System

Summary

Add a pool to an IBM Storwize file system.

Description

This task adds storage by adding a pool to a file system.

Inputs

Input	Description	Mappable To Type	Mandatory
File System Name	Choose an IBM file system to which to add the pool	ibmStorwizeFileSystem Identity	Y
Name	Enter the name of the file system pool to create. The pool name must be unique to the file system.	gen_text_input	Y
Storage Pool	Choose a storage pool	ibmStorwizePoolIdentity	Y
Size(GB)	Enter the storage pool size in GB	gen_text_input	Y

Outputs

Output	Description	Type
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_IDENTITY	Output of File System identity	ibmStorwizeFileSystem Identity
IBM_STORWIZE_OUTPUT_ ACCOUNT_NAME	Output of infra account name	accountName
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_NAME	Output of File System Name	gen_text_input
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_POOL_NAME	Output of File System Pool Name	gen_text_input
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_SIZE	Output of File System Size	gen_text_input
IBM_STORWIZE_OUTPUT_ DATACENTER	Output of datacenter name on which selected operation was performed	datacenterName

Create IBM Storwize Compressed File System

Summary
Description
Inputs

Input	Description	Mappable To Type	Mandatory
Account Name	Select account name	accountName	Y
File System Name	The name of the file system to be created	gen_text_input	Y
Owner	Sets the owner of the path. It can be a user name, or a combination of domain and user name. for example, admin1 or domain1\admin1. You can specify 'root' as an owner. Changes are possible only while the filesystem is empty.	gen_text_input	Y
Owning group	Sets the owner of the path. It can be a user name, or a combination of domain and user name. for example, admin1 or domain1\admin1. You can specify 'root' as an owner	gen_text_input	Y
File system pool settings	File system pool settings		
System(non compressed)	Select storage pool	ibmStorwizePool Identity	Y
Size(GB)	Enter storage pool size	gen_text_input	Y
System(metadata)	Select storage pool	ibmStorwizePool Identity	Y
Size(GB)	Enter storage pool size	gen_text_input	Y
In order to enable metadata replication Please select a pool from System(metadata)			
System(compressed)	Select storage pool	ibmStorwizePool Identity	Y
Size(GB)	Enter storage pool size	gen_text_input	Y

Outputs

Output	Description	Type
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_IDENTITY	Output of File System identity	ibmStorwizeFileSystem Identity
IBM_STORWIZE_OUTPUT_ ACCOUNT_NAME	Output of infra account name	accountName

Output	Description	Type
IBM_STORWIZE_OUTPUT_ DATACENTER	Output of datacenter name on which selected operation was performed	datacenterName
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_NAME	Output of File System Name	gen_text_input
IBM_STORWIZE_OUTPUT_ FILESET_JUNCTIONPATH_ IDENTITY	Output of file set junction path identity	ibmStorwizeFileSet JunctionPathIdentity

Create IBM Storwize File System

Summary

Create IBM Storwize File System

Description

This task creates a cluster file system. A file system is the physical structure that an operating system uses to store and organize files on a storage system. The block size and replication affects file system performance. The minimum supported physical capacity (size on disk) for a filesystem is 100GB. This provides the base file system and global namespace from which shares and exports can be created. Therefore, each file system comprises a structure of directories or paths that are based on a single root. Any of these directories can be used to create a share or export.

Inputs

Input	Description	Mappable To Type	Mandatory
Account Name	Select Account name	accountName	Y
File System Name	The name of the file system to be created.	gen_text_input	Y
Owner	Sets the owner of the path. It can be a user name, or a combination of domain and user name. for example, admin1 or domain1\admin1. You can specify 'root' as an owner. Changes are possible only while the filesystem is empty.	gen_text_input	Y
File System Pool Name	Enter file system pool name		Y
Storage Pool	Select storage pool	ibmStorwizePool Identity	Y
File System Size	Enter file system size	gen_text_input	Y
File System Size Units	File System Size Units		

Outputs

Output	Description	Type
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_IDENTITY	Output of File System identity	ibmStorwizeFileSystem Identity
IBM_STORWIZE_OUTPUT_ ACCOUNT_NAME	Output of infra account name	accountName
IBM_STORWIZE_OUTPUT_ DATACENTER	Output of datacenter name on which selected operation was performed	datacenterName
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_NAME	Output of File System Name	gen_text_input
IBM_STORWIZE_OUTPUT_ FILESET_JUNCTIONPATH_ IDENTITY	Output of file set junction path identity	ibmStorwizeFileSet JunctionPathIdentity

