

# hadoop fs

The `hadoop fs` command runs a generic filesystem user client that interacts with the MapR filesystem (MapR-FS).



On the Windows client, make sure that the `PATH` contains the following directories:

- `C:\Windows\system32`
- `C:\Windows`

If they are not present, the `hadoop fs` command might fail silently.

## Syntax

```
hadoop [ Generic Options ] fs
    [-cat <src>]
    [-chgrp [-R] GROUP PATH...]
    [-chmod [-R] <MODE[,MODE]... | OCTALMODE> PATH...]
    [-chown [-R] [OWNER][:[GROUP]] PATH...]
    [-copyFromLocal <localsrc> ... <dst>]
    [-copyToLocal [-ignoreCrc] [-crc] <src> <localdst>]
    [-count[-q] <path>]
    [-cp <src> <dst>]
    [-df <path>]
    [-du <path>]
    [-dus <path>]
    [-expunge]
    [-get [-ignoreCrc] [-crc] <src> <localdst>]
    [-getmerge <src> <localdst> [addnl]]
    [-help [cmd]]
    [-ls <path>]
    [-lsr <path>]
    [-mkdir <path>]
    [-moveFromLocal <localsrc> ... <dst>]
    [-moveToLocal <src> <localdst>]
    [-mv <src> <dst>]
    [-put <localsrc> ... <dst>]
    [-rm [-skipTrash] <src>]
    [-rmr [-skipTrash] <src>]
    [-stat [format] <path>]
    [-tail [-f] <path>]
    [-test [-ezd] <path>]
    [-text <path>]
    [-touchz <path>]
```

## Parameters

### Command Options

The following command parameters are supported for `hadoop fs`:

Parameter	Description
<code>-cat &lt;src&gt;</code>	Fetch all files that match the file pattern defined by the <code>&lt;src&gt;</code> parameter and display their contents on <i>stdout</i> .

<code>-fs [local   &lt;file system URI&gt;]</code>	<p>Specify the file system to use.</p> <p>If not specified, the current configuration is used, taken from the following, in increasing precedence:  <code>core-default.xml</code> inside the hadoop jar file  <code>core-site.xml</code> in <code>\$HADOOP_CONF_DIR</code></p> <p>The <code>local</code> option means use the local file system as your DFS.</p> <p><code>&lt;file system URI&gt;</code> specifies a particular file system to contact. This argument is optional but if used must appear first on the command line. Exactly one additional argument must be specified.</p>
<code>-ls &lt;path&gt;</code>	<p>List the contents that match the specified file pattern. If path is not specified, the contents of <code>/user/&lt;currentUser&gt;</code> will be listed.</p> <p>Directory entries are of the form <code>dirName (full path) &lt;dir&gt;</code> and file entries are of the form <code>fileName(full path) &lt;rn&gt;.size</code> where <code>n</code> is the number of replicas specified for the file and <code>size</code> is the size of the file, in bytes.</p>
<code>-lsr &lt;path&gt;</code>	<p>Recursively list the contents that match the specified file pattern. Behaves very similarly to <code>hadoop fs -ls</code>, except that the data is shown for all the entries in the subtree.</p>
<code>-df [&lt;path&gt;]</code>	<p>Shows the capacity, free and used space of the filesystem.</p> <p>If the filesystem has multiple partitions, and no path to a particular partition is specified, then the status of the root partitions will be shown.</p>
<code>-du &lt;path&gt;</code>	<p>Show the amount of space, in bytes, used by the files that match the specified file pattern. Equivalent to the Unix command <code>du -sb &lt;path&gt;/*</code> in case of a directory, and to <code>du -b &lt;path&gt;</code> in case of a file.</p> <p>The output is in the form <code>name (full path) size (in bytes)</code>.</p>
<code>-dus &lt;path&gt;</code>	<p>Show the amount of space, in bytes, used by the files that match the specified file pattern. Equivalent to the Unix command <code>du -sb</code>. The output is in the form <code>name(full path) size (in bytes)</code>.</p>
<code>-mv &lt;src&gt; &lt;dst&gt;</code>	<p>Move files that match the specified file pattern <code>&lt;src&gt;</code> to a destination <code>&lt;dst&gt;</code>. When moving multiple files, the destination must be a directory.</p>
<code>-cp &lt;src&gt; &lt;dst&gt;</code>	<p>Copy files that match the file pattern <code>&lt;src&gt;</code> to a destination. When copying multiple files, the destination must be a directory.</p>
<code>-rm [-skipTrash] &lt;src&gt;</code>	<p>Delete all files that match the specified file pattern. Equivalent to the Unix command <code>rm &lt;src&gt;</code>.</p> <p>If enabled, the <code>-skipTrash</code> option bypasses trash and immediately deletes <code>&lt;src&gt;</code>.</p>
<code>-rmr [-skipTrash] &lt;src&gt;</code>	<p>Remove all directories that match the specified file pattern. Equivalent to the Unix command <code>rm -rf &lt;src&gt;</code>.</p> <p>If enabled, the <code>-skipTrash</code> option bypasses trash and immediately deletes <code>&lt;src&gt;</code>.</p>
<code>-put &lt;localsrc&gt; ... &lt;dst&gt;</code>	<p>Copy files from the local file system into fs.</p>
<code>-copyFromLocal &lt;localsrc&gt; ... &lt;dst&gt;</code>	<p>Identical to the <code>-put</code> command.</p>
<code>-moveFromLocal &lt;localsrc&gt; ... &lt;dst&gt;</code>	<p>Same as <code>-put</code>, except that the source is deleted after it's copied.</p>
<code>-get [-ignoreCrc] [-crc] &lt;src&gt; &lt;localdst&gt;</code>	<p>Copy files that match the file pattern <code>&lt;src&gt;</code> to the local name. <code>&lt;src&gt;</code> is kept.</p> <p>When copying multiple files, the destination must be a directory.</p>