Create IBM Storwize Migration-ILM File System

Summary
Description
Inputs

Input	Description	Mappable To Type	Mandatory
Account Name	Select account name	accountName	Y
File System Name	The name of the file system to be created	gen_text_input	Y
Owner	Sets the owner of the path. It can be a user name, or a combination of domain and user name. for example, admin1 or domain1\admin1. You can specify 'root' as an owner	gen_text_input	Y
Owning group	Name or GID of group that is associated with the path	gen_text_input	Y
File system pool settings	File system pool settings		
System	Select storage pool	ibmStorwizePoolIdentity	Y
Size(GB)	Enter storage pool size	gen_text_input	Y
System(metadata)	Select storage pool	ibmStorwizePoolIdentity	
Size(GB)	Enter storage pool size	gen_text_input	
In order to enable metadata replication Please select a pool from System(metadata)	help		
Name	The name of the file system pool to create. The pool name must be unique per file system.	gen_text_input	Y
Compressed	Select compression for the file system pool .You cannot change compression settings after a file system pool has been created.	IBMStorwizeBooleanList	Y
Storage Pool	Select storage pool	ibmStorwizePoolIdentity	Y
Size(GB)	Enter storage pool size	gen_text_input	Y
Migration start threshold (%)	Migration will start automatically when the used capacity of file system exceeds this value	gen_text_input	Y

Input	Description	Mappable To Type	Mandatory
Migration stop threshold (%)	Migration will stop automatically when the used capacity of file system decreases below this value	gen_text_input	Y
Enable Migration Schedule	In addition to the migration being started automatically when the start threshold is reached .you may also choose to schedule the migration policy to run at a specific time.	IBMStorwizeBooleanList	
Type	Select the type	ibmStorwizeMigration ILMFileSystemSchedule TypeList	Y
Time Of Day	Select the time of day	ibmStorwizeMigration ILMFileSystemSchedule TimeList	Y
Day of the Week	Select the day of the week	ibmStorwizeMigration ILMFileSystemDayOf WeekTimeList	Y
Days of the week	Select the days of the week	IBMSTORAGEDaysOf Week	
Day of Month	Select the day of month	ibmStorwizeMigration ILMFileSystemDayOf MonthTimeList	Y
Days of Month	Select the days of month	IBMSTORAGEDaysOf Month	

Outputs

Output	Description	Type
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_IDENTITY	Output of File System identity	ibmStorwizeFileSystem Identity
IBM_STORWIZE_OUTPUT_ ACCOUNT_NAME	Output of infra account name	accountName
IBM_STORWIZE_OUTPUT_ DATACENTER	Output of datacenter name on which selected operation was performed	datacenterName
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_NAME	Output of File System Name	gen_text_input
IBM_STORWIZE_OUTPUT_ FILESET_JUNCTIONPATH_ IDENTITY	Output of file set junction path identity	ibmStorwizeFileSet JunctionPathIdentity

Delete IBM Storwize File System

Summary

Delete IBM Storwize File System

Description

This task deletes all the disks that belong to this file system, along with the data on it. Before deleting a file system, you must unmount it from all nodes. The file system cannot be deleted until all the exports are deleted.

Inputs

Input	Description	Mappable To Type	Mandatory
FileSystem	Choose a file system to delete	ibmStorwizeFileSystemIdentity	Y

Outputs

Output	Description	Type
IBM_STORWIZE_OUTPUT_FILE_SYSTEM_IDENTITY	Output of File System identity	ibmStorwizeFileSystemIdentity
IBM_STORWIZE_OUTPUT_ACCOUNT_NAME	Output of infra account name	accountName
IBM_STORWIZE_OUTPUT_DATACENTER	Output of datacenter name on which selected operation was performed	datacenterName
IBM_STORWIZE_OUTPUT_FILE_SYSTEM_NAME	Output of File System Name	gen_text_input

Edit IBM Storwize File System

Summary

Edit an IBM Storwize file system.

Description

This task edits a file system. The size of the file system can be modified by adding or removing disks associated with it. Reducing the size of the file system can cause loss of data.