<code>-getmerge &lt;src&gt; &lt;localdst&gt;</code>	Get all the files in the directories that match the source file pattern and merge and sort them to only one file on local fs. <code>&lt;src&gt;</code> is kept.
<code>-copyToLocal [-ignoreCrc] [-crc] &lt;src&gt; &lt;localdst&gt;</code>	Identical to the <code>-get</code> command.
<code>-moveToLocal &lt;src&gt; &lt;localdst&gt;</code>	Not implemented yet.
<code>-mkdir &lt;path&gt;</code>	Create a directory in specified location.
<code>-tail [-f] &lt;file&gt;</code>	Show the last 1KB of the file. The <code>-f</code> option shows appended data as the file grows.
<code>-touchz &lt;path&gt;</code>	Write a timestamp in <code>yyyy-MM-dd HH:mm:ss</code> format in a file at <code>&lt;path&gt;</code> . An error is returned if the file exists with non-zero length.
<code>-test [-ezd] &lt;path&gt;</code>	If file { exists, has zero length, is a directory then return 0, else return 1.
<code>-text &lt;src&gt;</code>	Takes a source file and outputs the file in text format. The allowed formats are <code>zip</code> and <code>TextRecordInputStream</code> .
<code>-stat [format] &lt;path&gt;</code>	Print statistics about the file/directory at <code>&lt;path&gt;</code> in the specified format.  Format accepts filesize in blocks ( <code>%b</code> ), filename ( <code>%n</code> ), block size ( <code>%o</code> ), replication ( <code>%r</code> ), and modification date ( <code>%y</code> , <code>%Y</code> ).
<code>-chmod [-R] &lt;MODE[,MODE]...   OCTALMODE&gt; PATH...</code>	Changes permissions of a file. This works similar to shell's <code>chmod</code> with a few exceptions.  <code>-R</code> modifies the files recursively. This is the only option currently supported.  <code>MODE</code> is same as the mode used for <code>chmod</code> shell command. The only letters recognized are <code>rwXt</code> . That is, <code>+t, a+r, g-w, +rwx, o=r</code>  <code>OCTALMODE</code> is specified in 3 or 4 digits. If 4 digits, the first may be 1 or 0 to turn the sticky bit on or off, respectively. Unlike shell command, it is not possible to specify only part of the mode For example, <code>754</code> is same as <code>u=rwx, g=rx, o=r</code>  If none of 'augo' is specified, 'a' is assumed and unlike shell command, no <code>umask</code> is applied.
<code>-chown [-R] [OWNER][:[GROUP]] PATH...</code>	Changes owner and group of a file. This is similar to shell's <code>chown</code> with a few exceptions.  <code>-R</code> modifies the files recursively. This is the only option currently supported.  If only owner or group is specified then only owner or group is modified.  The owner and group names may only consists of digits, alphabet, and any of <code>-.@/'</code> i.e. <code>[-./a-zA-Z0-9]</code> .  The names are case-sensitive.
	 <b>Warning</b> <b>WARNING:</b>  Avoid using <code>'</code> to separate user name and group though Linux allows it. If user names have dots in them and you are using local file system, you might see surprising results since shell command <code>chown</code> is used for local files.
<code>-chgrp [-R] GROUP PATH...</code>	This is equivalent to <code>-chown ... :GROUP ...</code>
<code>-count[-q] &lt;path&gt;</code>	Count the number of directories, files and bytes under the paths that match the specified file pattern. The output columns are: <ul style="list-style-type: none"> <li>• <code>DIR_COUNT FILE_COUNT CONTENT_SIZE FILE_NAME</code> or</li> <li>• <code>QUOTA REMAINING_QUOTA SPACE_QUOTA REMAINING_SPACE_QUOTA</code></li> <li>• <code>DIR_COUNT FILE_COUNT CONTENT_SIZE FILE_NAME</code></li> </ul>

-help [cmd]

Displays help for given command or all commands if none is specified.

## Generic Options

The following generic options are supported for the `hadoop fs` command: `-conf <configuration file>`, `-D <property=value>`, `-fs <local|file system URI>`, `-jt <local|jobtracker:port>`, `-files <file1,file2,file3,...>`, `-libjars <libjar1,libjar2,libjar3,...>`, and `-archives <archive1,archive2,archive3,...>`.

For more information on generic options, see [Generic Options](#).