Inputs

Input	Description	Mappable To Type	Mandatory
FileSystem	Choose a file system to edit	ibmStorwizeFileSystem Identity	Y
File System Pool	Choose a file system pool to resize	gen_text_input	Y
File System Size(GB)	Enter a file system pool size in GB. The minimum supported physical capacity (size on disk) for a file system is 100GB.	gen_text_input	Y

Outputs

Output	Description	Type
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_IDENTITY	Output of File System identity	ibmStorwizeFileSystem Identity
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_NAME	Output of File System Name	gen_text_input
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_POOL_NAME	Output of File System Pool Name	gen_text_input
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_SIZE	Output of File System Size	gen_text_input
IBM_STORWIZE_OUTPUT_ ACCOUNT_NAME	Output of infra account name	accountName
IBM_STORWIZE_OUTPUT_ DATACENTER	Output of datacenter name on which selected operation was performed	datacenterName

Mount IBM Storwize File System

Summary

Mount IBM Storwize File System

Description

This task mounts the specified file system on all interface and management nodes or a specified subset of nodes. If nodes are specified, those nodes must be suspended. If a cluster is not denoted, the default cluster is used. Any exports deactivated due to previous unmounts will be reactivated once all non-suspended nodes have mounted the file system again.

Inputs

Input	Description	Mappable To Type	Mandatory
File System	Choose a file system to mount	ibmStorwizeFileSystemIdentity	Y

Outputs

Output	Description	Type
IBM_STORWIZE_OUTPUT_FILE_SYSTEM_IDENTITY	Output of File System identity	ibmStorwizeFileSystemIdentity
IBM_STORWIZE_OUTPUT_ACCOUNT_NAME	Output of infra account name	accountName
IBM_STORWIZE_OUTPUT_DATACENTER	Output of datacenter name on which selected operation was performed	datacenterName
IBM_STORWIZE_OUTPUT_FILE_SYSTEM_NAME	Output of File System Name	gen_text_input

Remove Pool from IBM Storwize File System

Summary

Remove a pool from an IBM Storwize file system.

Description

This task removes a pool from a file system. Disk removal is possible down to the supported size of File system.

Inputs

Input	Description	Mappable To Type	Mandatory
FileSystem	Choose an IBM file system from which to remove the pool	ibmStorwizeFileSystem Identity	Y
File System Pool	Choose a file system pool to remove	gen_text_input	Y

Outputs

Output	Description	Type
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_IDENTITY	Output of File System identity	ibmStorwizeFileSystem Identity
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_NAME	Output of File System Name	gen_text_input
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_POOL_NAME	Output of File System Pool Name	gen_text_input
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_SIZE	Output of File System Size	gen_text_input
IBM_STORWIZE_OUTPUT_ ACCOUNT_NAME	Output of infra account name	accountName
IBM_STORWIZE_OUTPUT_ DATACENTER	Output of datacenter name on which selected operation was performed	datacenterName

Replicate IBM Storwize File System

Summary

Replicate IBM Storwize File System

Description

This task declares a source cluster and a local path for replication.

Inputs

Input	Description	Mappable To Type	Mandatory
FileSystem	Choose an IBM file system to replicate	ibmStorwizeFileSystem Identity	Y
Source cluster ID	Enter the ID value of the source cluster that uses the local path for replication	gen_text_input	Y

Outputs

Output	Description	Type
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_IDENTITY	Output of File System identity	ibmStorwizeFileSystem Identity
IBM_STORWIZE_OUTPUT_ ACCOUNT_NAME	Output of infra account name	accountName
IBM_STORWIZE_OUTPUT_ DATACENTER	Output of datacenter name on which selected operation was performed	datacenterName
IBM_STORWIZE_OUTPUT_FILE_ SYSTEM_NAME	Output of File System Name	gen_text_input

Unmount IBM Storwize File System

Summary

Unmount IBM Storwize File System

Description

This task unmounts the specified file system on all nodes or a specified set of nodes. If nodes are specified, those nodes must be suspended. Suspended nodes may need to be restarted with the stopcluster command in order to enable the file system to fully unmount. If unmounting on non-suspended nodes, all active exports must be deactivated before unmounting.

Inputs

Input	Description	Mappable To Type	Mandatory
FileSystem	Choose a file system to unmount	ibmStorwizeFileSystemIdentity	Y

Outputs

Output	Description	Type
IBM_STORWIZE_OUTPUT_FILE_SYSTEM_IDENTITY	Output of File System identity	ibmStorwizeFileSystemIdentity
IBM_STORWIZE_OUTPUT_ACCOUNT_NAME	Output of infra account name	accountName
IBM_STORWIZE_OUTPUT_DATACENTER	Output of datacenter name on which selected operation was performed	datacenterName
IBM_STORWIZE_OUTPUT_FILE_SYSTEM_NAME	Output of File System Name	gen_text_input